Preaching to the Choir: A Problem of Participatory Interventions

Rebecca Hanson, University of Florida
Dorothy Kronick, University of California
Tara Slough, New York University

Scholars and policymakers alike have endorsed dialogue as a remedy for the global crisis in police–community relations. But the community members who choose to engage in dialogue with police officers, we find, are those who trust the police to begin with (i.e., those who are hardest to impress and easiest to disappoint). In a large-scale field experiment in Medellín, Colombia, we discover that those who most trusted the police at baseline were twice as likely to attend police–community meetings as those who least trusted the police. We document similar patterns using survey data from 23 countries: people who most trust a given institution (e.g., city government) are the most likely to show up at that institution's public meetings. This preaching-to-the-choir problem undermined the effect of the intervention in Medellín and, we argue, poses an under-recognized threat to related initiatives across the globe.

The people who show up for public meetings with their city councillors, school board members, or local police officers are different from the people who stay home. In Cambridge, Massachusetts, for example, participants in community meetings about housing "differ starkly from the broader population" (Einstein, Glick, and Palmer 2019, 97); they are older, whiter, and much more likely to be homeowners. In Brazil, the organizations that support participatory budgeting disproportionately attract women (Wampler 2012, 349); at community meetings in Sierra Leone and the Democratic Republic of the Congo, men speak at least twice as often as women (Casey, Glennerster, and Miguel 2012, 1798; Humphreys, de la Sierra, and Van der Windt 2015, 12). People who attend public meetings differ not only in their demographic and socioeconomic characteristics but also in their political attitudes (Fiorina 1999). In their landmark study of community policing in Chicago, for example, Skogan and Hartnett (1999, 153) observe a "strong establishment bias" among attendees at police–community meetings: relative to neighbors, participants were much more likely to report that police were doing a good job.

We argue that this widely recognized fact has underrecognized implications. In many settings, public officials convene town hall meetings in hopes of building trust in government. But if the people who show up are those who trust the government to begin with, then the meetings may not exceed attendees' expectations—and may even disappoint. This preaching-to-the-choir problem can undermine efforts to win hearts and minds at public meetings.

We investigate this difficulty in two ways. First, we use a large-scale field experiment in Medellín, Colombia, to study the effects of police–community meetings on trust in police.

Rebecca Hanson (r.hanson@ufl.edu) is an assistant professor at the University of Florida, Gainesville, FL 32611. Dorothy Kronick (kronick@berkeley.edu) is an assistant professor at the University of California, Berkeley, Berkeley, CA 94720. Tara Slough (tara.slough@nyu.edu) is an assistant professor at New York University, New York, NY 10012.

This research was conducted in compliance with relevant laws and was approved by IRBS at the University of Florida (#IRB201701883) and the University of Pennsylvania (#827913). Support for this research was provided by the UK Foreign, Commonwealth & Development Office (205133), the Center for Effective Global Action Economic Development and Institutions initative, and the Corporación Andina de Fomento. The preanalysis plan for this experiment is available at https://osf.io/dny5t. Replication files are available in the *JOP* Dataverse (https://dataverse.harvard.edu/dataverse/jop). The empirical analysis has been successfully replicated by the *JOP* replication analyst. An online appendix with supplementary material is available at https://doi.org/10.1086/732983.

Published online March 7, 2025.

  https://doi.org/10.1086/732983

We find that the meetings did not build trust in the police or boost crime reporting (according to preregistered intent-to-treat estimates)—despite the fact that police held hundreds of meetings, reaching thousands of residents. One reason for the null effect on trust, we discover, is that the intervention failed to overcome the problem of preaching to the choir: people who appreciated the police to begin with were much more likely to show up. They were hard to impress and easily disappointed.

Second, we establish that the problem of preaching to the choir in public meetings extends far beyond Medellín and also beyond police to other government institutions. Using data from coordinated experiments on police–community meetings in other countries (Blair et al. 2021), we document a similar pattern across four of the six cases. And using publicly available surveys from 23 countries in Latin America, we find that people who trust a given institution are much more likely to show up at that institution's public meetings. People who trust city government, for example, attend town hall meetings at higher rates than their distrustful neighbors in 106 of the 115 country–survey waves in these data.

In our case, the preaching-to-the-choir problem helps explain the null results of a field experiment. But the problem is not specific to null results. Indeed, as we show in a stylized model, it would also plague a participatory intervention in which every single event provided a maximally positive signal about the government. In that case, even if all participants held rosy views of the government to begin with (relative to nonparticipants), the intervention could still "work," in the sense of measurably improving perceptions. Yet it might not actually work: the costs of the intervention could easily exceed the benefits of converting supporters into devotees.

This finding has implications both for theory and for policymakers. One theme of recent literature on law enforcement and state security agencies in Latin America is that well-meaning policies often do not work as intended because police officers and/or criminal groups respond strategically, in unanticipated ways (Acemoglu et al. 2020; Calderón et al. 2015; Castillo and Kronick 2020; Cruz and Durán-Martínez 2016; Dipoppa 2021; Kronick 2020; Lessing 2017; Magaloni, Franco-Vivanco, and Melo 2020; Trejo and Ley 2020). We add that unanticipated responses from ordinary citizens (in this case, residents considering attending police–community meetings) can similarly undermine interventions that might otherwise improve welfare.

Literature on participatory governance often considers the implications of who participates for resource allocation, sometimes finding that participatory forums are inclusive and improve resource allocation, other times finding that they amplify vocal minorities and distort policy outcomes (Falleti and Riofrancos 2018; Goldfrank 2007; Grossman 2014; Mayka 2019; McNulty 2013; Wampler 2008; Yoder 2020). We build on this work by considering the implications of who participates for a different outcome: trust in government. Just as who participates affects resource allocation, it is key to understanding the effect of participatory interventions on beliefs.

We also build on previous work about community policing. One debate over community policing turns on whether (and when) community policing constitutes a meaningful transformation of enforcement and police service provision, as opposed to window dressing used to placate vocal constituencies without altering police agencies' priorities. Blair et al. (2021, 1), summarizing the results of the meta-study of which our experiment forms part, conclude that community policing "does not, at least immediately and on its own, lead to major improvements in citizen–police relations" in the Global South; instead, "structural reforms to the police may be needed." We show that one impediment to effective community-policing interventions—preaching to the choir—lies outside of police agencies themselves and is likely not unique to the Global South.

Our findings contribute to literature on street-level bureaucrats and trust in government. Street-level bureaucrats, especially police officers, powerfully shape trust in the criminal justice system and in government as a whole (Hough et al. 2010; Tyler 2006; Tyler and Jackson 2013). And while previous work identifies how seemingly innocuous selection into engagement with bureaucrats affects their incentives and thereby service delivery (Slough 2021a, 2021b), we show how it mutes the effectiveness of a participatory government program, identifying preaching to the choir as a phenomenon that paradoxically erodes the efficacy of participatory interventions.

Finally, we build on a methods literature about estimating treatment effects in experiments with selection (de Benedictis-Kessner et al. 2019; Knox et al. 2019). We make two contributions. First, while it is widely recognized that the intent-to-treat effect (ITT) is attenuated relative to the local average treatment effect (LATE) among compliers, which is also the case here, the LATE is often thought to be larger than the unobserved average treatment effect (ATE)—because, for example, people who choose to take a given drug or to participate in training are those most likely to benefit from it.[1] Here we note that selection likely runs in the other direction in participatory interventions: compliers

1. By the "unobserved ATE," we mean the true (not estimated) difference between treatment and control outcomes under perfect compliance (i.e., in a world in which everyone assigned to treatment were treated and everyone assigned to control were untreated).

typically have the smallest positive individual treatment effects on beliefs about the police—or even negative individual treatment effects—rather than the largest positive ones. Participatory interventions are therefore an interesting class of treatments in which we suppose the LATE to be smaller than the ATE. Second, much of the methods literature views selection as a problem of inference: how to learn about the population average effect of a drug if the trial administers it only to volunteers or how to learn about the population average effect of neutral media coverage if the experiment shows it to partisans who would not otherwise watch. We instead view selection as a problem of policy (as in Mummolo and Peterson [2017] on voter guides). Our objective is not to estimate the effect of police–community meetings on trust in a counterfactual world of mandatory attendance. Rather, we estimate the extent to which actual selection into meetings undermines the policymakers' objective: to build trust between police and the policed and, thereby, to improve outcomes for the whole community.

THREE DECADES OF RESEARCH ON COMMUNITY POLICING

Twenty-five years after "community policing" became a watchword in policy and academic circles (fig. A1), a vast literature has investigated its successes and failures. Google Scholar indexed more than 75,000 academic publications mentioning "community policing" between 1970 and 2022; in a recent systematic review, Blair et al. (2021) identify 276 studies that evaluate community policing interventions, of which 26 focus specifically on town hall meetings (as in this study). In Latin America, the rise of community policing in the 1990s coincided with democratization in some countries and participation-oriented constitutional reform in others—both of which seemed to jibe with the community policing ethos (Ungar and Salomón 2012).

The literature on community policing documents tremendous variation in outcomes. At best, community policing has reshaped police–community relations for the better. In Chicago, for example, the police held more than 45,000 beat meetings during the 15-year heyday of the community policing program (1995–2010), approximately one per month per police beat (Skogan 2022). More than half a million Chicagoans attended. Confidence in the police improved. At worst, community policing "reinforces existing schisms and inequalities" (Brogden and Nijhar 2013, quoted in Müller 2010, 23). In Brazil, for example, González and Mayka (2023) find that police–community meetings amplify the voices of citizens who favor repressive policing. And at median, perhaps, community policing has proven "purely cosmetic" (González 2016): an understaffed office that exists "so that mayors and chiefs have something to point to when questioned" (Skogan 2022) or as "an effort of symbolic policing" (Müller 2010).

