Body-Worn Cameras in Policing: Benefits and Costs

Morgan C. Williams, Jr.*
Robert. F. Wagner Graduate School of Public Service, NYU

Nathan Weil
University of Chicago Crime Lab

Elizabeth Rasich
University of Chicago Crime Lab

Jens Ludwig
Harris School of Public Policy, University of Chicago

Hye Chang
University of Chicago Crime Lab

Sophia Egrari
University of Chicago Crime Lab

March 25, 2021

Abstract

Body-worn cameras (BWCs) are an increasingly common tool for police oversight, accountability, and transparency, yet there remains uncertainty about their impacts on policing outcomes. This paper reviews what we know about the benefits of BWCs and how those benefits compare to the costs of this new technology. We make two contributions relative to existing research. First, we update prior meta-analyses of studies of the impacts of BWCs on policing outcomes to incorporate the most recent, and largest, studies carried out to date in this literature. This additional information provides additional support for the idea that cameras may affect a number of policing outcomes that are important from a social welfare perspective, particularly police use of force. Second, we carry out a benefit-cost analysis of BWCs, as financial barriers are often cited as a key impediment to adoption by police departments. Our baseline estimate for the benefit-cost ratio of BWCs is 4.95. Perhaps as much as one-quarter of the estimated benefits accrue to government budgets directly, which suggests the possibility that this technology could, from the narrow perspective of government budgets, even pay for itself.

*Acknowledgements: This report was funded by a grant from Citadel to the Council on Criminal Justice and the University of Chicago Crime Lab. As part of the University of Chicago Crime Lab's ongoing work to build and implement an early warning system for the Chicago Police Department as part of the city's consent decree with the Illinois Attorney General, the Crime Lab received two $5,000 grants from Axon to support travel of subject matter experts to Chicago. For helpful comments and discussions on this report, we thank Nancy Lavigne and the CCJ Police Task Force as well as Anthony Braga, Aaron Chalfin, Philip Cook, Barry Friedman, and Emily Owens. All interpretations and any errors are of course our own.

1 Introduction

Body-worn cameras (BWCs) have become an increasingly common tool for police oversight, accountability, and transparency in the United States. The technology has the potential to help deter police misconduct by better monitoring officer behavior out in the field. On the other hand, there could be unintended consequences if, for example, by creating a formal video record of civilian infractions, officers respond by curtailing discretion and increasing formal enforcement actions.

Measuring the effects of BWCs on policing outcomes is complicated in practice. The commonly preferred approach of randomized controlled trials (RCTs) for BWCs usually involves assigning the new technology to some officers (or shifts) but not others within a department. However, with BWCs there could be 'spillover' in the effects of cameras within departments, for example, when an officer with a BWC shows up at the same scene as officers without cameras, or if a given officer spends some shifts with a camera and other shifts without. Spillover effects like this in general would be expected to lead an RCT to understate an intervention's impacts. In addition, the sample sizes of existing BWC RCTs, even with studies carried out in large cities, have in practice limited statistical power to detect impacts on the most important policing outcomes such as police use of force. Population-level data on large numbers of jurisdictions, and what happens when they adopt BWCs at different points in time, could help overcome both of these challenges. But the tradeoff is the use of quasi-experimental research designs that may be more vulnerable to questions about internal validity.

This paper seeks to make two contributions to the research on this new technology. First, we seek to update previous reviews of the available literature on the effects of BWCs on policing outcomes to incorporate newer studies in this literature, which also tend to be among the largest studies that have been carried out to date. Second, we examine how these benefits compare to the costs of BWCs since financial barriers are often cited by police departments as a barrier to adoption. For example, as part of the debate about whether to adopt BWCs in Portland, the mayor's office declared, "the budget is tight, but the Mayor intends to revisit the use of body cameras when financial conditions permit" (Iboshi, 2021).

The best review to date of the BWC literature is a meta-analysis carried out for the Campbell Collaboration by Lum et al. (2020), which offers a circumspect conclusion: "Overall, there remains substantial uncertainty about whether BWCs can reduce officer use of force, but the variation in effects suggests there may be conditions in which BWCs could be effective. BWCs also do not seem to affect other police and citizen behaviors in a consistent manner, including officers' self-initiated activities or arrest behaviors, dispatched calls for service, or assaults and resistance against police officers. BWCs can reduce the number of citizen complaints against police officers, but it is unclear whether this finding signals an improvement in the quality of police-citizen interactions or a change in reporting" (p. 1).

We extend this review by now incorporating data from two recent BWC studies. The first of these is the largest RCT carried out to date, by Braga et al. (2020), which also assigns BWCs at the precinct-level within New York City, thereby mitigating spillover effects. The second study, by Kim (2020), applies a difference-in-differences research design to data from over two thousand police departments nationwide as well as a separate analysis of more granular data from the state of New Jersey to estimate the effects of BWC adoption on police use of force. The main consequence of incorporating these new studies into the Lum et al. (2020) meta-analytic dataset is to produce a larger and more precisely estimated reduction in complaints against police (-16.9%), and an estimated -9.6% reduction in police use of force, which is 40% larger in absolute value than the one reported in Lum et al. (2020) but not statistically significant at conventional levels (p=.152).

We next try to quantify the value to society from these benefits and how they compare to the costs of BWCs. We focus on the expected value of the benefits and costs of this technology, rather than the risk of a type I error, following the approach argued for in Cook and Ludwig (2006). We go beyond earlier benefit-cost analyses such as those found in PERF (2018) and Braga et al. (2017) by incorporating valuations of a broader set of policing outcomes, including willingness-to-pay estimates for reduced police use of force.

Our baseline estimate for the benefit-cost ratio of BWCs is 4.95. Potentially up to one-quarter of the benefits accrue to government budgets directly, versus to the rest of society, so it is possible that even from an overly narrow perspective of the consequences for government budgets directly, BWCs might pay for themselves. The main limitation of our analysis is that while police use of force is by far the most important outcome in terms of societal impact, as suggested by its dominance in the benefit-cost analysis, it is also an outcome for which estimation of impacts is challenging given the available data. Incorporating the large new study by Kim (2020) helps improve our understanding, but learning more about BWC effects on this outcome in particular would seem to be a top priority for additional research on police reform.

While we have focused on the *average* effects of BWCs across departments as they currently implement them, the actual impacts will surely vary depending on how these cameras are used. For policy purposes, we would ideally wish to measure how aspects of BWC adoption or other features of the local context moderate impacts on policing outcomes. Unfortunately, that winds up asking more of the available data than it can currently accommodate.

Among the set of police departments that have adopted BWCs, the data seem to suggest they can be helpful on average, although by themselves they are clearly not a panacea. As we consider the possible impacts as this technology scales further, it is possible that our estimates represent an upper bound on what the effects of BWCs could achieve in other jurisdictions if the early-adopting agencies willing to participate in RCTs are the ones that are most motivated to try to implement cameras well (what Allcott (2015) calls ‘site-selection bias’). On the other hand, with most new technologies, human beings cannot immediately grasp the full range of uses to which they might be applied (in 1943, for instance, Thomas Watson, the then-president of IBM, said ‘I think there is a world market for maybe five computers’). So it is possible that the benefits of cameras may grow over time as departments better understand their full potential. This is, in short, a policing reform technology that we currently believe generates benefits in excess of costs but that is also worth continued monitoring and evaluation.

2 Background and conceptual framework

In response to calls for increased police oversight, a growing number of law enforcement agencies throughout the country have integrated body-worn camera (BWC) technology into their policing practices. In 2013, 32% of local police departments employed some form of BWC technology with nearly a quarter of patrol officers receiving BWCs on average (Reaves, 2015).[1] By 2016, just three years later, that figure had already increased to 47% (Hyland, 2018). Many states now require varying degrees of BWC usage by local law enforcement agencies.

A primary motivation for BWC adoption is to help better monitor police behavior, and hence deter police misconduct (Ariel et al., 2015; Gaub et al., 2020a). A substantial amount of police work takes place in the field and away from immediate supervision or oversight. Citizen complaints against officers can make supervisors aware of problems, but the share of total police misconduct that results in a complaint is not well understood. Even when a complaint is filed, the investigation can wind up boiling down to the word of the officer against that of the complainant. Knowing this could make some civilians reluctant to file complaints in the first place, so body-worn cameras have the potential to increase the likelihood that any officer misconduct is appropriately detected, investigated, and sanctioned.[2]

Research from psychology on the effects of *power* on human judgment and decision-making suggests a different type of mechanism through which BWCs could change police behavior. For example, Keltner et al. (2003) shows that being in a position of power can lead to a number of behavioral changes including reduced inhibition, less attention devoted to other less-powerful people within the social environment, focusing more narrowly on achieving one's own goals and objectives, and believing that one's actions will not be met with 'interference or serious social consequences.' The prospect of having interactions recorded and subject to review by a third party, and possible sanction if there is misconduct, could potentially change an officer's perceived power dynamic when interacting with a civilian. That is, beyond the standard mechanism of deterrence and the officer's conscious understanding of increased sanctions in response to misconduct, BWCs could change officer behavior at a subconscious level as well.

BWCs could, in principle, also serve as a useful management or teaching tool for officers. Owens et al. (2018), for example, find that having supervisors sit down with an officer and review a recent police-civilian interaction can lead to subsequent changes in officer behavior. BWCs have the potential to provide new opportunities for that type of reflective exercise and may also enable supervisors to identify which officers to prioritize for similar interventions.

Because BWCs may record the behavior of *all* parties to a policing interaction, they have the potential to change the behavior of *civilians* in these encounters as well. Officers in departments that choose to adopt BWCs often wind up supporting this technology because they believe it helps deter civilian misbehavior and frivolous or retributive complaints.

As with any technology, there is also the risk of unintended consequences. For example, BWCs create a formal record of whatever officers encounter out in the field when the cameras are turned on, which will include civilian infractions. This could make some officers wary of unwanted scrutiny

[1] The deployment of BWC technology generally involves patrol officers who constitute one-third of sworn personnel and often experience the greatest level of civilian contact (Gaub et al., 2020b). Non-patrol specialized units (e.g., special investigations, canine, and special weapons and tactics (SWAT) units) are in many departments currently exempt from BWC requirements, with the rationale that these units participate in sensitive matters such as surveillance or dealing with confidential informants.

