# Preaching to the Choir: A Problem of Participatory Interventions

Rebecca Hanson* Dorothy Kronick† Tara Slough‡

December 7, 2022§

**Abstract**

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 our intervention and, we argue, poses an under-recognized threat to related initiatives across the globe.

*Assistant Professor, University of Florida, r.hanson@ufl.edu

†Assistant Professor, University of California, Berkeley, kronick@berkeley.edu

‡Assistant Professor, New York University, tara.slough@nyu.edu

§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, and Michael Weintraub.

The people who show up for public meetings with their city councilors, 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 DRC, 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 socio-economic 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 under-recognized 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. We find that the meetings did not build trust in the police or boost crime reporting (according to pre-registered 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.[1]

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 (Calderón et al., 2015; Cruz and Durán-Martínez, 2016; Lessing, 2017; Magaloni, Franco-Vivanco, and Melo, 2020; Kronick, 2020; Castillo and Kronick, 2020; Trejo and Ley, 2020; Acemoglu et al., 2020; Dipoppa, 2021). We add that unanticipated responses from ordinary citizens (in this case, residents considering attending police–community meetings) can likewise undermine interventions that might otherwise improve welfare.

Similarly, literature on participatory governance often considers the implications of *who participates* for resource allocation, sometimes finding that participatory fora are inclusive and improve resource allocation, other times finding that they amplify vocal minorities and distort policy outcomes (Goldfrank, 2007; Wampler, 2008; Mansuri and Rao, 2013; McNulty, 2013; Grossman, 2014; Falleti and Riofrancos, 2018; Mayka, 2019; 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.

Our paper is closely related to Gonzalez and Mayka (2022), who study police–community meetings in São Paulo, Brazil. Using qualitative meeting observation and quantitative coding of an original collection of minutes from 793 meetings, Gonzalez and Mayka find that attendees at police–community meetings often call for "police repression of marginalized groups."[2] Their con-

[1] These differences are significantly different from zero in a two-tailed test in 57 of 115 country-survey waves.

[2] Specifically, a meeting includes "demands for repression" if (a) participants asked police to "do something" about

clusion that police–community meetings in São Paulo serve to amplify the voices of relatively pro-police citizens is consonant with our finding that police–community meetings in Colombia draw attendees who are relatively pro-police. On the other hand, when we apply their coding rules to a sample of our meeting minutes from Medellín, we do not observe comparable rates of "demands for police repression."[3] In our context, then, a related pattern of selection into police–community meetings has different political consequences.

We also build on build on previous work about community policing (e.g. Greene, 2000; Skogan and Hartnett, 1999; Ungar and Arias, 2012; González, 2022). The debate over community policing has turned largely 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 fundamentally shifting 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;" 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, 2022*a*,*b*), 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.

a marginalized group, in a way that implied coercion rather than (say) service provision, or (b) participants asked police to use physical force against a marginalized group. Using this coding rule, they find that 18% of meetings include "demands for repression" against youth; 15% of meetings include demands for repression against drug users; 6% against people experiencing homelenessness; and smaller but nontrivial proportions of meetings against street vendors and other marginalized groups.

[3]For example, only 3.1% of meetings could be coded as including a call for police action against youth (compared to 18% in São Paulo). See Appendix A9 for details.

Finally, we contribute to a methods literature about estimating treatment effects in experiments with selection (de Benedictis-Kessner et al., 2019; Knox et al., 2019). Much of this work 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 wouldn't 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.

# 1 How participation-based interventions change beliefs

In theory, non-enforcement 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 Figure 1). Similarly, non-enforcement contact could boost officers' opinions of the community and reorder policing priorities. And just as revised beliefs might change the behavior of members of the community, revised beliefs might change the behavior of police officers themselves, ideally improving police services and (perhaps) security outcomes, which would then further cultivate rosy views all around, continuing the virtuous cycle.

Non-enforcement contact between police and the policed, often in the form of police–community meetings (or *beat meetings*), is a common feature of community policing (see Blair et al., 2021, for a systematic review). It typically works (or doesn't) through changing beliefs (or not) (Figure 1). 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,

Figure 1: Beliefs: One Link in a Virtuous Cycle

Police–community meetings 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 paper, intervene at other points in the cycle.