Beyond documenting variation, researchers have worked to explain it. The factor most often cited as a determinant of the success of community policing initiatives is commitment: where politicians, bureaucrats, and police leadership embrace community policing; where they secure adequate funding for officer training and foot patrols and public meetings; where they institutionalize community policing so as to shield it from changing political winds; and where they work with community leaders to promote citizen buy-in, community policing is more likely to succeed. Research on community policing in Latin America underscores this point (Arias and Ungar 2009).

Even within initiatives backed by political and fiscal commitment, the literature documents influential variation in institutional design. In the 1980s and 1990s, as González (2016) documents, Bogotá, São Paulo, and Buenos Aires created mechanisms for nonenforcement contact between police and the community. But while the role of Bogotanos was, in practice, limited to sounding a physical alarm, Paulistas and especially Porteños had more influence: in Buenos Aires, participants' written evaluations of the police were even used as a criterion in promotion decisions. Similarly, Vásquez (2012, 52) observes that, though the Colombian National Police instituted many features of a community-oriented policing model—such as organizing patrol into small territorial beats, assigning officers to the same beats for more than a year, and investing in foot patrol—patrol officers lack the discretion required for a true "Chicago-style community police force," in which beat cops decide which activities to prioritize. Below, we provide a theory of change in which police–community meetings would build trust even under this institutional constraint.

A third factor thought to affect community policing outcomes is the crime rate. Skogan (2022, 389) shows that attendance at beat meetings in Chicago fell when crime declined, likely because fear of crime was one of the reasons that people participated in the first place. Ungar and Arias (2012, 414), among others, note that fear of reprisals from gangs can deter participation—which might imply lower attendance rates and inefficacious community policing initiatives in higher-crime contexts. A fourth factor is technology. In Chicago, for example, beat meetings originally served as a venue for connecting participants with other municipal services. But the advent of 311 in the 2000s all but obviated the need for this intermediation (Skogan 2022); Chicagoans could call 311 directly, dampening interest in beat meetings.

Scholars focused on all four of these factors—commitment, institutional design, crime, and technology—would likely be

optimistic about Medellín's police–community meetings. Police supervisors were committed to the intervention in principle; in practice, as we describe in more detail below, implementation was largely successful. The meetings were inclusive in the sense that all community members were invited to attend; for González (2016), inclusivity helps build trust. The homicide rate in Medellín had declined, but security remained a major concern for citizens; moreover, the police refrained from holding meetings in the neighborhoods where armed groups were most likely to pose a threat of reprisal. As for technology, Colombia did introduce a new crime-reporting smartphone app shortly before the start of the police–community meetings; to the extent that it served as a substitute for direct engagement with the police, it may have depressed demand for beat meetings. Ex post, however, we observe that attendance was not likely the principal barrier to the success of the intervention.

The branch of community policing literature that is most relevant to our study focuses on correlates of individual participation in participatory security. Many scholars have observed that participants tend to be older than the general population and that they are more likely to be women. Police–community meetings often attract "sectors of a community pre-disposed to support the police in general" or "the most conservative elements of society," "who often favor harsh crackdowns" (Ungar and Arias 2012). González and Mayka (2023), studying police–community meetings in São Paulo, conclude that the incidence of "calls for repression" from attendees is alarmingly high.[2]

This work might predict that, in Medellín, the most privileged and/or prorepression citizens could coopt police–community meetings, using them as an opportunity to advocate for violent policing. We instead investigate a more innocuous—and perhaps more pervasive—consequence of self-selection into police–community forums: the attendance of relatively pro-police members of the community leaves officers to preach to the choir, weakening any trust-building benefits of the meetings. Even in the absence of harmful calls for repression, then, selection can undermine community policing initiatives. We develop this idea in the following section.

HOW PARTICIPATION-BASED INTERVENTIONS CHANGE BELIEFS

In theory, nonenforcement contact between police and the policed could set in motion a virtuous cycle. Friendly, voluntary conversations with officers might lead people to trust their local officers, thereby encouraging crime reporting, which could then facilitate enforcement and strengthen security outcomes, fostering yet more community trust and cooperation (see fig. 1). Similarly, nonenforcement contact could boost officers' opinions of the community and reorder policing priorities, ideally improving police services and (perhaps) security outcomes—which would then further cultivate rosy views all around, continuing the virtuous cycle. Other common features of community policing intervene at other points in the cycle. Instituting (or expanding) foot patrol, for example, may directly change beliefs and directly reshape policing practices. Internal reforms or directives, in contrast, such as eliminating arrest or citation quotas, affect police behavior without directly reshaping community opinions or officers' beliefs about the policed.[3]

The community policing literature reviewed in the previous section considers many outcomes: crime reporting, police activities, and crime rates. We focus on the first outcome that Medellín's police–community meetings were designed to affect: beliefs.[4] Police–community meetings alone are unlikely to have salutary effects unless they reshape the beliefs of community members and/or police officers.[5] We therefore focus on the conditions under which police–community meetings change people's beliefs about the police: beliefs are a necessary (if insufficient) part of the potential virtuous cycle, and the part most proximate to our intervention. In this framework, there are two ways in which police–community meetings could build trust in police: one in which the meetings allay mistaken distrust toward officers and another in which the meetings themselves constitute a police service that people inherently value.

Police–community meetings can build trust by correcting biased beliefs. Suppose that a population begins with unduly negative views of police, which is to say, misplaced or mistaken

2. Applying their coding rules to a sample of our meeting minutes from Medellín, we estimate a much lower rate of "demands for police repression," suggesting that a related pattern of selection into police–community meetings has different political consequences (app. A9).

3. Of course, public news of internal reforms could directly shape community members' views of police. But, as Ba and Rivera (2019) establish, such changes often take place under the radar.

4. Another reason for this focus is statistical power. We prespecified our expectation that the intervention was not likely to measurably change the security environment within the time frame of an evaluation (one year). Following our preanalysis plan, we do not consider crime outcomes.

5. In principle, behavior change could occur in the absence of a change in beliefs: it could stem purely from revised incentives, for example (see González [2016] on Buenos Aires for an example). But our intervention did not meaningfully affect officer incentives. Moreover, given the relatively limited discretion of patrol officers in Colombia (Vásquez 2012), we expect any citizen learning about the police to occur in the meetings themselves, rather than as a result of changed officer behavior outside of the meetings.

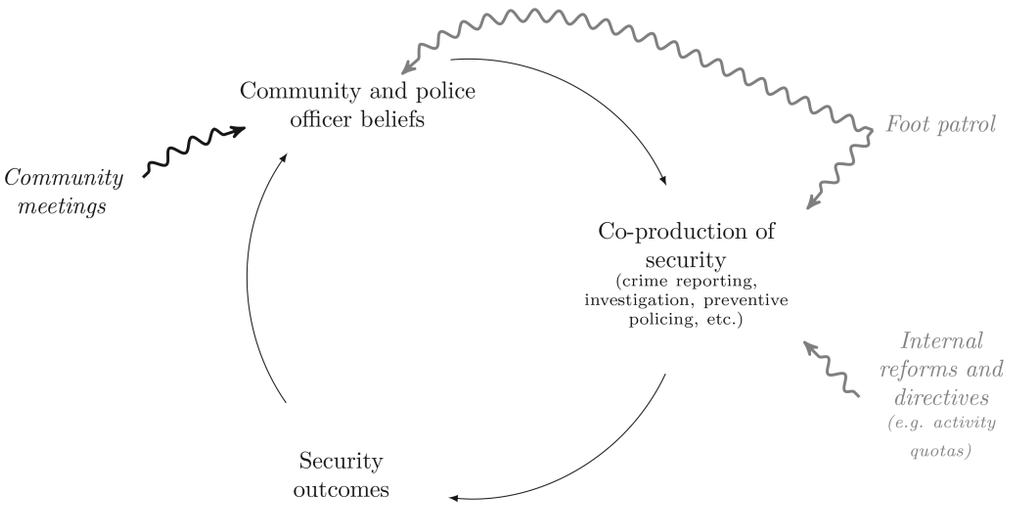

Figure 1. Beliefs: one link in a virtuous cycle. Police–community meetings can work (or not) through changing the beliefs of members of the community and/or of police officers. Other common features of community policing, shown in gray because they are not part of the policy studied in this article, intervene at other points in the cycle.

distrust. This might occur if, for example, policing improved quickly while beliefs lagged behind or if the press were to overreport police misconduct and underreport instances of good policing. In these scenarios and others, police–community meetings might provide an opportunity for people to learn that officers are more trustworthy than previously thought. We view trust as cognitive (as in Bhattacharya, Devinney, and Pillutla [1998] and Hardin [2002]) and posit that beliefs about the the police are formed through Bayesian updating. Updating unduly negative beliefs struck us as a plausible outcome in the context of Medellín, where the police had recently implemented sweeping changes (González 2019). Indeed, it was this intuition that motivated the design of the intervention.

Police–community meetings might also (or instead) build trust simply by providing a service that people inherently value. If people were to desire more opportunities for nonenforcement interaction with their local officers—interaction for its own sake, whether to air grievances or to express appreciation (i.e., to be heard or be seen)—then police–community meetings could directly make police more worthy of trust. In other words, police–community meetings might hold consumption value for invitees and attendees.

Paradoxically, the fact that police–community meetings can provide consumption value may undermine their ability to correct biased beliefs. This could occur if, as stands to reason, those who most trust the police might especially appreciate the opportunity to spend leisure time in the company of officers.[6] If prior beliefs determine individual-specific consumption values, and if consumption values drive attendance, then we would expect those who trust the police at baseline to be more likely to attend than those who do not.

It is this correlation—between prior beliefs and attendance—that can weaken the corrective effects of police–community meetings on biased prior beliefs. No matter what the signal provided by the meetings, attendees with the most positive prior beliefs are more likely to revise their beliefs downward (at best, they can maintain their maximally positive prior); attendees with the most negative prior beliefs are the most open to revising upward. The stronger the positive correlation between baseline beliefs and attendance, the smaller the effects of the intervention. This is the problem of preaching to the choir.