[2] Any reduction in police misconduct that might result would obviously be enormously important to society for its own sake, but in addition could have valuable secondary effects as well by, for instance, improving police-community trust in ways that facilitate improved investigation of unsolved shootings.

if they use their discretion to 'give someone a break.' This could potentially result in an increase in formal enforcement actions against civilians (e.g., citations or arrests).

BWCs also create the possibility of behavioral responses that may complicate the measurement and interpretation of their estimated 'effects' within the data (Lum et al., 2020). For example, an apparent increase in proactive policing activity in response to BWC adoption could reflect a change in proactive policing *or* an increase in how officers record such activity. Another example cited by Lum et al. (2020) is the possibility of a chilling effect of BWCs on the willingness of civilians to file complaints against the police. In principle, the opposite reporting effect on complaints is logically possible as well. Prior to BWCs, some civilians may have been reluctant to file complaints if they believed these would ultimately just come down to their word against the officer's. The presence of video documenting the interaction could thus encourage some people to file complaints who would not have otherwise.

Each of these potential effects is likely contingent on how a police department implements BWCs in practice. For example, if departments do not enforce rules that officers must turn on their cameras for every civilian interaction or if they fail to hold officers accountable for misconduct that is captured on video, this technological intervention could ultimately produce little change in policing outcomes. Departments that have adopted BWCs to date also vary enormously in other policies and practices such as how often or extensively footage is reviewed, what share of officers in a department is required to use cameras, whether a civilian should be notified that they are being recorded at the start of an interaction, and the extent to which BWC footage is made accessible to the general public. For policy purposes understanding how these specific practices moderate the effects of BWCs is of great importance. But as we discuss below, BWC studies to date have largely focused on measuring the *average* effect of BWC adoption on policing outcomes because the uncertainty intervals are so large we can often barely tell what the average effect itself is. Looking at how impacts vary across implementation practices asks even more of the data, not only by reducing the data available to estimate these types of sub-group effects, but also because the search through the space of possibly relevant sub-groups requires multiple-testing adjustments that further compromise statistical power. Examining how BWC impacts vary by implementation approach is a key priority for future research as the data on this technology expands over time.

3 Evidence on the Effects of BWC Technology

In this section, we review what is known about the impacts of BWCs on policing outcomes. We combine the meta-analytic dataset assembled by Lum et al. (2020) with two additional studies that came out after the publication of Lum et al. (2020), namely the large-scale RCT of BWCs in New York City reported in Braga et al. (2020) and the natural experiment evaluated in Kim (2020) based on both national officer-involved homicide data and use of force data for the state of New Jersey.[3]

[3]While there are a few additional studies that have come out since Lum et al. (2020) other than Braga et al. (2020) and Kim (2020), we do not include them in our review because they are either not directly on topic or do not meet our standards of evidence. We specifically exclude studies that rely on panel-data research designs but do not present evidence that the parallel trends assumption is satisfied as well as studies without sufficient information about the estimated sampling uncertainty behind their impact estimates (for example, by not presenting standard errors). For more information on the methodological criteria for our review, please see Feigenberg et al. (2021).

3.1 Prior studies

Meta-analysis pools together results from many studies to identify general tendencies in the available research. Idiosyncratic variation in results across studies due to statistical noise can be overcome through this pooling and averaging. The approach works particularly well in settings like medicine, where FDA requirements ensure that most of the studies being reviewed are well-executed RCTs and the key health outcomes of interest are often fairly standardized. The approach becomes more complicated in research literatures where there can be more heterogeneity across studies in their research designs, underlying quality, and measurement of key concepts of interest (outcomes and interventions). For example, measurement of police use of force can vary greatly from study to study; handcuffing suspects is considered a use of force in many studies, while others restrict use of force to injury-causing incidents or fatal incidents only. Different methods of measuring use of force can make a given department seem more or less extreme in its use of force practices and in at least some departments shows them to be at completely opposite ends of the spectrum (Geller et al., 2020).[4] In these cases, meta-analysis has the potential to obscure as much as it reveals.

In what follows, we aim to give readers a sense of the strengths and limitations of the studies in the BWC literature. We focus on the very best of these studies, including particularly those that use RCT designs. But even RCTs in this application have some important limitations. For example, consider the study by Ariel et al. (2015), widely viewed as the first rigorous evaluation of BWC adoption. Conducted within the police department in Rialto, California in 2012, their research design included 54 officers who on a shift-by-shift basis were randomized to either 'treatment' (BWC) or 'control' (no BWC) across a total of 988 shifts. One potential challenge of this design is the possibility that any behavioral changes induced in officers during BWC shifts could 'spill over' to their other (non-BWC) shifts as well. This could lead the design to, all else equal, understate any beneficial effects of BWCs on policing outcomes.

Other RCT designs raise different versions of this spillover concern. For example, Yokum et al. (2019), one of the largest RCTs in this literature carried out with the Washington, DC Metropolitan Police Department (MPD), involved a staggered roll-out of BWC technology for 1,922 MPD officers across seven districts over the June 2015-December 2016 period. Nearly one-third of observed calls during this period involved both a treatment and control officer responding to the same call. This creates the possibility that the control officer's behavior could be changed by the prospect of being recorded by another officer's camera. Such a contamination or 'spillover effect' could again bias the estimates in the direction of understating any effects of BWCs if the presence of a BWC alters the behavior of both treatment and control officers.

Braga et al. (2019) shows exactly how large these spillover effects can be. Implemented through a Boston Police Department 2016 pilot program, the study covered 10 district stations matched

[4]Geller et al. (2020) examined different measurement methods – including circumstances of the use of force (e.g., at a stop, at an arrest), the severity of the use of force (e.g., injury-causing, fatality-causing) or outcome (e.g. an excessive force complaint was filed afterward) – in 11 police departments. They analyzed how distinct measurement methods affected the department's position on a rank-ordered list of departments' force severity and found that the rank-order for a given department could change significantly across different measurements. For example, in their anonymized sample, department "K" was tied for the least severe department when use of force was measured as lethal force but was ranked the most severe department when use of force was measured as all recorded incidents of use of force. (See Table 4 of Geller et al. (2020) for more detail.)

into five pairs, with one district from each pair randomly assigned to BWC treatment. Then across each of the 'treatment' districts, 140 officers were randomly assigned BWCs and 141 officers were randomly assigned to the control group within these 'treatment' districts. They find that comparing treatment officers to control officers within the treatment districts implies a 50% reduction in civilian complaints. But for present purposes, the most striking fact is that control officers within treatment districts *also* experienced a 38.3% reduction in civilian complaints compared to control officers in control districts. This evidence of very large spillover effects implies that RCTs that only randomize officers to treatment and control within the same district may substantially understate the impacts of BWCs.

The fact that we can plausibly sign the bias with these RCTs due to spillover effects makes these studies asymmetrically informative. Spillovers likely work in the direction of understanding BWC impacts, so null results could simply reflect that any impact of BWCs themselves is offset or masked by spillover effects that push the impact estimate towards zero. However, evidence of impacts is informative. Spillovers make these impact estimates lower-bounds for the true effects; put differently, the actual impact under this logic would be even larger still.

The other challenge with even the best of the RCTs in this literature is limited sample size, which in turn limits the statistical power to detect BWC impacts. Consider for instance:

- In the Ariel et al. (2015) study mentioned above, the point estimate, taken at face value, implies that BWC implementation led to a large (52%) reduction in the use of force. Yet the standard errors are so large that even a decline of this magnitude is not statistically significant at the conventional 5 percent cutoff level.
- In White et al. (2018)'s RCT with 149 officers the point estimate for impacts on use of force is -22%, but the confidence interval is so wide we cannot rule out true impacts that are large in the opposite direction (95%CI: -61% to 55%); the same is true for their impact on citizen complaints against the police, equal to -81% (95% CI: -97% to 22%).
- Even one of the largest RCTs by Yokum et al. (2019) winds up suffering from very large confidence intervals. For example, their point estimate for the effect of BWCs on use of force is positive 9% (an increase) but with a 95% confidence interval that allows for anything from large declines (-12%) to large increases (+30%).

We see similar statistical power challenges in many of the other studies in this literature, as shown in Figures 1 and 2, which present the point estimates and confidence intervals for two particularly important outcomes from a benefit-cost perspective (civilian complaints and police use of force) from Lum et al. (2020). The large confidence intervals around the estimates from most of these studies provide one natural candidate explanation for at least part of the variation we see across studies in not just the *magnitude* of the estimated impacts of BWCs on policing outcomes, but even the *sign* of the estimated impact (beneficial vs. adverse). These large confidence intervals also make it difficult to determine how much of the variability in impact estimates is due to something about how BWCs were implemented in a given study (or other features of the local study context), versus simply reflecting idiosyncratic statistical noise in the data. Meta-analysis, by averaging together the results of multiple studies of the (hopefully) same intervention, at least helps improve statistical power to understand average effects across contexts and studies.

Another challenge with many of these studies is that the conditions under which RCTs are carried out tend to differ from how many departments may wind up implementing BWCs at scale. For example, some RCTs, because they are structured as pilot demonstrations, may only involve officers who *volunteered* to participate or use of BWCs. Braga et al. (2018)'s study of the March 2014 adoption of BWC technology within the Las Vegas Metropolitan Police Department notes that the department could not enforce BWC usage among officers due to a provision of their contract with the police union. Although the volunteer and non-volunteer officers within the randomized treatment groups are often similar on observable characteristics, there is still a concern that intervention impacts could vary by unobserved, not just observed, personal attributes. That is, whatever unmeasured attributes are associated with explaining the variation in why some people volunteer or 'comply' with the intervention and others do not could also be relevant for moderating the intervention's impacts.[5] This problem does not compromise the internal validity of these studies, but it does mean that we cannot necessarily generalize these findings to what would happen under a department-wide policy that mandated BWC use by *all* officers.