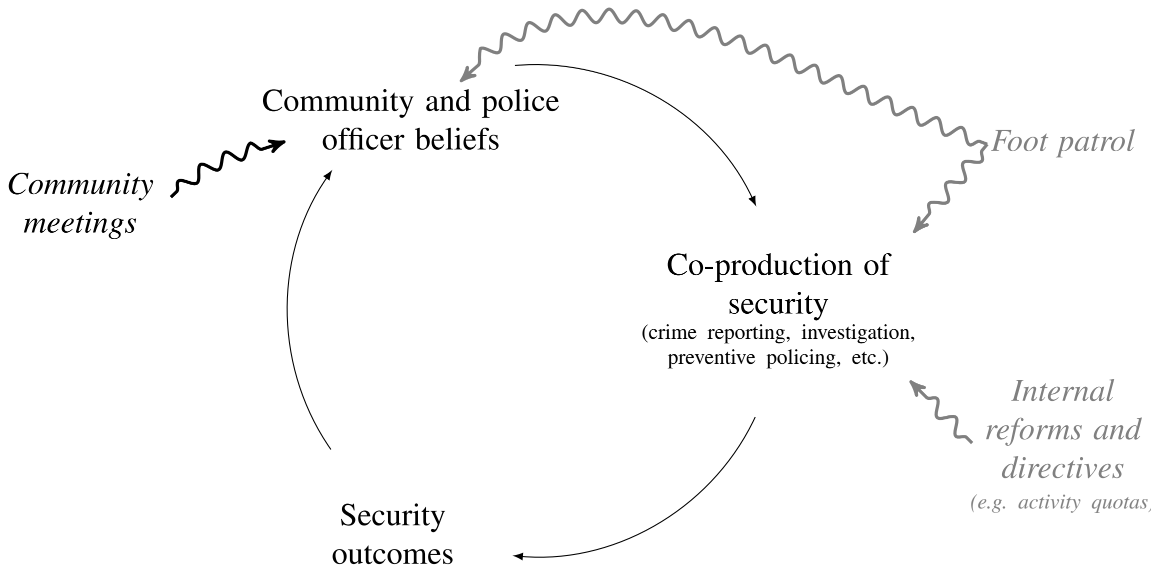


affect police behavior without directly reshaping community opinions or officers' beliefs about the policed.[4]

We suggest that, where policy intervenes in a *cycle* (virtuous or otherwise), researchers might productively study the *first* outcome that the policy affects. That is one reason for our focus on the effect of police–community meetings on *beliefs* (see Figure 1), rather than behavior or security outcomes.[5] Police-community meetings alone are unlikely to have salutary effects unless they reshape the beliefs of community members and/or police officers.[6] 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 any potential virtuous cycle, and the part most proximate to

[4]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.

[5]Another reason 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 pre-analysis plan, we do not consider crime outcomes.

[6]In principle, behavior change could occur in the absence of a change in beliefs: it could stem purely from revised incentives, for example. But our intervention did not meaningfully affect incentives.

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 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 (Esberg and Mummolo, 2018). 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 fundamentally cognitive (consistent with Hardin, 2002; Bhattacharya, Devinney, and Pillutla, 1998), and we posit that beliefs about the trustworthiness of 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 (Gonzalez, 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 non-enforcement 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, prior beliefs determine the consumption value that people derive from police–community meetings: those who most trust the police might especially appreciate the opportunity to spend leisure time in the company of officers.[7] If prior beliefs determine individual-specific consumption values, and

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

if consumption values drive attendance, then we would expect prior beliefs to be correlated with attendance: those who trust the police at baseline might 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. If people who trust the police to begin with are more likely to participate in meetings, then we are less likely to observe positive changes in 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 susceptible to revising upward. The stronger the positive correlation between baseline beliefs and attendance, the smaller the effects of the intervention. Even if the population as a whole holds unduly negative beliefs about the police, and even if police–community meetings provide an informative signal about police trustworthiness, positive selection into attendance will dampen the corrective effects of the intervention. This is the problem of preaching to the choir.

While our project is focused on community-police meetings and citizen attitudes toward police, we argue that 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. The very establishment of a 311 line could have one or both of these effects, regardless of the consequences for service delivery. Indeed, selection into sharing information via 311—like that documented by Slough (2022*a*)—could exaggerate or attenuate any attendant changes in citizen beliefs. While the structures underlying virtuous cycles are specific to each type of intervention, these two connections to beliefs—and the interaction between them—are general.