These dynamics generalize to other participatory interventions that seek to change citizens' beliefs about the state. Consider, for example, 311-type telephone reporting systems. Whether or not 311 lines improve municipal services, they might build trust in government through one or both of the mechanisms outlined above: providing consumption value (the value of voice), and/or correcting negatively biased beliefs about government responsiveness. But selection into sharing information via 311—like that documented by Slough

6. The correlation could run the other way: those most distrustful of the police might especially appreciate meetings, as a forum for lodging complaints.

(2021a)—could exaggerate or attenuate any attendant changes in citizen beliefs. The preaching-to-the-choir problem is not specific to policing.

CONTEXT

The puzzle of police–community relations in Colombia is that trust in police deteriorated even as citizen security improved and police killings declined. Between 1998 and 2018, Colombia's homicide rate plummeted from 75 to 35 per 100,000 (fig. 2, left panel); the improvement in security was yet more dramatic in Antioquia (of which Medellín is the capital), where the homicide rate dropped nearly 80% since the early 2000s.[7] And while we lack a reliable multidecade time series on police use of lethal force, repeated snapshots suggest a significant decline. As recently as the early 1990s, the Colombian police committed hundreds if not thousands of extrajudicial killings every year (Amnesty International 1994); in 2015, 2016, and 2017, there were just 85, 73, and 72 known victims of police killings in Colombia, respectively (Correa, Forné, and Cano 2019, 62). By any measure—absolute number of victims, victims per 100,000 population, victims as a fraction of all homicides, or victims relative to number of police killed in the line of duty—these numbers imply that the Colombian police are considerably less lethal than their counterparts in Mexico, Venezuela, Brazil, El Salvador, or even the United States (Correa et al. 2019; Hanson and Zubillaga 2021; Osse and Cano 2017). Moreover, especially since 2010, the Colombian police have embraced many aspects of community policing, with salutary effects on the quality of police services (García, Mejia, and Ortega 2013).

Yet confidence in the police has not improved. The proportion of survey respondents expressing confidence in the police[8] declined from 56% (in 2004) to 42% (in 2018), according to survey data from the Latin American Public Opinion Project (LAPOP; fig. 2, right panel).

One possible explanation is that neither citizen security nor officers' abuse of force figure heavily in people's estimation of the police. Perhaps people evaluate the police primarily based on outcomes other than crime rates and/or physical abuse, for example, petty corruption. If officers' bribe-seeking increased over the time period shown in figure 2, trust in the police could have fallen as a result. Another possible explanation is that people paradoxically value *mano dura*, as several studies have argued (Holland 2013). Were that the case, the decline in police use of lethal force might actually drive more negative assessments of the police.

A more likely explanation, in our view, is that beliefs are sticky: a few years of reform—even a half-generation of reform—might be no match for the impressions created by generations of violence and repression. Recent, high-profile incidents of human rights violations committed by the police "riot-control" unit and, at a much larger scale, by the Colombian military (Acemoglu et al. 2020) might contribute to the persistence of mistrust. Were it indeed the case that beliefs have lagged behind the reality of reform in the Colombian police, police–community meetings might provide a valuable opportunity for people to learn about their local officers.[9]

This conclusion echoes the expectations of our counterparts in the Medellín police. The Metropolitan Police of the Aburrá Valley is a metropolitan division of the National Police of Colombia. At our first meetings with them—together with officials in the Security Secretariat of the Medellín mayor's office, which has some jurisdiction over local policing—they were keenly aware of low levels of trust in police, as expressed in recent local surveys. They had already held several public meetings in which high-level police supervisors would meet with residents from large swaths of the city, but they were interested in (as they put it) "going micro": refocusing municipal and police attention on problems specific to small neighborhoods. Neighborhood-specific town hall–style meetings between officers and residents fit this vision. Our counterparts' hope and expectation was that these meetings would improve perceptions of the police by providing new information about local patrol officers and about the institution as a whole.

RESEARCH DESIGN AND DATA

Treatment assignment

To study the effect of police–community meetings on attitudes toward the police, we randomly assigned police beats in Medellín to meetings (treatment) or no meetings (control), and we measured attitudes using baseline and endline surveys of residents.[10] We included 347 police beats in our sample, excluding the remaining 66 because of insufficient residential populations (e.g., the airport); our sample contains approximately 96% of Medellín residents. Because police beats are fairly large—the median population was 5,348 in the 2005

7. The drop in the homicide rate was even sharper in Medellín: from a record-setting 350 per 100,000 to approximately 35 per 100,000.

8. Specifically, choosing 5, 6, or 7 on a 7-point scale in response to the question, "How much do you trust the police? Use any number between 1 (not at all) and 7 (very much)."

9. Indeed, contrary to work that views misperceptions as resistant to new information, Esberg and Mummolo (2018, 3) conclude that "citizens would hold more accurate beliefs if they encountered relevant information [about crime rates], but common news reporting practices . . . may undermine the uptake of facts."

10. In a second treatment arm, analyzed separately, we distributed informational flyers; this treatment was cross-randomized with police–community meetings.

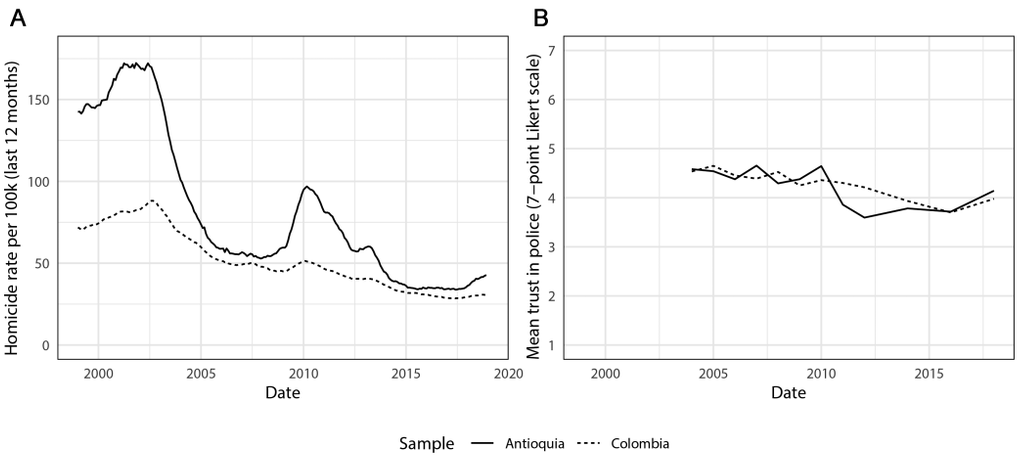

Figure 2. Citizen security dramatically improves, trust in police does not. The left panel shows the number of homicides per 100,000 per year in Colombia (dashed line) and Antioquia (of which Medellín is the capital, solid line), using vital statistics data published by the National Administrative Department for Statistics (DANE). The right panel shows the proportion of survey respondents expressing trust in police according to surveys from the Latin American Public Opinion Project.

census (IQR: [2,734, 8,339])—we defined "prioritized neighborhoods" within each beat. Specifically, we distributed meeting invitations and conducted our surveys in the set of inhabited, contiguous city blocks closest to the centroid of each beat (fig. A3). These prioritized neighborhoods each contained approximately 1,200 residents, or 400 households.

To ensure balance, we randomized within blocks of police beats. Each block contained four police beats that (a) belong to one of 11 police-station groups[11] and (b) share the same treatment status in a simultaneous intervention conducted by other researchers (Blattman et al. 2022). Within each block of four beats, we randomly assign two beats to the police–community meetings treatment and two to control. As a result, our treated and control units are balanced on population, household structure, and various measures of socioeconomic status (table A1).

Intervention

In neighborhoods assigned to treatment, the study team and the police held three police–community meetings over a period of approximately nine months. The research budget paid for research-assistant time, project-coordinator time, surveys, and meeting invitations but did not fund police-officer time or provide any resources to the police. To advertise the meetings, a survey firm distributed 350 invitation flyers (fig. A4) door-to-door among the approximately 400 households within each treated neighborhood. The study team also hung a poster at the meeting location, advertising the meeting to passersby, and worked with community leaders to spread the word via WhatsApp. Meetings generally took place in the late afternoon or early evening and ran anywhere between one to three hours, at the end of which the study team distributed light refreshments. We describe the content of the meetings in more detail below.

Ethics

A field experiment involving the police entails special ethical considerations. For one, creating opportunities for nonenforcement contact between officers and citizens implies creating opportunities for contentious interactions or even conflict. This risk was smaller in our context than in many others; as noted above, despite high homicide rates, Colombia has relatively low levels of police violence. Consultation with local researchers and nongovernmental organizations confirmed this impression. We also implemented safeguards to monitor and mitigate the risk of harmful conflict in the meetings, including, but not limited to, (1) involving the local elected community council (Junta de Acción Comunal [JAC]) in the process of planning, publicizing, and implementing meetings and (2) reviewing and responding to detailed meeting reports from our team. For another, traveling to local community centers to attend meetings potentially posed risks to attendees, to our study team, and to police officers themselves. We mitigated this risk by following guidance from police chiefs and from our study team about when and where to cancel meetings (we discuss cancellations in more detail below) and by choosing meeting locations near residents' homes and within police officers' beats. We address additional ethical considerations in appendix A3.

11. There are 14 police stations, but we grouped six small stations into three groups.

Data

We measure residents' attitudes toward the police using a panel survey. At baseline, between January and April 2018, a local survey firm conducted an in-person survey of 5,205 residents of the prioritized neighborhoods (15 residents per neighborhood). At endline, between September and December 2019, the firm was able to recontact 2,434 of these original interviewees and survey 1,210 new interviewees in treatment and control neighborhoods. There is no evidence that attrition was related to treatment assignment (table A2, fig. A5). Our survey instrument was written in concert with the other studies in Blair et al. (2021). Table A5 describes how we operationalize our principal outcomes.