None of this is intended as criticism for the excellent studies in this literature; almost no policy evaluation is completely ideal. This is intended instead to give readers a sense of the strengths and weakness of the individual studies that go into the meta-analysis, which of these can be overcome by pooling studies together as in meta-analysis (and which can't), and what the two new studies described below might be able to add to our understanding.

3.2 New studies

The largest RCT of BWCs to date was recently released by Braga et al. (2020). The authors examine a BWC pilot program among New York City Police Department (NYPD) officers in 40 NYC precincts associated with the highest number of complaints according to the New York City Citizen Complaint Review Board (CCRB). These precincts were divided into 20 within-borough pairs based on a number of demographic factors and local policing characteristics, with one precinct within each pair assigned to the BWC treatment and the other to the control group. Randomization at the *precinct* level in this manner addresses the spillover problem described above. Their study sample ultimately included 3,889 officers, more than twice as many as in one of the other large RCTs in this area, in Washington DC by (Yokum et al., 2019).

Braga et al. (2020) find that BWC implementation led to a 21.1% decrease in complaints against officers filed with the CCRB. This study also reports a statistically insignificant 1.9% increase in use of force among the officers given BWCs. It is noteworthy that even with a sample size of 3,889 officers, by far the largest RCT to date in this literature, the confidence interval around this estimated impact on police use of force is again enormous, ranging from -26% to +40%.

Braga et al. (2020) also tests for differences in the number of stops, arrests, summonses, domestic incident reports, and crime complaints in response to BWC adoption. Most of those estimates are relatively small in magnitude with large confidence intervals. The number of stops is the only outcome for which they find evidence of a statistically significant relationship to BWCs, a 38.8% increase among treatment officers. Community members stopped by officers with BWCs were also

[5] Or put differently the logic behind local average treatment effects is not limited to compliers who differ from non-compliers along observable dimensions (see for example (Angrist et al., 1996)).

more likely to be non-white compared to the stops made by officers without BWCs; those stops also tended to result in fewer summonses and were more likely to be ruled illegal upon review. It is possible this finding reveals an important adverse risk associated with BWC adoption. As Braga et al. (2020) note, it is also possible that these changes might instead reflect the effect of BWCs on the likelihood that officers report the street stops they make, rather than change the true number of stops that officers actually carry out.

The second new study we add to the analysis is by Kim (2020), who assembles data on police use of force from 2,380 police departments across the country, as well as more detailed data for jurisdictions in New Jersey specifically. The 'natural experiment' at the heart of this study arises from the staggered timing of BWC adoption starting around 2014 in the US. As Kim (2020) notes, the departments that chose to adopt are systematically different from the non-adopters so the primary focus is on comparing trends across the adopting agencies that adopt at different times. The research design capitalizes on the implementation lag between the decision to adopt BWCs and when the agency is actually able to begin deploying BWCs out in the field. This lag reduces the risk that the difference-in-differences design confounds the effects of BWC adoption with those of other events in the jurisdiction that may have prompted local policymakers to decide to adopt BWCs. The identification strategy also assumes that other reforms that departments might adopt would be unlikely to have exactly the same adoption lag as BWCs themselves.

The use of monthly data and large samples enables Kim (2020) to focus narrowly on the time right before and after adoption. Event-study graphs visually suggest that trends in use of force prior to adoption are generally similar for the slightly-earlier versus slightly-later adopters and that the break in trend in policing outcomes seems to occur right around the time of BWC adoption. But it should be said that the confidence intervals in these event study graphs are sizable. The trade-off associated with having additional statistical power to detect impacts on low-base-rate outcomes such as police use of force is that natural experiment studies typically cannot provide the same internal validity guarantees as randomized experiments. Put differently, this study (like the RCTs on BWCs) is not perfect, although its strengths and weaknesses are different from and complementary to those of the RCTs discussed above.

The one outcome for which there are measures available from large numbers of police departments in the national data is the one with the lowest base rate and the one of particular social importance: fatal police use of force. Kim (2020) finds evidence of a statistically significant 41% reduction in officer-involved homicides among BWC adopters. Even with over a thousand agencies contributing to the estimates the confidence intervals are sizable; while the confidence interval excludes zero, the reduction could be as small as -2% and as large as -76%. Part of the challenge may come from the base rate of the outcome itself. In this sample of jurisdictions, the police homicide rate overall is 0.6 per 100,000 population; by comparison, the overall homicide rate in the US over this time period has averaged about 5 or 6 per 100,000. There is no statistically significant change in arrests or crime from BWC adoption at the conventional statistical cutoffs.

Kim (2020) supplements these national findings with data on a broader range of use of force outcomes (including fatal and non-fatal incidents) from New Jersey specifically, using essentially the same research design. That analysis suggests BWC adoption was associated with a statistically significant 20% decline in general use of force among BWC adopters in the state. We might worry that this estimate confounds the effects of BWCs on use of force with officer tendency to report use of force. In an effort to mitigate the risk of this type of measurement error, Kim (2020) also conducts an analysis considering only those incidents of force that resulted in injury, an outcome for which there could potentially be less room for an officer to fail to report the event. Here, the analysis suggests a statistically significant 42% reduction in use of force associated with BWC adoption.

3.3 Meta-Analysis

To assess the effects of BWCs on policing outcomes we build on the recent meta-analysis of Lum et al. (2020) by updating their meta-analytic dataset. Their own meta-analysis includes estimates from 30 studies of BWCs on police and civilian behavior. Most, but not all, of these studies were carried out in the U.S. The authors searched the Global Policing Database for both published and unpublished papers to address the concern that studies with null findings may be less likely to be published in peer-reviewed journals (so-called 'publication bias').[6] The meta-analysis set a relatively low floor for the methodological requirements of a study to be included. The review considered all studies that had any sort of comparison group, as well as some time-series studies that did not have a control group so long as they had at least two years of data and 24 data points for both pre-intervention and post-intervention periods. Studies evaluating the effects of BWC technology among correctional officers, private security, or court officers (e.g., bailiffs) were not included.

Since many studies reported results for multiple versions of each policing outcome, Lum et al. (2020) selected the most general type—for example, overall complaint rates rather than excessive force complaints—to include as the key result to report. A series of codes were created to capture study features such as research design (e.g., RCT or a quasi-experimental design), unit of analysis (e.g., individual officer or shifts), and the way the outcomes were measured (e.g., complaints, sustained complaints, or complaints described as involving use of force). The average effect across all studies (and the standard error for that average effect) is calculated using a random effects model (see, for example, DerSimonian and Kacker (2007)). Results were then presented not just for all studies averaged together but also for subgroups of studies by 'type' (i.e., RCTs vs. quasi-experimental studies).

We present the main results from Lum et al. (2020) for each policing outcome in Table 1. Use of force and complaints are the only outcomes reported in more than half of the thirty studies included in the Lum et al. (2020) meta-analysis. The meta-analysis average effect for these two outcomes is -6.8% and -16.6%, respectively, where only the average effect on complaints is statistically significant at the 95% level. Of the other outcomes reported in Lum et al. (2020), only assaults on officer/officer injuries/resistance and arrests reflect data from more than 10 underlying studies. The average effect on arrests is small and imprecise, whereas the average effect on assaults on officer/officer injuries/resistance suggests if anything a modest increase in response to BWCs of 15.9% (95% confidence interval: -4.9% to 41.3%, p=.143). An important caveat is that this outcome is drawn from only 15 studies and covers a very heterogeneous mix of underlying events.

[6] More details about the exact parameters of the search can be found in Section Five of Lum et al. (2020).

As shown in Table 1, the average effects on each of the other outcomes are small and, with the exception of non-traffic citations (the average effect for which is drawn from only two studies), imprecisely estimated.

We focus on updating the meta-analysis from Lum et al. (2020) for the two outcomes where the new studies Kim (2020) and Braga et al. (2020) contribute the most useful new information and for which we are able to calculate valuations of averted outcomes in our benefit-cost analysis below: civilian complaints and police use of force (fatal, those involving injuries, and overall). Incorporating these new studies does not substantially change the estimated effects of BWC adoption on the other main outcomes reviewed in Lum et al. (2020), such as assaults on officers, arrests, officer-initiated calls for service, dispatched calls for service, traffic stops or tickets, field interviews or stop and frisk, incident reports, response times, non-traffic citations, or time on scene.[7]

Table 2 presents the results of incorporating these new studies into the meta-analysis from Lum et al. (2020). We begin with complaints. Because Kim (2020) does not examine this outcome, the new information here comes entirely from the addition of Braga et al. (2020). Our confidence intervals here reflect statistical sampling uncertainty but cannot quantify the conceptual uncertainty associated with the possibility that these estimates could reflect in part some change in the willingness of residents to file complaints in response to police misconduct. The first row in Panel A of Table 2 reproduces the estimate reported by Lum et al. (2020), which, if taken at face value, (setting possible reporting effects aside for now) would imply a decline in complaints from BWC adoption of -16.6% (95% confidence interval: -30.0% to -0.7%, p=.042). Adding Braga et al. (2020) does not change the point estimate much, equal now to -16.9%, but given that this NYC RCT is the largest to date, we do see a gain in the precision of our estimate (with a confidence interval now -28.2% to -3.8%, p=.013).

For police use of force, the meta-analysis by Lum et al. (2020) presents an estimated impact of -6.8% but with a very wide 95 percent confidence interval (-19.5% to +7.9%, p=.347). Our preferred updated estimate (second row of Panel B in Table 2) combines the data from Lum et al. (2020) with the results from Braga et al. (2020) and two separate estimates from Kim (2020), the estimated impact on fatal police use of force from the national data (the only outcome that can be examined in the national data) and the estimated impact on police use of force with injury from the New Jersey data. The results imply a decline in police use of force of -9.6%, a point estimate that is over 40% larger in absolute value than the figure reported in Lum et al. (2020). The confidence interval around this estimate is also now about 8% smaller than in Lum (-21.3% to +3.8%). The p-value (likelihood of seeing an estimate this size of the true effect were zero) is now p=0.152, or put differently, the odds of seeing an estimate as large as -9.6% if there were no actual impact of BWCs on use of force is about 1 in 7.