## 2 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 (Figure 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.[8] And while we lack a reliable multi-decade 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 extra-judicial 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: p. 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 (Osse and Cano, 2017; Correa, Forné, and Cano, 2019; Hanson and Zubillaga, 2021). Moreover, especially since 2010, the Colombian police have embraced many aspects of community policing, with salutary effects on the quality of police services (Garcia, Mejia, and Ortega, 2013).

Yet confidence in the police has not improved. The proportion of survey respondents expressing confidence in the police[9] declined from 56% (in 2004) to 42% (in 2018), according to survey data from the Latin American Public Opinion Project (Figure 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. In other words, perhaps people evaluate the police primarily based on outcomes other than crime rates and/or physical abuse. Petty corruption, for example. If officers' bribe-seeking increased over the time period shown in Figure 2, trust in the

[8]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.

[9]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)."

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.

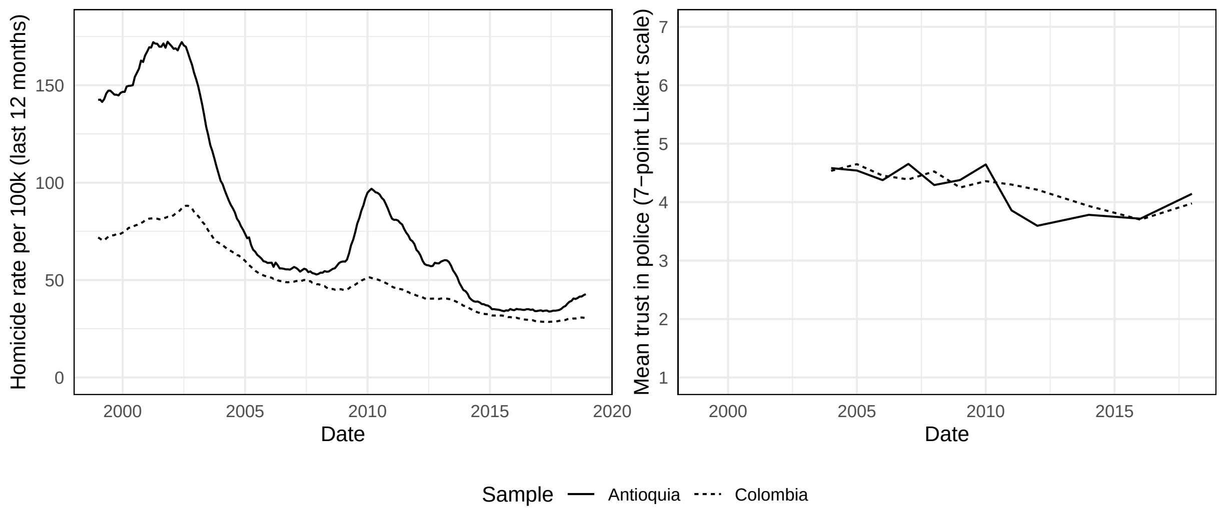


police could have fallen as a result. Another possible explanation is that people paradoxically value *mano dura*, as several studies have argued (Holland, 2013). Were that the case, the decline in police use of lethal force might actually drive more negative assessments of the police.

A more likely explanation, in our view, is that opinions of the police are shaped in part by high-profile negative news—especially about the "riot-control" unit that responds to protests—and by knowledge of grave human rights violations perpetrated by the Colombian military (Acemoglu et al., 2020). Given that the Colombian police remain housed in the Ministry of Defense, that the command hierarchy mimics that of the military, that police uniforms resemble those of the army, and, critically, that the police worked alongside the military during much of Colombia's civil war, it would not be surprising if news about the armed forces affected beliefs about the police. Were that the case, police–community meetings might provide a valuable opportunity for people to learn about their local officers. Indeed, contrary to work that views misperceptions as resistant to new information, Esberg and Mummolo (2018) 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" (3).

This conclusion echoes the expectations of our counterparts in the Medellín police. The Metropolitan Police of the Aburrá Valley, or MEVAL, 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.

## 3 Research Design and Data

**Sampling and 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 (the airport and parts of downtown, for example). Our sample contains approximately 96% of Medellín residents.