We also collected two types of qualitative data. First, we conducted one focus group, five interviews, and preliminary fieldwork prior to the start of the intervention. Second, we collected detailed notes on each of the 519 police–community meetings. One of the coauthors trained research assistants to take short-form field notes at the meetings and write long-form field notes after the meetings. The objective of this documentation was to characterize interactions between officers and citizens, list topics of discussion, and capture certain illustrative conversations. Figure A2 summarizes the study timeline.

Estimation

Our primary prespecified estimand is the ITT: the effect of assignment to police–community meetings on attitudes toward the police. We estimate the ITT effect using

$$Y_{ijb}^{\text{Post}} = \beta Z_{jb} + \boldsymbol{\gamma}_b + \epsilon_{ijb}, \tag{1}$$

where Y_{ijb}^{Post} is the endline survey response of person i in police beat j in block b, Z_{jb} is an indicator for assignment to the police–community meetings treatment, and γ_b is a vector of block fixed effects. Some specifications also include baseline measures of the survey response Y_{ijb}^{Pre}.

In addition to this prespecified analysis, we develop a stylized model in order to quantify the consequences of positive citizen selection into meetings. We noted above that positive selection into meetings attenuates the actual effect of the intervention (i.e., the estimand); it also attenuates our estimates of both the ITT and the ATE on the treated (ATT) for two reasons. The first is censoring. If enough people express prior beliefs at the maximum of the measurement scale, then we cannot observe improvements in trust. The second is mean reversion. In the presence of mean reversion—negative changes in beliefs for those with the most positive priors, positive changes for those with the most negative priors—positive selection into meetings will mechanically attenuate our estimate of the ATT. Censoring and mean reversion are problems of measurement; in that sense, they pose challenges for the researcher but not for the policymaker.

Part of our discussion considers the ATT, for which we use only data from treated neighborhoods to estimate

$$(Y_{ijb}^{\text{Post}} - Y_{ijb}^{\text{Pre}}) = \beta A_{ijb} + \epsilon_{ijb}, \tag{2}$$

where Y_{ijb}^{Post} is the endline survey response of person i in police beat j in block b, Y_{ijb}^{Pre} is the response for that same individual at baseline, and A_{ijb} is an indicator for (self-selected) participation in the police–community meetings.[12] Some specifications also include Y_{ijb}^{Pre} on the right-hand side, for reasons that we discuss below.

RESULTS

Using ITT estimates from (prespecified) equation 1, we find that the police–community meetings did not affect trust in the police as an institution, trust in beat officers, perceptions of police quality, perceptions of security, or perceptions of community relations (*convivencia*). We report these null findings in figure 3 and in tables A6–A8. The point estimates are substantively small—generally less than 0.05 standard deviation and always less than 0.1 standard deviation—and statistically indistinguishable from zero. We consider three possible explanations for these null findings: imperfect reach (or low compliance), varied meeting quality, and positive citizen selection into meetings.

Imperfect reach

Imperfect reach is perhaps the most obvious potential explanation for our null results. If the police–community meetings reached only a small fraction of the population, we would not expect them to change attitudes or beliefs in the population at large. In fact, in our case, the meetings reached a nontrivial fraction of residents of treated neighborhoods. We can see this, first, simply by dividing the total number of attendees (as counted by our study team at each meeting) by the population of all treated neighborhoods; using this measure, we find that 8% of the adult population attended a meeting. This figure is remarkably close to the estimate we obtain from our panel survey: at endline, 8.2% of survey respondents in treated neighborhoods reported having attended a police–community meeting.

To compare meeting exposure in treated neighborhoods to meeting exposure in control neighborhoods, we use our panel survey. The survey asks (1) whether the respondent has heard

12. Results are robust to the inclusion of block fixed effects, γ_b. We exclude these fixed effects from the primary ATT specifications because there are some blocks without variation in A_{ijb}. By omitting the fixed effects, we compare the same effective sample across specifications.

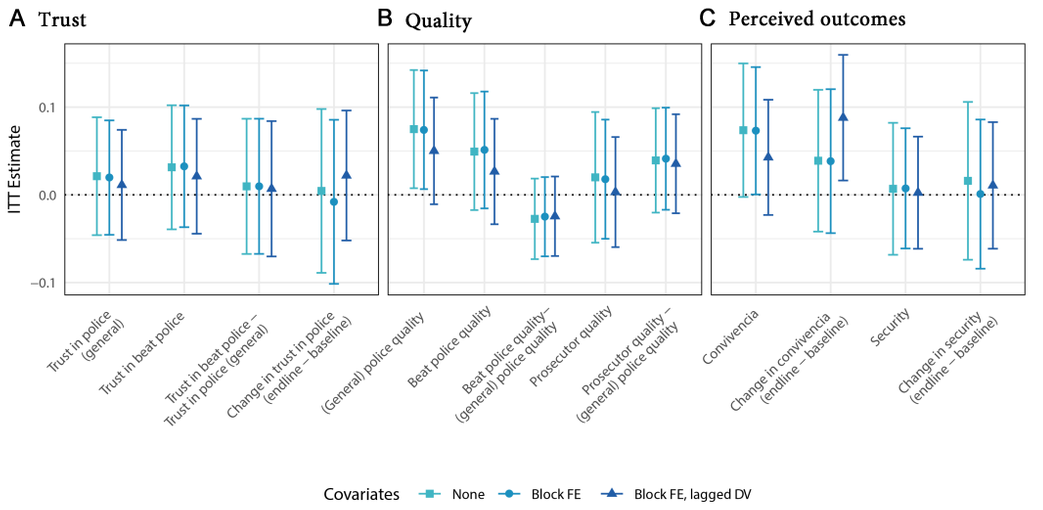

Figure 3. Intent-to-treat effects are small and indistinguishable from zero. Estimates of (prespecified) equation 1 with 95% confidence intervals (see also tables A6–A8).

of "meetings or activities" organized by the police in their neighborhoods in the past year and (2) how many such events the respondent attended. We code both responses as binary variables capturing (1) any exposure to police–community meetings and (2) any participation in police–community meetings, respectively. We then estimate versions of equation 1 in which the dependent variable is one of these two measures of compliance (exposure or attendance). Table 1 reports the results. Columns 1–4 reveal that, as expected, assignment to treatment is uncorrelated with exposure to or attendance at police–community meetings at baseline. Columns 5–8 show that the intervention was indeed implemented in a meaningful way: at endline, residents of treated neighborhoods were nearly twice as likely (as residents of control neighborhoods) to report hearing about or attending police–community meetings. Table 1 and figure A6 also suggest that there were spillovers in exposure to police–community meetings but not in attendance.

Varied meeting quality

A second source of attenuation in the effect of the intervention was heterogeneity in the signal provided by the meetings

Table 1. The Intervention Meaningfully Increased Exposure to Police–Community Meetings

	Baseline Exposure		Baseline Attendance		Endline Exposure		Endline Attendance	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Meetings	−.003	.005	−.003	.003	.163***	.216***	.041***	.052***
	(.007)	(.010)	(.007)	(.010)	(.018)	(.022)	(.009)	(.011)
Block FE	✓	✓	✓	✓	✓	✓	✓	✓
Observations	5,205	2,434	5,205	2,434	3,644	2,434	3,644	2,434
Clusters	347	347	347	347	347	347	347	347
Control Mean	.054	.054	.051	.052	.264	.291	.059	.063
Sample	Baseline	Panel	Baseline	Panel	Endline	Panel	Endline	Panel

Note. Estimates of equation 1. Columns 1–4 confirm balance. Columns 5–8 reveal that the intervention substantively and measurably affected exposure to and attendance at police–community meetings in treated neighborhoods. Standard errors clustered at the level of police beat. FE, fixed effects.
* $p < .05$.
** $p < .01$.
*** $p < .001$.

themselves. In other words, police officers looked good in some meetings and terrible in others. We use the qualitative notes described above to characterize variation in the tone of the police–community meetings in two ways: first, by recounting illustrative interactions from several specific meetings and, second, by using the National Research Council emotion lexicon (in Spanish) and the improved Spanish Opinion Lexicon (Molina-González et al. 2014) to create a quantitative measure of meeting sentiment, which we validate by comparing it to hand-coded sentiment values for a small sample of meetings. Both approaches suggest that meeting sentiment is not only highly variable but also, contrary to our expectations, highly idiosyncratic.

Consider first selected examples of variation in meeting quality and sentiment. One of the more positive meetings literally ended in a group hug. On a Tuesday afternoon in April 2019, 14 residents gathered in a community center to meet the two patrol officers assigned to their beat. The officers inspired confidence early in the meeting by distributing stickers with their phone numbers and by asking attendees to call them then and there, in order to check that residents had saved the officers' numbers correctly and to check that the phones were working. "When the officers' phones rang, a number of women were very pleased and said that they thought that the officers would pick up" when residents called with requests for service in the future, wrote the research assistant who observed the meeting. After that auspicious beginning, one of the more vocal residents began sharing tips about possible criminal activity in the neighborhood; she pointed officers to an internet café that seemed to operate at all hours, information that the officers wrote down. Toward the end of the meeting, after several additional productive interactions, the project coordinator asked attendees if they had heard about Colombia's "hug a police officer" campaign. "Before I had finished explaining it," she wrote, the attendees were coming up to hug the officers "and give them blessings and kisses." One of the officers thanked the project coordinator on his way out, saying how happy it made him to feel appreciated and noting that he would never forget it.