The Lum et al. (2020) meta-analysis combines RCTs with quasi-experimental studies that have research designs of varying credibility. We might wonder to what degree the less-credible quasi-experimental studies affect the overall estimated effect on this particularly important outcome.

[7]To conduct the updated meta-analysis we replicate the method laid out in Lum et al. (2020). First, in order to compare across studies, we convert each estimate (effect size and standard error) to a logged incident rate ratio (LIRR) or logged relative incident rate ratio (LRIRR). Next, we fit a random-effects model which takes as inputs the effect size and variance for each study using the *metafor* package in R to fit our random-effects meta-analysis model and use the Restricted Maximum-Likelihood (REML) estimator for tau (which represents heterogeneity across study estimates).

If we use only the RCTs reviewed in Lum et al. (2020) combined with Braga et al. (2020) the estimated impact on citizen complaints is slightly larger in absolute value than what we get from including non-RCTs, -18.8% (95 percent confidence interval -30.6% to -5.1%, p=.009). For use of force using just the RCTS from Lum et al. (2020) with Braga et al. (2020) and the national fatal use of force estimates and NJ non-fatal use of force with injury estimates from Kim (2020), the estimated average effect is now larger in absolute value -11.5% but somewhat less precisely estimated (confidence interval from -25.7% to +5.4%, p=.169).

4 Benefit-Cost Analyses of Body-Worn Cameras

Finally, we turn to the question of what these data mean to police departments or municipalities deliberating about whether to adopt BWCs. There is inevitably some uncertainty in this exercise, not only because of uncertainty about how BWCs affect policing outcomes as expressed in their natural units (complaints, etc.), but also because of additional uncertainties about the exact valuation of these impacts to society. Reasonable people may disagree about which parameters are the right ones to use, so we aspire to present our benefit-cost analyses as transparently as possible for a wide range of parameter values to enable readers to select the estimate they find most compelling. Our baseline estimate for the benefit-cost ratio of BWCs is 4.95, with a defensible range from 0.95 to 26.51.

4.1 Benefit-cost methodology

Given the uncertainties of the available evidence, the question of whether to adopt BWCs for a department requires some way to trade off (or balance) risk and reward. This in turn raises what Cook and Ludwig (2006) called the 'null hypothesis problem.'

The usual convention of scientific peer review focuses on controlling the probability of type I errors (false positives) at 5%. Scientists are willing to live with many instances of type II errors (false negatives), where some successful intervention is mistakenly judged to not work, in exchange for a low risk of false positives. From a policy perspective, this privileges current practice over innovation. As Cook and Ludwig (2006) argue, we might use a different standard for policy decisions: rather than setting the standard purely based on controlling the risk of type I error, the standard might be closer to something like the expected value of benefits and costs for the policies being evaluated. In our case, for instance, our meta-analysis review suggests there is a 1 in 7 chance that there is no effect of BWCs on police use of force, given the average estimate of -9.6%. If the value to society from a reduction in police use of force from BWC adoption is large in relation to costs, which we argue below is the case, would one really want their local mayor or police chief to *not* adopt BWCs but rather stick with the *status quo*?

This then raises the question of quantifying benefits and costs. Benefit-cost analysis requires measuring both benefits and costs in some common metric so that their magnitudes can be compared. The usual convention is dollars. Sometimes this is misunderstood to mean that benefit-cost analysis 'only cares about money,' or only accounts for things that can be bought and sold in markets or measured on a government agency budget. But the necessity of converting costs and benefits into a common metric that enables everything to be compared - usually dollars - does *not* mean that ‘only money matters’ or more generally that only impacts on things bought and sold in markets is of societal concern. Benefit-cost analysis properly conceived also tries to assign dollar values to outcomes not traded in markets, ideally using the *ex-ante* willingness to pay (WTP) by residents who are being asked to support some new policy that is under consideration (see for example (Cook et al., 2000)).

The methods used for this valuation exercise have largely come from the environmental economics literature; Kniesner and Viscusi (2020) provide an excellent review of the underlying conceptual framework and discussion of the available empirical methods, including their strengths and limitations.

- One measurement approach is to look at how people make risk-money trade offs in their daily lives; for example, when people choose between two jobs that are similar except one job entails a higher risk of exposure to some health or safety risk, what wage premium is associated with the more dangerous job? When people are choosing a home, what is the additional house price premium associated with living in a neighborhood with a relatively lower risk from some health or safety threat? This is the so-called *revealed preference* approach because it relies on actual behavior. However, this approach has the conceptual limitation of aiming at the wrong target - it only captures people’s willingness to pay for changes to their own safety (or, in the case of home decisions, those of their households), and so does not include altruistic valuations of changes in health or safety to others in society. This approach also has the practical challenge of often being vulnerable to the problem of confounding from omitted variables (other hedonically-relevant job or house features that are correlated with risk).
- The alternative approach is to ask people to self-report their willingness to pay for reduced exposure to risk in some hypothetical scenario, the so-called *contingent valuation* approach. This approach has the advantage of conceptually aiming at the right target, but the downside of relying on survey responses to hypothetical questions, not actual behavior.

The US federal government, in carrying out benefit-cost analyses of its own policies, programs, and regulations, itself relies on some combination of these two approaches.

Many people correctly note that one normative question raised by this approach flows from the observation that WTP will depend on income. Sunstein (2004) argues that if the public in some jurisdiction is being asked to financially support the costs of the new policy themselves, taking their WTP values at face value is the correct approach; in contrast, if the costs are being supported by some third party (like the federal government) some adjustment should be made to account for income disparities in WTP valuations. In what follows, we assume jurisdictions would be asked to cover the costs of adopting BWCs themselves since in the United States local government is currently responsible for covering most of the costs of police services.

4.2 Costs

Monetization is most straightforward for those things that are bought and sold in the marketplace, where we can use market prices directly. Even in this best-case scenario, it turns out there is a sizable range in estimates of those costs.

- Braga et al. (2017) document the actual costs incurred per camera per year at the Las Vegas Metropolitan Police Department, which is the cost figure also used in the benefit-cost analysis of Kim (2020). This figure includes the amortized costs of the cameras, storage costs, software licenses, IT infrastructure, training, and personnel costs associated with responding to FOIA requests. The estimated cost per camera per year is equal to $1221 (adjusted to 2020 dollars).
- The Police Executive Research Forum (PERF) studied the costs of adopting BWCs in three departments: Phoenix, Arizona, Mesa, Arizona, and Dallas, Texas (PERF (2018)). The figures include annual costs of the camera equipment itself (from either amortized purchase costs or annual contract costs), equipment maintenance costs, data storage, labor for administration and IT, and labor for responding to FOIA requests where relevant (Dallas requires FOIA requesters to pay those costs). The cost estimates per camera per year range from $1221 to $3219 (adjusted to 2020 dollars).[8]

Given these estimates, we show how our benefit-cost analysis changes if we use the lowest and highest defensible cost estimates implied by this literature:[9]

- $1221 per camera per year
- $3219 per camera per year

Because all of our benefit calculations are measured in terms of adverse policing outcomes averted per 100,000 population, we report our cost figures at this level as well. We initially assume that departments commit to having one camera for each officer. If there are 800,000 sworn law enforcement officers in the US[10] to serve a population of 328 million people, then there are .00244 officers per capita or 244 per 100,000.[11] In this case the cost of BWC per year per 100,000 residents at our two cost estimates equals:[12]

- $297,900 per 100,000 residents
- $785,400 per 100,000 residents

Of course not every department has (or needs) one camera per officer. Officers working different shifts can (and in many departments, do) share cameras. So we also present estimates that assume 0.5 cameras per officer, which would make BWC costs per year equal to:

- $149,000 per 100,000 residents per year
- $392,700 per 100,000 residents per year

[8]https://www.policeforum.org/assets/BWCCostBenefit.pdf

[9]The cost figures above capture most of the cost line items listed by the US Department of Justice's Bureau of Justice Assistance BWC 'cost calculator' for departments considering procuring cameras. https://bwctta.com/resources/bwc-cost-and-storage-estimator

[10]https://nleomf.org/facts-figures/law-enforcement-facts

[11]BLS (2019) report that there are 665,280 patrol officers in the US. But many of the BWC reforms underway around the country, as with the new law passed recently in Illinois, require BWC usage by all sworn officers, not just those assigned to patrol, and so we use the estimate of 800,000 total sworn around the country instead.

[12]Estimates are rounded to the nearest 100.

4.3 Benefits

We begin with the exercise of trying to monetize the benefits from averted resident complaints against police. We use two estimates from the literature:

- Braga et al. (2017) use data from the Las Vegas Metropolitan Police Department and note that BWCs reduce the costs of investigating complaints (from reduced investigation time) by the equivalent of $6,882 (adjusted to 2020 dollars).
- Ariel et al. (2015) use data from several jurisdictions including Berkeley, California, Minneapolis, Minnesota, Pittsburgh, Pennsylvania, and London, United Kingdom to calculate the expected costs of compensation to community members, together with administrative costs associated with oversight, and report a cost per complaint (that is, a benefit from each complaint averted) of about $25,400 (adjusted to 2020 dollars).

We prefer the latter estimate because (as we understand it) the former captures only the cost savings to the government agencies themselves, while the latter aspires to also capture the harms to society from the underlying event that led to the complaint being filed in the first place. In our estimates below we also show what happens when we use the alternative estimate instead.

Our quantification of the value to society from averted fatal police use of force relies on the large existing literature on what economists call the 'value per statistical life,' or VSL, which is the approach the federal government uses in carrying out benefit-cost analyses of its own activities (Kniesner and Viscusi, 2020). We report results using three different candidate values:

- $8.4 million, the average contingent valuation estimate for the U.S. taken from Kniesner and Viscusi (2020) (adjusted to 2020 dollars).
- $10.6 million, which as Kniesner and Viscusi (2020) note is the default value used by the U.S. Department of Health and Human Services (adjusted to 2020 dollars).
- $13.5 million, which is the average of the best revealed-preference studies taken from Kniesner and Viscusi (2020) (adjusted to 2020 dollars).

Estimating the value of an averted non-fatal police use of force is more complicated. One reason is that these uses of force are so heterogeneous, ranging from handcuffing suspects (counted as use of force in many studies) to the use of the officer's hands or feet against a civilian up to and including use of pepper spray, tasers, or even non-fatal shootings. Even restricting attention to just police uses of force that result in injury still leaves us with a very heterogeneous mix of events. A second complication is that to the extent to which the literature provides estimates for WTP values relevant for inter-personal violence, they are for crime events among civilians (a civilian offender and a civilian victim). It is possible society views the harms from inter-personal violence differently when it is out-of-policy use of force by police officers against civilians; this is an important limitation we cannot overcome with the data currently available. The building block of our estimation approach is the WTP per averted aggravated assault taken from Cohen and Piquero (2009), updated to 2020 dollars to account for inflation, which equals $109,000. Aggravated assaults themselves are a heterogeneous mix of events that vary by what instrument is used in the attack and the severity of the victim's injuries. To account for the intrinsic uncertainty we also show how our estimates vary when we adjust this figure plus or minus 50% so that the values we use are as follows:

- \$54,500
- \$109,000
- \$163,500

It is possible that even 0.5 times the estimated costs to society from aggravated assaults could be too large on average. In our analysis below we show how our estimates change when we exclude the monetized benefits altogether from reductions in non-fatal police use of force.

The final parameter values we need are population base rates for the policing outcomes averted. The meta-analysis reports impacts in proportional terms ('BWCs cause an X% reduction in adverse policing outcome Y'), so we need to convert these into averted events per 100,000 population for comparison to estimated BWC costs.

We use two estimates for the base rate for *incidence of complaints*:

- 31 citizen complaints per year per 100,000 residents, based on Hickman (2006)[13]
- 66 citizen complaints per 100,000 residents from Braga et al. (2020)[14]

Our estimates for the base rate of *non-fatal police use of force* are as follows:

- 30 uses of force *with injury* per 100,000 per year is the NJ estimate from Kim (2020)
- 144 uses of force per 100,000 per year, from the New Jersey overall police use of force estimate from Kim (2020)[15]
- 198 uses of force per 100,000 residents per year, which comes from the 0.81 incidents of non-fatal use of force per officer per year that is the control group average in the Washington, DC RCT of Yokum et al. (2019)

We use two different estimates for estimating the base rate for *fatal use of force*:

- 0.3 fatal police uses of force per 100,000 per year is the number for 2013 from Sinyangwe et al. (2021), where we use data from 2013 to draw from the period prior to large-scale BWC adoption in order to reduce the risk of endogeneity.
- 0.6 fatal police uses of force per 100,000 residents per year is from the national data from Kim (2020)

[13]Hickman (2006) reports 0.129 complaints per officer annually, which reflects the nationwide average incidence of complaints for officers in local departments. Assuming 244 officers per 100,000 residents this implies 31 complaints per 100,000.

[14]The control group in Braga et al. (2020)'s NYC sample had 0.27 complaints per officer prior to the introduction of BWCs. If we assume 244 officers per 100,000, this implies 0.27*244 = 66 complaints per 100,000 population. We get very similar estimates if we use data for Orlando from Jennings et al. (2015), which reports 0.3 complaints per officer in the control group prior to the introduction of BWCs, and Yokum et al. (2019), which reports 0.28 complaints per officer in the control group (during the study period).

[15]The figures reported in Kim (2020) from NJ are quarterly data, and turn out to be reported on a per 10,000 population, not per 1,000 population figure as reported in the January 25, 2021 version of the paper. Personal communication, Jens Ludwig with Taeho Kim, January 31, 2021.

4.4 Estimated Benefit-Cost Ratios

Table 3 presents our baseline estimates for the ratio of monetized social benefits from BWC to social costs. Each row itemizes a different benefit from BWCs - fatal police use of force, non-fatal police use of force, and civilian complaints against police. The second column presents our baseline estimate for the monetized value of that outcome. The third column presents our estimate for the incidence of that outcome per 100,000 population, followed by the estimated effect (from our meta-analysis above) and then the sum of these three columns for each row (for each benefit category). Below we present sensitivity analyses for changing each of these parameter values, but for this baseline calculation, we have generally selected mid-points from within the range of values discussed above. Under this baseline calculation the benefits of adopting BWC per 100,000 population equal approximately $1.9 million, of which (in this baseline calculation):

- 15.6% comes from averted fatal police uses of force
- 77.5% comes from averted non-fatal uses of force
- 6.8% comes from a reduction in citizen complaints to the police.[16]

The overall baseline estimate for the benefit-cost ratio of BWCs equals 4.95.[17] While our best estimate suggests that BWCs generate benefits to society as a whole that are larger than the costs, a different question we might ask is: from the narrower perspective of the government itself, what are the effects of each dollar invested in BWCs on government budgets specifically? Answering this question requires some additional assumptions and adds some additional uncertainty because there is no completely ideal way to break out what share of the benefits accrue to society as a whole versus to government. But our 'best guess' estimate here is that about $528,000 of the $1.944 million total benefit from adopting BWCs per 100,000 population accrues to government agency budgets.[18] Recognizing the intrinsic uncertainty of these estimates, if they are correct, they would imply that BWCs might even pay for themselves from the perspective of public-sector budgets.

[16]Percentages add to 99.9% rather than 100% due to rounding.

[17]We use the mid-point estimates for the cost of cameras, $392,700 per 100,000 people, value per statistical life (fatal police uses of force), $10.6M, non-fatal force incidents, $109,000, base rate of non-fatal police uses of force, 144 per 100,000, and our baseline estimates for the effect of BWCs on use of force suggested by the updated meta-analysis, -9.6% (we initially assume the same proportional effect on fatal and non-fatal uses of force but revisit this assumption in our sensitivity analysis), and the effect of BWCs on complaints suggested by the updated meta-analysis, -16.9%. We have two plausible base rate estimates for police fatal use of force, we conservatively use the lower of the two, 0.3 per 100,000 population; similarly, we have two estimates for the base rate of complaints, we conservatively use the lower of the two, 31 per 100,000. Both of those choices will have the effect of reducing benefit-cost ratios relative to the alternative choices. For our estimates for the cost of complaints, we use the higher of the two values we consider (but the one we believe is most conceptually justified since it aspires to measure impacts on society, not just on government budgets), $25,400.

[18]The estimates from Ariel et al. (2015), benefit-cost supplement, suggest that the benefits of averted citizen complaints against the police come from reduced investigation costs (a benefit limited to the government budget) together with reduced payouts or settlements (which is a transfer from the public-sector budget to society and so relevant to the public-sector budget), so we count the entire inflation-adjusted $25,400 per complaint as a benefit to government per averted complaint. In sum, we estimate that reductions in complaints account for savings to government of $133,071 per 100,000 population. While we have no way to directly apportion the benefits from averted non-fatal and fatal police uses of force into changes in government budgets versus benefits to society as a whole, we do have some estimates for this type of breakdown for the costs of crime. So absent any alternative, we assume the division between benefits to society as a whole versus to government budgets is the same for non-fatal and fatal police uses of force

Our final table shows how our estimates for the ratio of benefits and costs from the perspective of society as a whole change under alternative choices of the key parameters. The columns show what happens when we adjust our assumptions about the value to society per averted fatal police use of force and non-fatal use of force. The rows show what happens when we alter our assumptions about the costs of cameras, the estimated effect we use for fatal and non-fatal police uses of force (including what happens if we assume no impacts at all on either type of use of force, or if we allow for the impacts of use of force to be different for fatal vs. non-fatal events), and if we adjust the base rates we assume for each adverse policing outcome. The smallest of the estimated benefit-cost ratios we present, which assumes the effect on non-fatal use of force is zero and uses the lower-bound estimate for the dollar value of fatal use of force, is 0.95. The largest estimate, which uses Kim (2020)'s estimated impact on non-fatal use of force from his New Jersey data (rather than averaging with the other studies) and the high-end estimate for the valuation to society from averted fatal and non-fatal use of force, equals 26.51.

Put differently, while reasonable people can disagree about what the best choices are for each of the parameters in our benefit-cost analysis, under most defensible choices for these parameters BWCs would seem to generate benefits to society in excess of costs. Given the number of parameters, we have not calculated benefit-cost ratios for every possible combination. But we hope we have provided enough information for readers to calculate benefit-cost ratios using whatever alternative parameter value combinations seem appropriate, beyond the set of results we have shown here. And as new parameter estimates arise as the BWC research literature moves forward, it will also be possible for readers to update these figures.

5 Conclusion

Our analysis suggests the ratio of benefits to society from adoption of BWC to the costs is on the order of 5 to 1. If our analysis is correct this implies that from society's perspective this is the equivalent of the ability to turn a \$1 bill into a \$5 bill. A different policy simulation comes from the possibility of re-arranging resources more narrowly within the criminal justice system itself. Our estimated benefit-cost analysis of 5:1 from body-worn cameras is substantially higher, for example, than the estimated benefit-cost ratio of 2:1 for additional spending on hiring more police Chalfin and McCrary (2018).

In benefit-cost analyses of federal policies and regulations, impacts on health and safety outcomes wind up dominating the analysis. The same turns out to be true in our benefit-cost analysis of body-worn cameras, with changes in police use of force accounting for the largest share of the benefits. However, this is the outcome for which there is also the most uncertainty about the size of the impact, and especially for non-fatal police use of force, the appropriate valuation

as for non-fatal civilian-on-civilian violence (aggravated assault) and fatal civilian-on-civilian violence (murder), recognizing this is an untestable assumption that may either overstate or understate the true government budget share for police use of force cases. With that caveat in mind, the estimates from Cohen and Piquero (2009), Table 5, suggest that around 25% of the costs of aggravated assaults accrue to the justice system and that around 6% of the costs of homicides accrue to the justice system. In that case, of the total \$1.944 million estimated benefit per 100,000 population from adoption of BWCs, from Table 3, these calculations would (roughly) imply that \$133,071 + (25%)x(\$1,506,816) + (6%)x(\$304,128) = \$528,000 shows up as a reduction in some government budget line-item.

of changes in that outcome from society's perspective. This would seem to be a particularly high priority for future research.