At the other end of the spectrum, in August of 2018, leaders of a local elected community council (JAC) berated the officers, accusing them over and over of corruption. The president of the JAC accused the officers of looking the other way when local businesses violated noise ordinances, "as long as they pay you off." One of the two officers responded by invoking a bad-apple narrative, noting that "that is not how all police officers behave," but the council president was not buying it. She was precisely the type of informed, connected, and powerful local citizen with whom the patrol officers would have liked to establish a good relationship, but nothing they said seemed to sway her. At one point, she even accused the officers of behaving like "a legalized armed gang." After the end of the meeting, the council president approached the representative of the research team and warned them about the patrol officers. "We'll see how things go, but I just don't have faith in them—especially *ese morenito*," she said. Figure 4 summarizes these interactions, succinctly illustrating the process by which police–community meetings might have changed beliefs and showing variation in the direction of change.

Our quantitative measures of meeting sentiment reveal not only that there was substantial heterogeneity but also that sentiment is uncorrelated with neighborhood-average trust in police (fig. A14). In other words, it is not the case that the most positive meetings generally took place in pro-police neighborhoods. Rather, the tone of the meetings was highly idiosyncratic. This heterogeneity in the signal from the meetings likely attenuated their effects.

Positive selection into police–community meetings[13]

Residents of neighborhoods assigned to police–community meetings were invited to attend the meetings. But residents themselves decided whether to attend. And trust in police is a very strong predictor of attendance, both in Colombia and in three of the five other countries that implemented coordinated experiments.

For Colombia, the left panel of figure 5 plots the rate at which respondents report hearing about police–community meetings at endline, as a function of their trust in police at baseline. While just 26% of respondents who least trust the police report hearing about the meetings, 43% of those who most trust the police hear about the meetings. Similarly, just 5% of those who least trust the police report attending meetings, compared to 12% of those who most trust the police (right panel of fig. 5). Nor are these differences driven by small numbers of people at the extremes of the distribution of trust in police (fig. A7).

In three of the other five sites in the community-policing meta-study—Liberia, Uganda, and the Philippines—we observe positive and substantively large differences in attendance

13. Positive selection attenuates the treatment effects themselves, as we discuss in this section. Censoring and mean reversion, in contrast, may attenuate our estimates of treatment effects. Censoring and mean reversion are issues of measurement: the former arises because of the scale of the survey question, and the latter because of idiosyncratic measurement error in survey responses. In that sense, censoring and mean reversion pose challenges to the researcher but not to the policymaker. We investigate censoring and mean reversion in apps. A9.1 and A9.2.

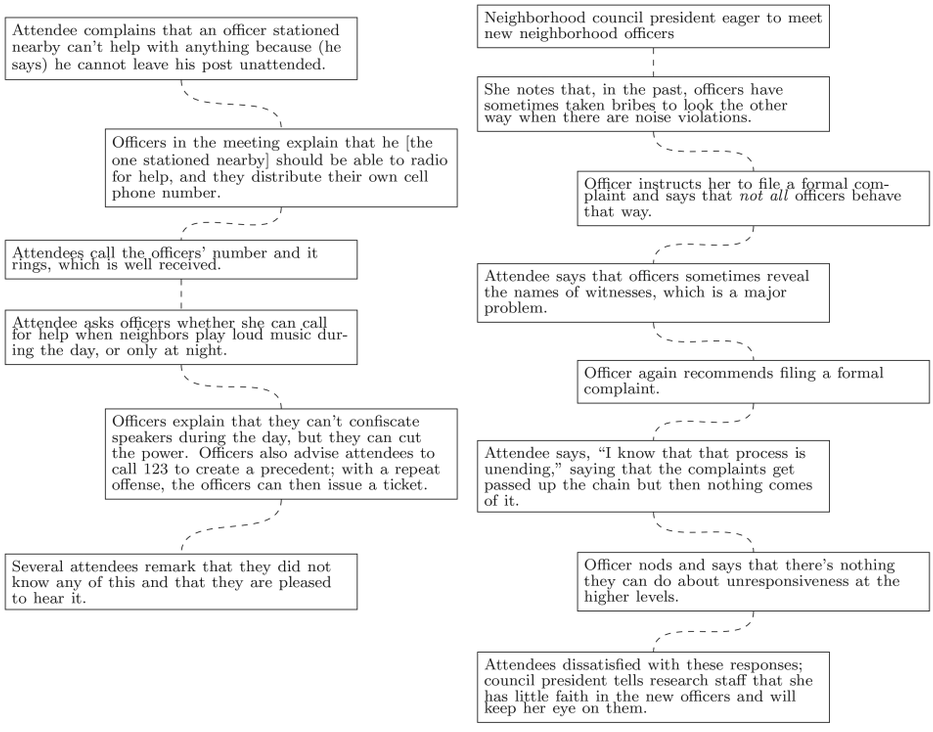

Figure 4. Positive and negative updating in police–community meetings. This figure illustrates how one meeting led to attendees improving their opinions of the police (left), while another resulted in negative updating (right): attendees left more distrustful of the police than they had been when they arrived.

rates between people with the lowest and highest levels of trust in police. In Liberia, average meeting attendance was 12 percentage points (47%) higher in treated communities with the highest levels of baseline trust than in treated communities with the lowest levels of baseline trust (38% vs. 26%), though the difference is not precisely estimated (table 2).[14] In Uganda, baseline trust was uncorrelated with attendance. But endline trust—which was measured closer in time to the meetings and which (per the ITT estimates) was unaffected by the treatment—reveals a similar pattern: in treated neighborhoods, 39% of respondents most trustful of the police attended meetings, compared to 29% of the least-trusting respondents. In Uganda, as in Colombia, attendance rates increase monotonically across levels of trust in police. In the Philippines, there was no baseline survey; at endline, 18% of the highest-trust respondents reported attending meetings, compared to 10% of the lowest-trust respondents, though the difference is not precisely estimated because there are so few respondents in the lowest-trust category.

In the two remaining Metaketa sites, Pakistan and Brazil, in contrast, we do not observe a positive correlation between baseline trust and attendance at police–community meetings. In Pakistan, attendance was nearly flat across baseline and endline trust categories. One possible explanation is that trust in police is simply more fluid in this context, given that the correlation between baseline and endline beliefs is zero (compared to 0.48 in Colombia or 0.38 in Uganda); if we were to interpret this fluidity as evidence that, in Pakistan, this particular survey question captures an ephemeral feeling rather than an enduring attitude, then it would be unsurprising that the responses are uncorrelated with behavior. In Brazil, we see no evidence of positive selection into meetings, though the low recontact rate makes it difficult to assess. Of the 493 recontacted respondents in treated units, just 24 reported the lowest level of trust at baseline, of whom four (16%) attended, compared to 19 of the 278 respondents (7%) with the highest level of trust.[15] In other words, the low recontact rate in Brazil means that our estimate of the trust–attendance gradient is noisy. To summarize, we observe positive selection into meetings in Colombia, Liberia, Uganda, and the Philippines, though not in Pakistan (where beliefs appear less sticky) or Brazil (where there is less data).

14. In Liberia, the survey was a repeated cross-section rather than a panel, meaning that we cannot study selection at the individual level.

15. If we use more of the data from Brazil by comparing respondents who report the top two trust categories to those who report the bottom two, the difference shrinks to −1.9 percentage points.

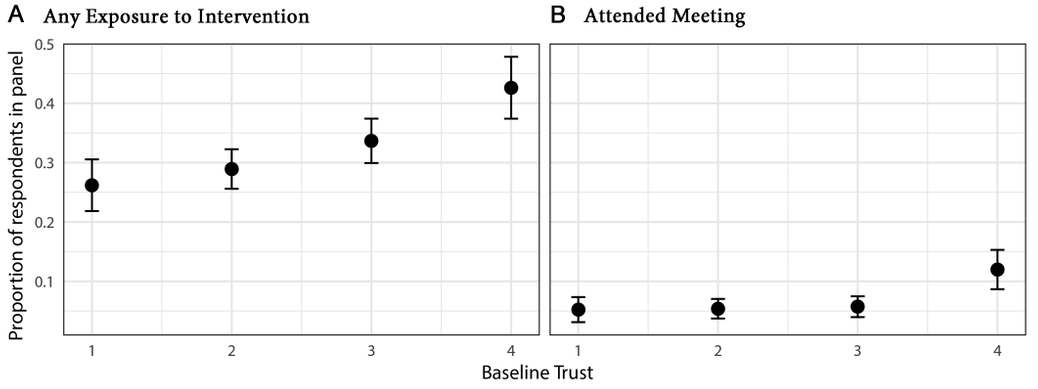

Figure 5. Baseline trust in police strongly predicts engagement. This figure plots the probability of hearing about police–community meetings (left) and attending meetings (right) as a function of baseline trust in police (increasing along the *x* axis).

This relationship between prior trust in police and selection into police–community meetings is not merely one of many correlations discovered ex post in our data. Using a Lasso model including 10 categories of predictors (treatment assignment, age, gender, and various measures of political attitudes and socioeconomic status), we find that baseline trust in police is second only to treatment assignment in predicting whether respondents hear about meetings (app. A7.3). In predicting whether respondents attend meetings, baseline ranks third, just behind treatment assignment and participation in the 2016 plebiscite (another form of engagement with the state). Moreover, controlling for a series of variables that could be correlated with trust in the police—age, education, political affiliation, migrant status, etc.—does not appreciably change the relationship between baseline trust and meeting participation (app. A7.3). This result suggests that our estimates of the relationship between participation and trust are not merely capturing associations between participation and some other characteristic.[16]

One way to get a sense of the magnitude of the preaching-to-the-choir problem is to consider the ATT (equation 2). Columns 2, 3, 5, and 6 of table 3 (panel A) reveal that, defining "attendees" as the endogenously treated population, the ATT is negative and on the order of 0.15 standard deviations. In other words, endline trust in police declined by 0.143 points (on a 1–4 scale) among people who did not attend meetings, and it declined more—0.332 points—among people who did attend meetings (using results from column 2). Were it not for positive selection, we might interpret this result as suggestive evidence that the meetings made people more distrustful of the police, or caused them to update negatively. But panel B of table 3 suggests otherwise. When we control for baseline beliefs—which is to say, when we consider the fact that attendees were much more likely to trust the police to begin with—then the (modified) ATT estimates are generally positive, if imprecisely estimated. Given where they started, the endline beliefs of people who heard about and/or attended meetings declined less than those of people who neither heard about nor attended meetings. The differences in coefficients between panel A (no accounting for positive selection) and panel B (accounting for positive selection) are substantively large and, in some cases, distinguishable from zero. The bottom line is that acknowledging positive selection leads to a strikingly different conclusion about the relationship between meeting attendance and change in attitudes toward the police.