Our analysis has also estimated the *average* effects of BWCs among the departments that have adopted this technology to date. It is very possible that these departments are not representative in terms of how this technology translates into changes in policing outcomes. For example, Allcott (2015) notes the possibility that early adopters, particularly those most willing to participate in a research study, may be the ones with the most pronounced beneficial impacts - or 'site-selection bias.' This will make follow-up analysis as the technology diffuses across the country particularly valuable.

Our analysis also inevitably relies on estimates of the effects of BWCs as departments currently deploy them. But those deployment practices may well change over time if professional associations, advocates, or the federal government push departments to adopt more standardized BWC policies and practices. It may also be the case that new insights into how to capitalize on the technology can lead to qualitatively new types of impacts. For example, in a study of discrimination in police-civilian interactions, Voigt et al. (2017) apply machine learning tools to data from BWCs. That was just a proof-of-concept study but highlights the potential for artificial intelligence to change and expand the ability of departments to process and monitor BWC video footage. The effects of BWCs on policing outcomes could be quite different in a world in which artificial intelligence increasingly makes it possible to mass-audit BWCs to detect police-civilian interactions that are worthy of further investigation. This is, in short, a policing reform technology that is worth continued monitoring and evaluation.

References

Allcott, H. (2015). Site selection bias in program evaluation. *The Quarterly Journal of Economics 130*(3), 1117–1165.

Angrist, J. D., G. W. Imbens, and D. B. Rubin (1996). Identification of causal effects using instrumental variables. *Journal of the American statistical Association 91*(434), 444–455.

Ariel, B. (2016). Police body cameras in large police departments. *The Journal of Criminal Law and Criminology (1973-)*, 729–768.

Ariel, B., W. A. Farrar, and A. Sutherland (2015). The effect of police body-worn cameras on use of force and citizens' complaints against the police: A randomized controlled trial. *Journal of quantitative criminology 31*(3), 509–535.

Ariel, B., A. Sutherland, D. Henstock, J. Young, P. Drover, J. Sykes, S. Megicks, and R. Henderson (2016a). Report: Increases in police use of force in the presence of body-worn cameras are driven by officer discretion: A protocol-based subgroup analysis of ten randomized experiments. *Journal of experimental criminology 12*(3), 453–463.

Ariel, B., A. Sutherland, D. Henstock, J. Young, P. Drover, J. Sykes, S. Megicks, and R. Henderson (2016b). Wearing body cameras increases assaults against officers and does not reduce police use of force: Results from a global multi-site experiment. *European journal of criminology 13*(6), 744–755.

Ariel, B., A. Sutherland, D. Henstock, J. Young, P. Drover, J. Sykes, S. Megicks, and R. Henderson (2017). "contagious accountability" a global multisite randomized controlled trial on the effect of police body-worn cameras on citizens' complaints against the police. *Criminal justice and behavior 44*(2), 293–316.

Blitz, M. J. (2015). Police body-worn cameras: Evidentiary benefits and privacy threats. *Advance 9*, 43.

BLS (2019). Occupational employment and wages. Technical report, U.S. Bureau of Labor Statistics.

Braga, A. A., L. M. Barao, G. M. Zimmerman, S. Douglas, and K. Sheppard (2019). Measuring the direct and spillover effects of body worn cameras on the civility of police–citizen encounters and police work activities. *Journal of Quantitative Criminology*, 1–26.

Braga, A. A., J. R. Coldren, W. H. Sousa, D. Rodriguez, and O. Alper (2017). *The benefits of body-worn cameras: New findings from a randomized controlled trial at the Las Vegas Metropolitan Police Department*. CNA Analysis & Solutions Arlington, VA.

Braga, A. A., J. MacDonald, and J. McCabe (2020, November). Body worn cameras, lawful police stops, and nypd officer compliance: A cluster randomized controlled trial. Technical report, Working paper, http://nypdmonitor.org/wp-content/uploads/2020/12/12th-Report.pdf.

Braga, A. A., W. H. Sousa, J. R. Coldren, and D. Rodriguez (2018). The effects of body-worn cameras on police activity and police-citizen encounters. *The Journal of Criminal Law and Criminology (1973-) 108*(3), 511–538.

Chalfin, A. and J. McCrary (2018). Are us cities underpoliced? theory and evidence. *Review of Economics and Statistics 100*(1), 167–186.

Cohen, M. A. and A. R. Piquero (2009). New evidence on the monetary value of saving a high risk youth. *Journal of Quantitative Criminology 25*(1), 25–49.

Cohen, M. A., R. T. Rust, S. Steen, and S. T. Tidd (2004). Willingness-to-pay for crime control programs. *Criminology 42*(1), 89–110.

Cook, P. J. and J. Ludwig (2000). *Gun Violence: The Real Costs*. Oxford University Press.

Cook, P. J. and J. Ludwig (2006). Aiming for evidence-based gun policy. *Journal of Policy Analysis and Management 25*(3), 691–735.

Cook, P. J., J. Ludwig, et al. (2000). *Gun violence: The real costs*. Oxford University Press on Demand.

Dehejia, R. H. and S. Wahba (2002). Propensity score-matching methods for nonexperimental causal studies. *Review of Economics and statistics 84*(1), 151–161.

Demir, M., R. Apel, A. A. Braga, R. K. Brunson, and B. Ariel (2020). Body worn cameras, procedural justice, and police legitimacy: a controlled experimental evaluation of traffic stops. *Justice quarterly 37*(1), 53–84.

DerSimonian, R. and R. Kacker (2007). Random-effects model for meta-analysis of clinical trials: an update. *Contemporary clinical trials 28*(2), 105–114.

Douglas, S. (2020). The effects of body-worn cameras on violent police victimization. *Policing: A Journal of Policy and Practice*.

Feigenberg, B., D. Fitzpatrick, J. Ludwig, and M. Williams Jr. (2021, January). Methods for research review. Technical report, Council on Criminal Justice and University of Chicago Crime Lab.

Gaub, J. E., D. E. Choate, N. Todak, C. M. Katz, and M. D. White (2016). Officer perceptions of body-worn cameras before and after deployment: A study of three departments. *Police quarterly 19*(3), 275–302.

Gaub, J. E., N. Todak, and M. D. White (2020a). The distribution of police use of force across patrol and specialty units: a case study in bwc impact. *Journal of Experimental Criminology*, 1–17.

Gaub, J. E., N. Todak, and M. D. White (2020b). One size doesn't fit all: The deployment of police body-worn cameras to specialty units. *International Criminal Justice Review 30*(2), 136–155.

Gaub, J. E. and M. D. White (2020). Open to interpretation: Confronting the challenges of understanding the current state of body-worn camera research. *American Journal of Criminal Justice*, 1–15.

Geller, A., P. A. Goff, T. Lloyd, A. Haviland, D. Obermark, and J. Glaser (2020, Oct). Measuring racial disparities in police use of force: Methods matter. *Journal of Quantitative Criminology*.

Goh, L. S. (2020). *THE POLICY EVALUATION OF MEASURES TO REDUCE POLICE USE OF FORCE*. Ph. D. thesis, University of Pennsylvania.

Grossmith, L., C. Owens, W. Finn, D. Mann, T. Davies, and L. Baika (2015). Police, camera, evidence: London's cluster randomised controlled trial of body worn video. *London: College of Policing*.

Harrer, M., P. Cuijpers, T. Furukawa, and D. Ebert (2019). Doing meta-analysis in r: A hands-on guide. *PROTECT Lab Erlangen*.

Headley, A. M., R. T. Guerette, and A. Shariati (2017). A field experiment of the impact of body-worn cameras (bwcs) on police officer behavior and perceptions. *Journal of Criminal Justice 53*, 102–109.

Henstock, D. and B. Ariel (2017). Testing the effects of police body-worn cameras on use of force during arrests: A randomised controlled trial in a large british police force. *European journal of criminology 14*(6), 720–750.

Hickman, M. J. (2006). Citizen complaints about police use of force. Technical report.

Hyland, S. (2018). *Body-worn cameras in law enforcement agencies, 2016*. US Department of Justice, Office of Justice Programs, Bureau of Justice . . . .

Iboshi, K. (2021, Mar). Portland city council reluctant to move forward on police body-worn cameras.

Jennings, W. G., L. A. Fridell, M. Lynch, K. K. Jetelina, and J. M. Reingle Gonzalez (2017). A quasi-experimental evaluation of the effects of police body-worn cameras (bwcs) on response-to-resistance in a large metropolitan police department. *Deviant behavior 38*(11), 1332–1339.

Jennings, W. G., M. D. Lynch, and L. A. Fridell (2015). Evaluating the impact of police officer body-worn cameras (bwcs) on response-to-resistance and serious external complaints: Evidence from the orlando police department (opd) experience utilizing a randomized controlled experiment. *Journal of criminal justice 43*(6), 480–486.

Katz, C. M., D. E. Choate, J. R. Ready, and L. Nuño (2014). Evaluating the impact of officer worn body cameras in the phoenix police department. *Phoenix, AZ: Center for Violence Prevention & Community Safety, Arizona State University*.

Keltner, D., D. H. Gruenfeld, and C. Anderson (2003). Power, approach, and inhibition. *Psychological review 110*(2), 265.

Kim, T. (2020). Facilitating police reform: Body cameras, use of force, and law enforcement outcomes. Technical report, Working Paper.

Kniesner, T. J. and W. K. Viscusi (2020). The value of a statistical life. *Oxford Research Encyclopedia of Economics and Finance*.

Koper, C. S., C. Lum, J. J. Willis, D. J. Woods, and J. Hibdon (2015). Realizing the potential of technology in policing: A multi-site study of the social, organizational, and behavioral aspects of implementing policing technologies. In *Report to the National Institute of Justice, US Department of Justice. Fairfax, VA: Center for Evidence-Based Crime Policy, George Mason University and Police Executive Research Forum*.

Koslicki, W. M., D. A. Makin, and D. Willits (2019). When no one is watching: evaluating the impact of body-worn cameras on use of force incidents. *Policing and society*, 1–14.

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), Article–number.

Lum, C., M. Stoltz, C. S. Koper, and J. A. Scherer (2019). Research on body-worn cameras: What we know, what we need to know. *Criminology & public policy 18*(1), 93–118.

Lum, C. M., C. S. Koper, L. M. Merola, A. Scherer, and A. Reioux (2015). *Existing and ongoing body worn camera research: Knowledge gaps and opportunities*. George Mason University.

McClure, D., N. La Vigne, M. Lynch, L. Golian, D. Lawrence, and A. Malm (2017). How body cameras affect community members' perceptions of police. *Results from a randomized controlled trial of one agency's pilot. Washington, DC: Urban Instititue*.

Merola, L., C. Lum, C. S. Koper, and A. Scherer (2016). Body worn cameras and the courts: A national survey of state prosecutors. *Fairfax, VA: George Mason University*.

Miller, L. and J. Toliver (2014). Implementing a body-worn camera. In *Police executive research forum and United States of America*.

Owens, E., D. Weisburd, K. L. Amendola, and G. P. Alpert (2018). Can you build a better cop? experimental evidence on supervision, training, and policing in the community. *Criminology & Public Policy 17*(1), 41–87.

PERF (2018). Cost and benefits of body-worn camera deployments. Technical report, Police Executive Research Forum.

Peterson, B. E. and D. S. Lawrence (2020). Do the effects of police body-worn cameras on use of force and complaints change over time? results from a panel analysis in the milwaukee police department. *Criminal Justice and Behavior*.

Peterson, B. E., L. Yu, N. La Vigne, and D. S. Lawrence (2018). The milwaukee police department's body-worn camera program. *Washington, DC: Urban Institute*.

Ramsey, C. H. and L. O. Robinson (2015, May). Final report of the president's task force on 21st century policing. Technical report, Office of Community Oriented Policing Services.

Reaves, B. A. (2015). Local police departments, 2013: Equipment and technology. *Washington, DC: Bureau of Justice Statistics*.

Sinyangwe, S., D. McKesson, and J. Elzie (2021). Mapping police violence. Technical report.

Sunstein, C. R. (2004). Are poor people worth less than rich people? disaggregating the value of statistical lives. *U Chicago Law & Economics, Olin Working Paper* (207), 04–05.

Tregle, B., J. Nix, and J. T. Pickett (2020). Body-worn cameras and transparency: Experimental evidence of inconsistency in police executive decision-making. *Justice Quarterly*, 1–23.

Voigt, R., N. P. Camp, V. Prabhakaran, W. L. Hamilton, R. C. Hetey, C. M. Griffiths, D. Jurgens, D. Jurafsky, and J. L. Eberhardt (2017). Language from police body camera footage shows racial disparities in officer respect. *Proceedings of the National Academy of Sciences 114*(25), 6521–6526.

White, M. D., J. E. Gaub, and N. Todak (2018). Exploring the potential for body-worn cameras to reduce violence in police-citizen encounters. *Policing: a journal of policy and practice 12*(1), 66–76.

Yokum, D., A. Ravishankar, and A. Coppock (2019). A randomized control trial evaluating the effects of police body-worn cameras. *Proceedings of the National Academy of Sciences 116*(21), 10329–10332.

6 Figures and Tables

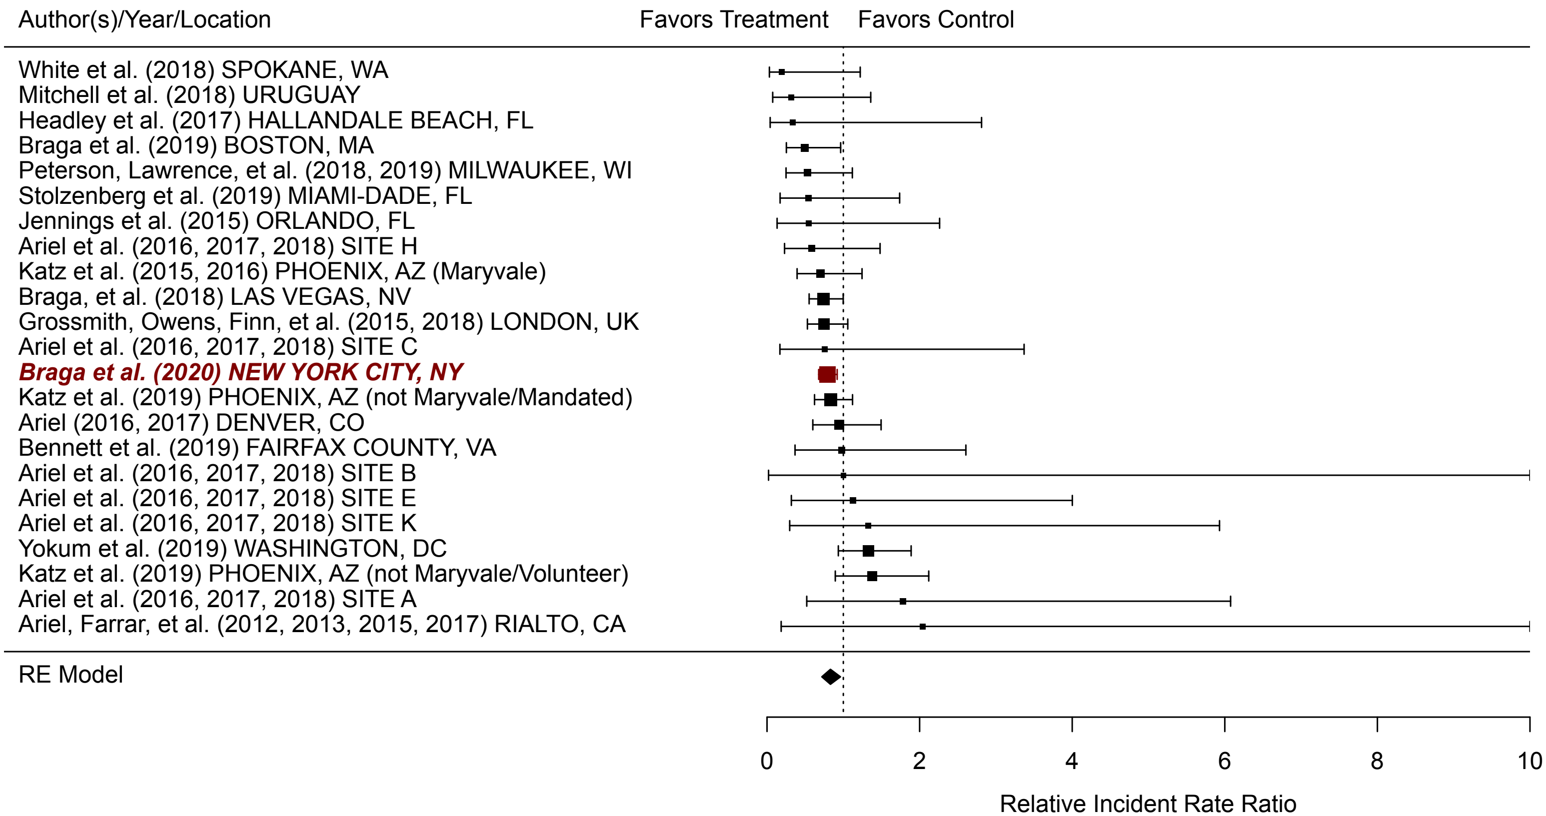

Figure 1: Lum et al. (2020) Meta-Analysis of BWCs and Complaints: Updated to Include Braga et al. (2020)

Note: Meta-analysis replication made possible using code and data publicly shared by Lum et al. (2020). We include each study from Lum et al. (2020) that reports an estimate for the effect of BWCs on complaints as well as one new study (Braga et al. (2020)) which is highlighted in red and italicized. Points to left of the vertical dashed line reflect studies which estimated a reduction in complaints in response to BWC adoption, whereas those to the right reflect estimated increases in complaints. Horizontal bands reflect 95% confidence intervals.