Comparing the influence of reach, signal, and selection

To gauge how these challenges influence the effect of police–community meetings, we develop a stylized model. The model starts with a population whose prior trust in police, π, is drawn from a discrete uniform distribution between 1 and 4 (mimicking the four-point trust in police measure from our survey). We vary both (a) the overall rate of meeting attendance among the half of the population invited to meetings and (b) the extent to which attendance covaries with prior beliefs. For simplicity, we consider only linear relationships between prior beliefs and attendance. We then define positive selection as

$$\theta = \frac{Pr(\text{Attend}|\pi = 4) - Pr(\text{Attend}|\pi = 1)}{Pr(\text{Attend})} = \frac{3\Delta}{Pr(\text{Attend})}. \quad (3)$$

16. A related concern might be that participants do not update not because of their prior (i.e., that they trust the police to begin with) but rather because of some other, correlated characteristic—in particular, exposure to the history of police violence—that makes their beliefs stickier. In fig. A11, we show that, in fact, neither older people nor internal migrants (both of whom likely have more exposure to violence) are less prone to change their beliefs about the police.

Table 2. Meeting Attendance is Positively Correlated with Trust

Country	P(Attend | Lowest Trust)	P(Attend | Highest Trust)	Difference
Liberia	.251	.379	.128
	(.054)	(.061)	(.102)
Uganda*	.286	.386	.100
	(.037)	(.032)	(.052)
Colombia	.084	.184	.100
	(.019)	(.029)	(.035)
Philippines*	.097	.179	.083
	(.054)	(.011)	(.057)
Pakistan	.058	.037	−.021
	(.018)	(.037)	(.042)
Brazil	.167	.068	−.098
	(.084)	(.019)	(.091)

Note. For all countries except Liberia, this table reports average meeting attendance in the highest-trust category and the lowest-trust category, in treated neighborhoods/communities. For Liberia, where we observe baseline trust at the community level (not the individual level), we report predicted attendance values at the minimum and maximum of observed community-level trust.

* In Uganda and the Philippines, we use endline rather than baseline trust. In Uganda, the endline survey was conducted closer in time to the meetings (see additional discussion in the text); in the Philippines, there was no baseline survey.

Table 3. Average Treatment Effect on the Treated among Individuals in Treated Neighborhoods

	Change in Trust in Police			Change in Police Quality Index		
	(1)	(2)	(3)	(4)	(5)	(6)
Panel A						
Any exposure	−.088		−.057	.020		.058
	(.064)		(.071)	(.059)		(.065)
Attended meeting		−.179	−.143		−.141	−.177
		(.110)	(.121)		(.094)	(.102)
Panel B						
Any exposure	.092		.094	.121*		.131*
	(.052)		(.056)	(.051)		(.056)
Attended meeting		.048	−.009		.032	−.049
		(.095)	(.102)		(.074)	(.082)
Baseline as covariate	Yes	Yes	Yes	Yes	Yes	Yes
Any exposure: panel A = panel B, p value	.03		.096	.195		.392
Attended mtg: panel A = panel B, p value		.118	.397		.195	.33
Untreated mean	−.155	−.158	−.155	.025	.046	.025
Untreated standard deviation	1.158	1.156	1.158	.935	.932	.935
Observations	1,191	1,191	1,191	1,206	1,206	1,206
Clusters	173	173	173	173	173	173

* $p < .05$.
** $p < .01$.
*** $p < .001$.

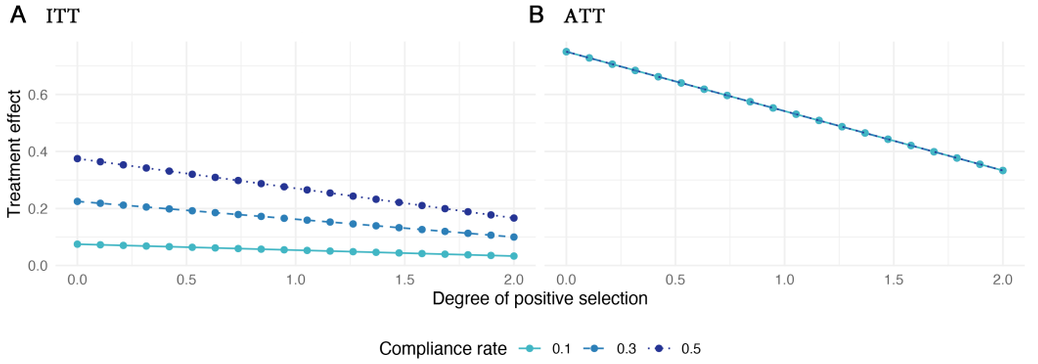

Figure 6. The problem of preaching to the choir. These figures plot how the ITT and the ATT change with the degree of positive selection into meetings under our model of selection. The *x* axes span θ, our measure of the degree of positive selection (equation 3); $\theta = 2$ implies that the difference in attendance rates between the maximum and minimum baseline trust categories is twice the average attendance rate.

In other words, θ captures the difference in attendance rates between those who most trust the police at baseline ($\pi = 4$) and those who least trust the police at baseline ($\pi = 1$), normalized by the overall attendance rate. Note that $\theta \in [0, 2]$: the maximum Δ is 1/3, in which case the attendance rate would increase from 0 (in the $\pi = 1$ group) to 1 (in the $\pi = 4$ group), and the overall attendance rate would be 0.5, making $\theta = 1/0.5 = 2$. We then assume that attendees (and only attendees) are exposed to a signal about police trustworthiness, allow them to update in a Bayesian manner, and calculate the ITT and the ATT.

We begin by considering treatment effects when the meetings provide maximally positive information about the police (i.e., when the signal to attendees is always 4). We choose this setup in order to characterize the role of positive selection even in an environment that is favorable to the intervention, which is to say, an environment in which we would expect the intervention to build trust. Figure 6 presents the results. When the overall attendance rate (compliance) is 50%, the darkest line in the left panel of figure 6 shows that, in the absence of positive selection (when the *x* axis value is 0, meaning that the attendance rate is uncorrelated with prior beliefs) the ITT is quite large: 0.38 on a 1–4 scale, which corresponds to 0.36 control-group standard deviations. But positive selection quickly shrinks this effect: moving from $\theta = 0$ to $\theta = 1$, the actual value observed in our data, the ITT falls 18%, from 0.38 to 0.31. Moving to the maximum possible extent of positive selection, $\theta = 2$, further shrinks the ITT to 0.25, a 43% reduction relative to the ITT with random selection. With lower rates of compliance, such as 30% (middle line) or 10% (light gray line), positive selection similarly attenuates the ITT. Positive selection also quickly attenuates the ATT (right panel of fig. 6).

This exercise emphasizes that positive selection into meetings reduces the estimated treatment effects even when the intervention (a) provides maximally positive information about police and (b) reaches a sizable fraction of the population. The consequences of positive selection are at least as perverse in an environment less favorable to the intervention. If prior beliefs were biased in favor of the police—if, for example, priors were uniformly distributed on [1, 4] and the meetings were to provide a signal lower than 2.5—positive selection would exaggerate the negative effect of treatment. It is, of course, impossible to quantify precisely how positive selection affected our actual results. But our estimated ATTs and stylized model suggest that, while imperfect reach and heterogeneous meeting quality did hamper the intervention, positive selection into police–community meetings also posed a significant challenge.

PREACHING TO THE CHOIR BEYOND POLICING AND BEYOND MEDELLÍN

The dynamic that we observe in our field experiment—that police–community meetings in Medellín attract people who trust the police to begin with—is not specific either to policing or to Medellín. We discover suggestive evidence of similar patterns in 116 surveys from 23 countries across Latin America and the Caribbean.

The LAPOP routinely asks respondents whether they attend city council meetings and whether they attend community-association meetings, as well as how much they trust city government and how much they trust their fellow community members. This allows us to estimate, for each survey, a measure of positive selection analogous to the one defined in equation 3 above: the (linearized) difference between attendance rates in the highest- and lowest-trust categories, normalized by the overall attendance rate.[17] In particular, we estimate

17. Substantive inferences are unchanged when calculating λ using rates of attendance in the highest and lowest trust groups, without linearizing. Since the LAPOP surveys measure trust in community and city

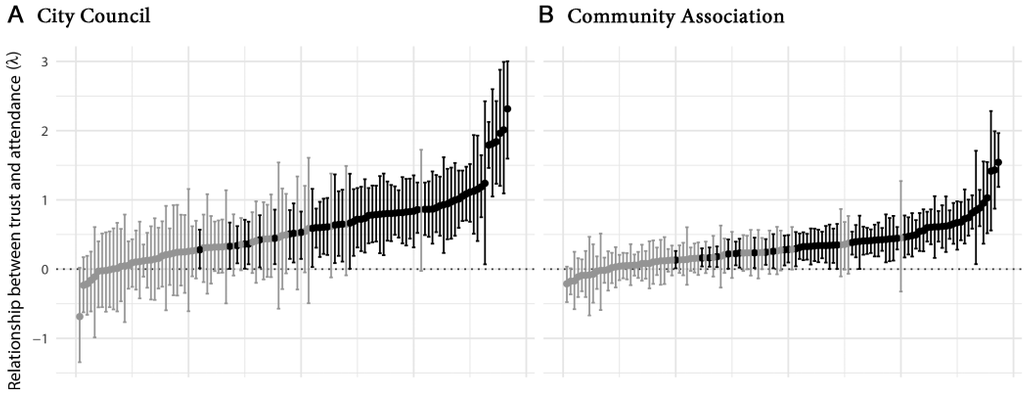

Figure 7. Do public meetings attract more fans than critics? Using data from the Latin American Public Opinion Project, the subfigure on the left plots (for each country–survey wave) a measure of positive selection into city council meetings (λ, equation 4). On the right, we report the same measure for attendance at neighborhood association meetings, given trust in community members. The segments mark 95% confidence intervals constructed by bootstrapping; the estimates in black are distinguishable from zero at $\alpha = 0.05$.

$$\lambda = \frac{Pr(\text{Attended}|\text{Trust} = \text{Highest}) - Pr(\text{Attended}|\text{Trust} = \text{Lowest})}{Pr(\text{Attended})}. \quad (4)$$

Larger positive values of λ indicate a stronger positive relationship between trust and public-meeting attendance.

There is a conceptual difference between the estimates of λ from the LAPOP survey data and the estimates of θ from our field experiment. The LAPOP estimates tell us about the relationship between trust (at the moment of the survey) and attendance at previous public meetings. The estimates from our field experiment tell us about the relationship between trust (at the moment of the survey) and attendance at subsequent public meetings. One might not expect the quantities to be similar. The relationship between trust and previous meeting attendance captures not only selection into meetings but also the effect of previous meetings on trust; if previous meetings built trust in government, the relationship between trust and previous meeting attendance could be strong and positive even in the absence of any positive selection into meetings. For that reason, ex ante, we might not expect the LAPOP estimates to be informative about our quantity of interest (i.e., the extent to which trust predicts meeting attendance). But ex post, we are able to observe in our data from Medellín that, in fact, the quantities are similar. Our baseline survey asks not only about trust but also about attendance at previous public meetings (larger town hall meetings that the police had occasionally convened prior to the intervention). Using our baseline measure of trust in police and the baseline measure of attendance at previous meetings, we estimate positive selection of $\hat{\lambda} = 1.10$ (95% CI: [0.66, 1.34]). This figure is statistically indistinguishable from the degree of selection into the police–community meetings held as part of the intervention (i.e., using baseline trust and endline attendance), which we estimate as $\hat{\lambda} = 1.01$ (95% CI: [0.43, 1.33]). In other words, the relationship between trust and past attendance is indistinguishable from the predictive effect of trust on attendance. If a similar pattern holds for other participatory institutions, our estimates of λ from LAPOP may in fact be informative about positive selection into public meetings.

With this in mind, we estimate the relationship between trust and public-meeting attendance in LAPOP survey data. Figure 7 plots estimates of λ (equation 4) for (a) trust in city government and attendance at town hall meetings and (b) trust in members of the community and attendance at community-association meetings. In 106 of 115 country–survey waves, people who most trust the city government were more likely to attend city council meetings; in 62 country–survey waves, the difference was at least 5 percentage points, a substantial increase over the average overall reported attendance rate of 11%. A similar pattern holds for trust in community members and attendance at the meetings of neighborhood associations. We estimate positive selection into neighborhood association meetings in 104 of 116 country–survey waves; in 78 waves, the increase in attendance from the minimum to the maximum trust category exceeds 5 percentage points, a substantial change relative to the average attendance rate of 28%. Even if only a fraction of these associations stem from positive selection (as opposed to salutary effects of meetings on trust), we would conclude that the preaching-to-the-choir problem is widespread in the region.

Colombia is not an outlier. Indeed, for community-association meetings, the mean λ across each of the 116

government on different Likert scales, we use a linearized version to maintain comparable samples across the two estimates of λ.

response waves is 0.32 (fig. A18), meaning that the change in attendance from the lowest- to highest-trust categories is 32% of the overall attendance rate; the mean across the seven surveys taken in Colombia is slightly lower ($\hat{\lambda} = 0.30$). Similarly, for city council meetings, the mean $\hat{\lambda}$ across all response waves is 0.57, compared to 0.33 in Colombia. Colombia's position in the distribution of estimates from Latin America in figure 7 suggests that, were policymakers elsewhere to pursue interventions similar to ours, they might get similar results.

The positive relationship between trust and attendance is not obvious ex ante. One might think that public meetings would attract more critics than fans, serving as forums for expressing frustration, airing grievances, or making demands. The consistency of the positive selection across countries and across government institutions suggests otherwise, indicating that the relationship between prior trust and participation in public meetings—and its consequences for efforts to build trust—are underappreciated and merit further study.

DISCUSSION AND POLICY IMPLICATIONS

Our analysis leads to a counterintuitive conclusion for policymakers who seek to build trust in government: public meetings may be more valuable where they are likely to attract critics and perhaps be contentious; they may be less valuable where they are likely to draw fans and proceed harmoniously. How might policymakers act on this insight?

One intuitive answer might be to hold public meetings only in neighborhoods where few residents trust the police. Our data suggest that this is unlikely to work in practice, especially in Latin America. In the United States, trust in police is positively correlated with socioeconomic status (Slough and Torreblanca 2023), which is observable at the neighborhood level (in the census). In Latin America, in contrast, trust in police is at most very weakly correlated with socioeconomic status, making it difficult for policymakers to predict trust in police at the neighborhood level. In our data from Medellín, we observe that trust varies much more within neighborhoods than across them (the intraclass correlation coefficient is just 0.008) and, moreover, that the effect of the intervention was no greater in low-trust neighborhoods (table A9). Even if police could somehow target low-trust neighborhoods, then, they might end up preaching to the most-trusting residents of low-trust neighborhoods.

We instead propose two alternatives. First, policymakers and police supervisors might consider targeting at the individual level: distributing meeting invitations to people with prior contact with the criminal justice system, for example, or holding meetings exclusively for youth attendees. Second, policymakers might first estimate the trust–attendance gradient in their jurisdictions (using widely available survey data, as in the LAPOP analysis above), and then, if the gradient is steep (i.e., if the highest-trust people are much more likely to attend), they might forgo public meetings entirely and pursue other investments instead.

A different approach to avoiding the preaching-to-the-choir problem suggests imposing rather than inviting nonenforcement contact with officers. Peyton, Sierra-Arévalo, and Rand (2019) and Karim (2020) describe community policing interventions in which police knock on residents' doors. But our measures of uptake suggest that many citizens prefer not to engage with police. Imposition of police presence entails risks for residents and may also entail risks for police officers. Our recommendations instead preserve the benefits of opt-in, voluntary nonenforcement contact while minimizing the costs of ineffective interaction with constituents who already trust the police.

CONCLUSION

We show that a large-scale field experiment evaluating police–community meetings in Medellín, Colombia, did not alter attitudes toward police, despite the fact that the police held more than 500 meetings to which more than 8% of the adult population of treated neighborhoods showed up. We attribute this failure in part to positive selection into meetings: those who liked the police to begin with were more likely to attend and less likely to be impressed by what they saw. Officers ended up preaching to the choir.

More broadly, we suggest a focus on who participates in participatory institutions. A central goal of these initiatives is to broaden the role of citizens in policymaking and implementation, or "flattening access" (Grossman, Humphreys, and Sacramone-Lutz 2020). But when selection into participation is nonuniform, participatory institutions may be highly inefficient (as in the present study) or even detrimental (Hanson 2018; Slough 2021a). Better understanding of this citizen selection into participation can inform the design of more effective public institutions.

ACKNOWLEDGMENTS

We thank field manager Luz Ángela Londoño for her extraordinary work with the police and other counterparts in Medellín. We are also grateful to the team of field coordinators and research assistants who helped schedule and take notes on the police–community meetings. For assistance coding and analyzing data, we thank Estefanía Bolívar Méndez, Francisco Sánchez, Carolina Torreblanca, and Mark Williamson. For comments, we thank Graeme Blair, Guy Grossman, Jonathan Mummolo, Michael Weintraub, and two anonymous reviewers.

REFERENCES

Acemoglu, Daron, Leopoldo Fergusson, James Robinson, Dario Romero, and Juan F Vargas. 2020. "The Perils of High-Powered Incentives: Evidence from Colombia's False Positives." *American Economic Journal: Economic Policy* 12 (3): 1–43.

Amnesty International. 1994. "Political Violence in Colombia: Myth and Reality." Report AMR 23/001/1994.

Arias, Enrique Desmond, and Mark Ungar. 2009. "Community Policing and Latin America's Citizen Security Crisis." *Comparative Politics* 41 (4): 409–29.

Ba, Bocar A., and Roman G. Rivera. 2019. "The Effect of Police Oversight on Crime and Allegations of Misconduct: Evidence from Chicago." Paper no. 2109, All Faculty Scholarship, Penn Carey Law, University of Pennsylvania, Philadelphia, Pennsylvania. https://scholarship.law.upenn.edu/faculty_scholarship/2109.

Bhattacharya, Rajeev, Timothy M. Devinney, and Madan M. Pillutla. 1998. "A Formal Model of Trust Based on Outcomes." *Academy of Management Review* 23 (3): 459–72.

Blair, Graeme, Jeremy M. Weinstein, Fotini Christia, Eric Arias, Emile Badran, Robert A. Blair, Ali Cheema, et al. 2021. "Community Policing Does Not Build Citizen Trust in Police or Reduce Crime in the Global South." *Science* 374 (6571): abd3446.

Blattman, Christopher, Gustavo Duncan, Benjamin Lessing, and Santiago Tobon. 2022. "State-Building on the Margin: An Urban Experiment in Medellín." Working paper no. 29692, National Bureau of Economic Research, Cambridge, MA.

Brogden, Mike, and Preeti Nijhar. 2013. *Community Policing*. London: Taylor & Francis.