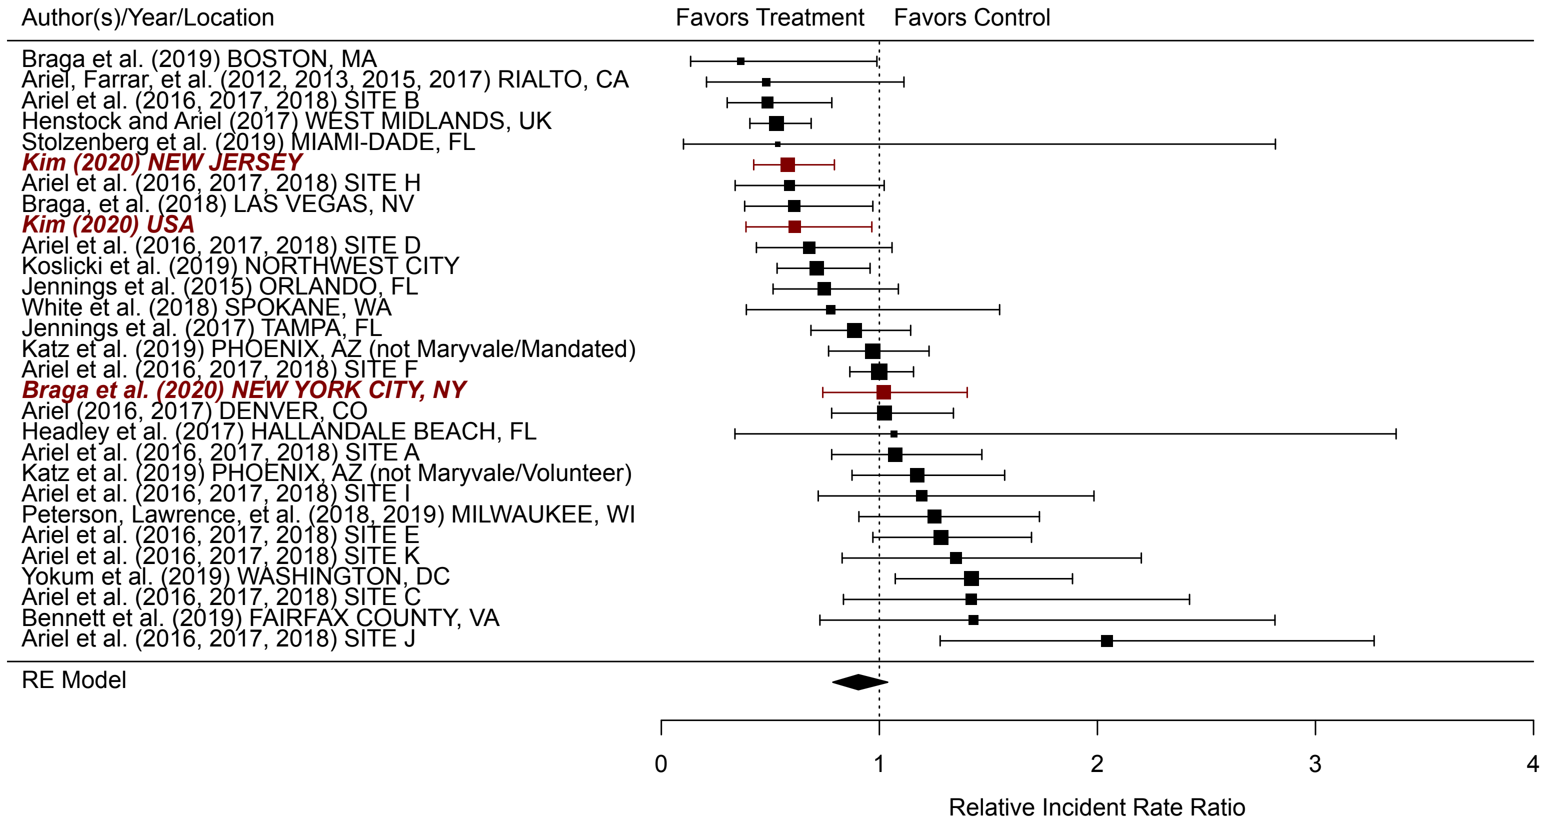

Figure 2: Lum et al. (2020) Meta-Analysis of BWCs and Use of Force: Updated to Include Braga et al. (2020) and Kim (2020)

Note: Meta-analysis replication made possible using code and data publicly shared by Lum et al. (2020). We include each study from Lum et al. (2020) that reports an estimate for the effect of BWCs on complaints as well as three new studies, one from Braga et al. (2020) and two from Kim (2020), which are highlighted in red and italicized. Points to left of the vertical dashed line reflect studies which estimated a reduction in use of force in response to BWC adoption, whereas those to the right reflect estimated increases in use of force. Horizontal bands reflect 95% confidence intervals.

Table 1: Body-Worn Cameras: Lum et al. (2020) Meta-Analysis

Construct	Number of Studies	Mean % Change	Lower CI	Upper CI	p
Use of force	26	-6.8	-19.5	7.9	.347
Complaints against officer	22	-16.6	-30.0	-0.7	.042
Assault on officer/officer injuries/resistance	15	15.9	-4.9	41.3	.143
Arrests	13	-3.9	-12.7	5.8	.416
Officer-initiated calls for service	8	3.8	-5.2	13.5	.422
Dispatched calls for service	6	2.6	-3.0	8.6	.365
Traffic stops or traffic tickets	5	-5.0	-33.1	34.8	.772
Field interviews or stop and frisk	4	-12.0	-37.5	24.0	.466
Incident reports	3	-8.0	-21.7	8.0	.309
Response time	3	-0.2	-1.7	1.3	.812
Non-traffic citations	2	6.4	5.8	7.1	.000
Time on scene	1	-4.6	-12.0	3.4	.251

Note: This table presents results from the meta-analysis of BWC studies reported in Lum et al. (2020). The results included here are equivalent to those presented in Table 5 of Lum et al. (2020) which reports the primary meta-analysis specification for each construct including all studies, rather than restricting the sample by study design or other characteristics.

Table 2: Body-Worn Cameras: Updated Meta-Analysis

Panel A: Complaints

Studies Included	Number of Studies	Mean % Change	Lower CI	Upper CI	p
Lum et al. (2020) *(All studies)*	22	-16.6	-30.0	-0.7	.042
Lum et al. (2020) *(All studies)*, Braga et al. (2020)	23	-16.9	-28.2	-3.8	.013
Lum et al. (2020) *(RCTs only)*, Braga et al. (2020)	16	-18.8	-30.6	-5.1	.009

Panel B: Use of Force

Studies Included	Number of Studies	Mean % Change	Lower CI	Upper CI	p
Lum et al. (2020) *(All studies)*	26	-6.8	-19.5	7.9	.347
Lum et al. (2020) *(All studies)*, Braga et al. (2020), Kim (2020) *(National fatal force)*, Kim (2020) *(NJ force w/ injury)*	29	-9.6	-21.3	3.8	.152
Lum et al. (2020) *(RCTs only)*, Braga et al. (2020), Kim (2020) *(National fatal force)*, Kim (2020) *(NJ force w/ injury)*	22	-11.5	-25.7	5.4	.169

Note: Meta-analysis replication made possible using code and data publicly shared by Lum et al. (2020). For each row in the table we replicate the meta-analysis procedure performed in Lum et al. (2020), first performing a simple replication using only those studies reported in Lum et al. (2020) and then incorporating data from three new studies, one in Braga et al. (2020) and two in Kim (2020). The first row of each panel, A and B, reports the simple replication of Lum et al. (2020) for each of two outcomes: complaints and use of force, respectively. The subsequent rows in each panel include studies from Lum et al. (2020) and new data from Braga et al. (2020) (for both complaints and use of force) and Kim (2020) (for only use of force). For specifications with Kim (2020), the second and third rows of Panel B, we include two separate outcomes: fatal use of force from the national study and use of force resulting in injury from the New Jersey study (see main text for additional details). Finally, the third row in each panel restricts the set of studies we include from Lum et al. (2020) to only include those carried out as RCTs.

Table 3: Baseline Benefit-Cost Analysis

	Social harm per event	Incidence per 100K residents	Estimated reduction from BWC	Value per 100K people
Panel A: Benefits				
Fatal force	$10,560,000	0.3	0.096	$304,128
Non-fatal force	$109,000	144	0.096	$1,506,816
Complaints	$25,400	31	0.169	$133,071
Total averted harm per 100K people per year				$1,944,015
Panel B: Costs				
Cost of BWC per 100K people per year				$392,700
Panel C: Total				
Benefit-cost ratio				4.95

Note: This table presents our baseline benefit-cost estimates accounting for the effect of BWCs on use of force and complaints as described in the main text. We estimate benefits and costs based on the expected annual incidence and social harm associated with each outcome. Estimated reductions are in proportional terms (e.g. a reduction of 0.1 for a given event would translate to a 10% reduction in the incidence of that event). All dollar values in 2020 dollars.

Table 4: Benefit-Cost Sensitivity Analysis

	Value per statistical life / Value per non-fatal force incident averted								
Changes relative to baseline:	8.4M / 54.5k	10.6M / 54.5k	13.5M / 54.5k	8.4M / 109k	10.6M / 109k	13.5M / 109k	8.4M / 163.5k	10.6M / 163.5k	13.5M / 163.5k
Panel A: Alternative camera cost estimates per 100K population									
$149,000 (0.5 cameras per officer, camera cost $1221)	7.56	7.99	8.57	12.62	13.05	13.62	17.68	18.10	18.68
$297,900 (1 camera per officer, camera cost $1221)	3.78	4.00	4.28	6.31	6.53	6.81	8.84	9.05	9.34
$785,400 (1 camera per officer, camera cost $3219)	1.44	1.52	1.63	2.39	2.48	2.58	3.35	3.43	3.54
Panel B: Modify fatal use of force estimate									
Effect = 0	2.26	2.26	2.26	4.18	4.18	4.18	6.09	6.09	6.09
Effect = 11.5% (alternative meta-analysis result)	2.99	3.19	3.45	4.91	5.10	5.37	6.83	7.02	7.28
Effect = 41% (Kim (2020) national data estimate)	4.87	5.56	6.50	6.79	7.48	8.42	8.71	9.40	10.34
Panel C: Modify non-fatal use of force estimate									
Effect = 0	0.95	1.11	1.33	0.95	1.11	1.33	0.95	1.11	1.33
Effect = 11.5% (alternative meta-analysis result)	3.25	3.41	3.63	5.55	5.71	5.93	7.85	8.01	8.23
Effect = 20% (Kim (2020) total use of force estimate from NJ data)	4.95	5.11	5.33	8.95	9.11	9.33	12.94	13.10	13.32
Effect = 42% (Kim (2020) use of force with injury from NJ data)	9.35	9.51	9.73	17.74	17.90	18.12	26.13	26.29	26.51
Panel D: Modify complaint effect									
Effect = 0	2.53	2.69	2.91	4.45	4.61	4.83	6.37	6.53	6.75
Panel E: Modify fatal use of force base rate									
0.6 per 100K (Kim (2020) national data)	3.48	3.81	4.24	5.40	5.72	6.16	7.32	7.64	8.08
Panel F: Modify non-fatal use of force base rate									
30 per 100k (Kim (2020) use of force with injury from NJ data)	1.35	1.51	1.73	1.75	1.91	2.13	2.15	2.31	2.53
198 per 100k (Yokum et al. (2019) study-period control group)	3.59	3.75	3.97	6.23	6.39	6.61	8.87	9.03	9.25
Panel G: Modify complaint base rate									
66 per 100k (Braga et al. (2020) pre-period control group)	3.25	3.41	3.63	5.17	5.33	5.55	7.09	7.25	7.47
Panel H: Modify complaint cost									
$6882 (Braga et al. (2017) estimate of labor time saved)	2.62	2.78	3.00	4.54	4.70	4.92	6.46	6.62	6.84

Note: This table presents the sensitivity analysis for our baseline benefit-cost ratio using alternative estimates for the cost of cameras and the benefits of use of force incidents averted as outlined in the main text. All dollar values in 2020 dollars.