Calderón, Gabriela, Gustavo Robles, Alberto Díaz-Cayeros, and Beatriz Magaloni. 2015. "The Beheading of Criminal Organizations and the Dynamics of Violence in Mexico." *Journal of Conflict Resolution* 59 (8): 1455–85.

Casey, Katherine, Rachel Glennerster, and Edward Miguel. 2012. "Reshaping Institutions: Evidence on Aid Impacts Using a Preanalysis Plan." *Quarterly Journal of Economics* 127 (4): 1755–812.

Castillo, Juan Camilo, and Dorothy Kronick. 2020. "The Logic of Violence in Drug War." *American Political Science Review* 114 (3): 874–87.

Correa, Catalina Pérez, Carlos Silva Forné, and Ignacio Cano. 2019. "Monitor del Uso de la Fuerza Letal en América Latina: Un Estudio Comparativo de Brasil, Colombia, El Salvador, México y Venezuela." Monitor Fuerza Letal, Centro de Investigación y Docencia Económicas, Aguascalientes, Mexico.

Cruz, José Miguel, and Angélica Durán-Martínez. 2016. "Hiding Violence to Deal with the State: Criminal Pacts in El Salvador and Medellin." *Journal of Peace Research* 53 (2): 197–210.

de Benedictis-Kessner, Justin, Matthew A. Baum, Adam J. Berinsky, and Teppei Yamamoto. 2019. "Persuading the Enemy: Estimating the Persuasive Effects of Partisan Media with the Preference-Incorporating Choice and Assignment Design." *American Political Science Review* 113 (4): 902–16.

Dipoppa, Gemma. 2021. "How Criminal Organizations Expand to Strong States: Migrant Exploitation and Political Brokerage in Northern Italy." Working paper.

Einstein, Katherine Levine, David M. Glick, and Maxwell Palmer. 2019. *Neighborhood Defenders: Participatory Politics and America's Housing Crisis*. Cambridge, UK: Cambridge University Press.

Esberg, Jane, and Jonathan Mummolo. 2018. "Explaining Misperceptions of Crime." Working paper.

Falleti, Tulia G., and Thea N. Riofrancos. 2018. "Endogenous Participation: Strengthening Prior Consultation in Extractive Economies." *World Politics* 70 (1): 86–121.

Fiorina, Morris P. 1999. "Extreme Voices: A Dark Side of Civic Engagement." In Theda Skocpol and Morris P. Fiorina, eds., *Civic Engagement in American Democracy*. Washington, DC: Brookings Institution, 395–425.

García, Juan, Daniel Mejia, and Daniel Ortega. 2013. "Police Reform, Training and Crime: Experimental Evidence from Colombia's Plan Cuadrantes." Working paper no. 10497, Centro de Estudios sobre Desarrollo Económico, Universidad de los Andes, Bogotá, Colombia.

Goldfrank, Benjamin. 2007. "The Politics of Deepening Local Democracy: Decentralization, Party Institutionalization, and Participation." *Comparative Politics* 39 (2): 147–68.

González, Yanilda. 2016. "Varieties of Participatory Security: Assessing Community Participation in Policing in Latin America." *Public Administration and Development* 36 (2): 132–43.

González, Yanilda. 2019. "The Social Origins of Institutional Weakness and Change: Preferences, Power, and Police Reform in Latin America." *World Politics* 71 (1): 44–87.

González, Yanilda, and Lindsay Mayka. 2023. "Policing, Democratic Participation, and the Reproduction of Asymmetric Citizenship." *American Political Science Review* 117 (1): 263–79.

Grossman, Guy. 2014. "Do Selection Rules Affect Leader Responsiveness? Evidence from Rural Uganda." *Quarterly Journal of Political Science* 9 (1): 1–44.

Grossman, Guy, Macartan Humphreys, and Gabriella Sacramone-Lutz. 2020. "Information Technology and Political Engagement: Mixed Evidence from Uganda." *Journal of Politics* 82 (4): 1321–36.

Hanson, Rebecca. 2018. "Deepening Distrust: Why Participatory Experiments Are Not Always Good for Democracy." *Sociological Quarterly* 59 (1): 145–67.

Hanson, Rebecca, and Verónica Zubillaga. 2021. "From Carceral Punitivism to Systematic Killing: The Necropolitics of Policing in Post-Chávez Venezuela." *Violence: An International Journal* 2 (1): 65–84.

Hardin, Russell. 2002. *Trust and Trustworthiness*. New York: Russell Sage Foundation.

Holland, Alisha C. 2013. "Right on Crime? Conservative party politics and *Mano Dura* Policies in El Salvador." *Latin American Research Review* 48 (1): 44–67.

Hough, Mike, Jonathan Jackson, Ben Bradford, Andy Myhill, and Paul Quinton. 2010. "Procedural Justice, Trust, and Institutional Legitimacy." *Policing: A Journal of Policy and Practice* 4 (3): 203–10.

Humphreys, Macartan, Raúl Sánchez de la Sierra, and Peter Van der Windt. 2015. "Social Engineering in the Tropics: A Grassroots Democratization Experiment in the Congo." Working paper, Center for Development Economics and Policy, Columbia University, New York.

Karim, Sabrina. 2020. "Relational State Building in Areas of Limited Statehood: Experimental Evidence on the Attitudes of the Police." *American Political Science Review* 114 (2): 536–51.

Knox, Dean, Teppei Yamamoto, Matthew A. Baum, and Adam J. Berinsky. 2019. "Design, Identification, and Sensitivity Analysis for Patient Preference Trials." *Journal of the American Statistical Association* 114 (528): 1532–46.

Kronick, Dorothy. 2020. "Profits and Violence in Illegal Markets: Evidence from Venezuela." *Journal of Conflict Resolution* 64 (7–8): 1499–1523.

Lessing, Benjamin. 2017. *Making Peace in Drug Wars: Crackdowns and Cartels in Latin America*. Cambridge, UK: Cambridge University Press.

Magaloni, Beatriz, Edgar Franco-Vivanco, and Vanessa Melo. 2020. "Killing in the Slums: Social Order, Criminal Governance, and Police Violence in Rio de Janeiro." *American Political Science Review* 114 (2): 552–72.

Mayka, Lindsay. 2019. *Building Participatory Institutions in Latin America: Reform Coalitions and Institutional Change*. Cambridge, UK: Cambridge University Press.

McNulty, Stephanie. 2013. "Participatory Democracy? Exploring Peru's Efforts to Engage Civil Society in Local Governance." *Latin American Politics and Society* 55 (3): 69–92.

Molina-González, M. Dolores, Eugenio Martínez-Cámara, M. Teresa Martín-Valdivia, and L. Alfonso Urena-López. 2014. "Cross-Domain Sentiment Analysis Using Spanish Opinionated Words." In Elisabeth Métais, Mathieu Roche, and Maguelonne Teisseire, eds., *Natural Language Processing and Information Systems*. Cham, Switzerland: Springer, 214–19.

Müller, Markus-Michael. 2010. "Community Policing in Latin America: Lessons from Mexico City." *Revista Europea de Estudios Latinoamericanos y del Caribe/European Review of Latin American and Caribbean Studies* 88:21–37.

Mummolo, Jonathan, and Erik Peterson. 2017. "How Content Preferences Limit the Reach of Voting Aids." *American Politics Research* 45 (2): 159–85.

Osse, Anneke, and Ignacio Cano. 2017. "Police Deadly Use of Firearms: An International Comparison." *International Journal of Human Rights* 21 (5): 629–49.

Peyton, Kyle, Michael Sierra-Arévalo, and David G. Rand. 2019. "A Field Experiment on Community Policing and Police Legitimacy." *Proceedings of the National Academy of Sciences* 116 (40): 19894–98.

Skogan, Wesley G. 2022. "Prospects for Reform? The Collapse of Community Policing in Chicago." *University of Chicago Law Review* 89 (2): 383–404.

Skogan, Wesley, and Susan Hartnett. 1999. *Community Policing, Chicago Style*. Oxford: Oxford University Press.

Slough, Tara. 2021a. "Oversight, Capacity, and Inequality." Working paper, New York University.

Slough, Tara. 2021b. "Squeaky Wheels and Inequality in Bureaucratic Service Provision." Working paper, New York University.

Slough, Tara, and Carolina Torreblanca. 2023. "Class and the Development of Trust in Police in Latin America." Working paper.

Trejo, Guillermo, and Sandra Ley. 2020. *Votes, Drugs, and Violence: The Political Logic of Criminal Wars in Mexico*. Cambridge, UK: Cambridge University Press.

Tyler, Tom R. 2006. "Psychological Perspectives on Legitimacy and Legitimation." *Annual Review of Psychology* 57:375–400.

Tyler, Tom, and Jonathan Jackson. 2013. "Future Challenges in the Study of Legitimacy and Criminal Justice." Working paper no. 264, Yale Law School, New Haven, CT.

Ungar, Mark, and Enrique Desmond Arias. 2012. "Reassessing Community-Oriented Policing in Latin America." *Policing and Society* 22 (1): 1–13.

Ungar, Mark, and Leticia Salomón. 2012. "Community Policing in Honduras: Local Impacts of a National Programme." *Policing and Society* 22 (1): 28–42.

Vásquez, J. C. Ruiz. 2012. "Community Police in Colombia: An Idle Process." *Policing and Society* 22 (1): 43–56.

Wampler, Brian. 2008. "When Does Participatory Democracy Deepen the Quality of Democracy? Lessons from Brazil." *Comparative Politics* 41 (1): 61–81.

Wampler, Brian. 2012. "Entering the State: Civil Society Activism and Participatory Governance in Brazil." *Political Studies* 60 (2): 341–62.

Yoder, Jesse. 2020. "Does Property Ownership Lead to Participation in Local Politics? Evidence from Property Records and Meeting Minutes." *American Political Science Review* 114 (4): 1213–29.