Because police beats are fairly large—the median population was 5,348 (in the 2005 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 police beat (Figure 3). 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

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

Figure 3: Definition of Units of Analysis

Within each police beat, shown here in black outlines, we distributed meeting invitations and conducted our surveys within *prioritized neighborhoods*, shown here in blue.

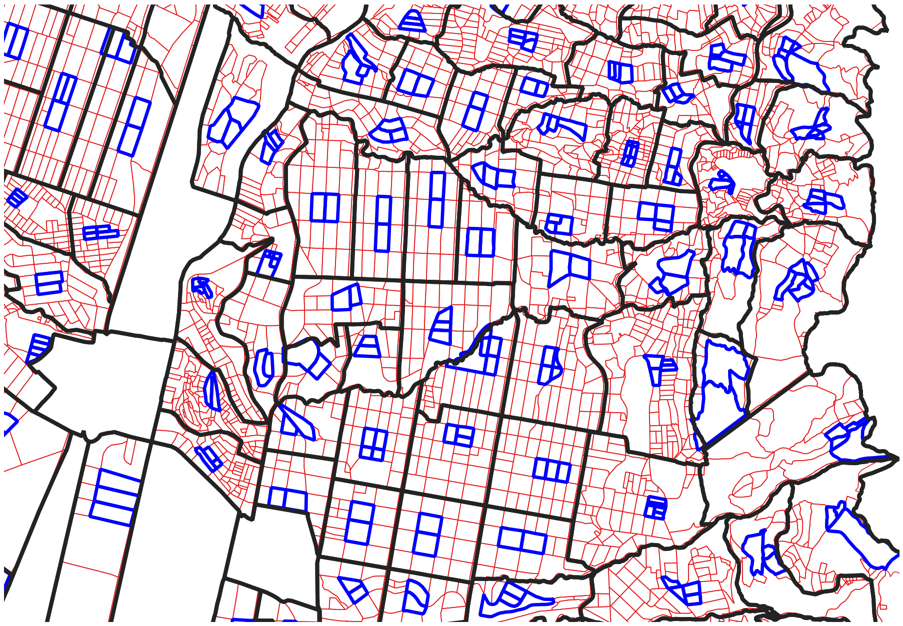

police beats that (a) belong to one of eleven police station groups[11] and (b) share the same treatment status in a simultaneous intervention conducted by other researchers (Blattman et al., 2022). Within each block of four beats, we randomly assign two beats to the police–community meetings treatment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

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 (Appendix Figure A2) door-to-door among the (approximately) 400 households within each treated neighborhood. The study team also hung a poster at the meet-

[11]There are fourteen police stations, but some of the stations are relatively small, and thus we grouped 6 of these into 3, for a total of 11 police station groups. The fourteen stations are Aranjuez, Belén, Buenos Aires, Candelaria, Castilla, Doce de Octubre, Laureles, Manrique, Poblado, Popular, San Antonio de Prado, San Javier, Santa Cruz, and Villa Hermosa; we grouped Aranjuez with Manrique, Buenos Aires with Vila Hermosa, and Popular with Santa Cruz.

ing 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 structure and content of the meetings in more detail below.

**Ethics.** A field experiment involving the police entails special ethical considerations. For one, creating opportunities for non-enforcement contact between officers and citizens implies creating opportunities for contentious interactions—perhaps even opportunities for conflict. This risk was smaller in our context than in many others; as noted above, despite relatively high homicide rates, Colombia has relatively low levels of police violence. Consultation with local researchers and NGOs 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 (JAC, or *Junta de Acción Comunal*) 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.

**Data.** We measure residents' attitudes toward the police using a panel survey. At baseline, between January and April, 2018, the survey firm Invamer conducted an in-person survey of 5,205 residents of the prioritized neighborhoods (15 residents per neighborhood). At endline, between September and December, 2019, Invamer was able to re-contact 2,434 of these original interviewees, as well as surveying 1,210 new interviewees in treatment and control neighborhoods. There is no evidence that attrition was related to treatment assignment (Table A2 and Figure A1). Our survey instrument was written in concert with the other studies in Blair et al. (2021). Appendix Table A5 describes how we operationalize our principal outcomes.

Figure 4: Study Timeline

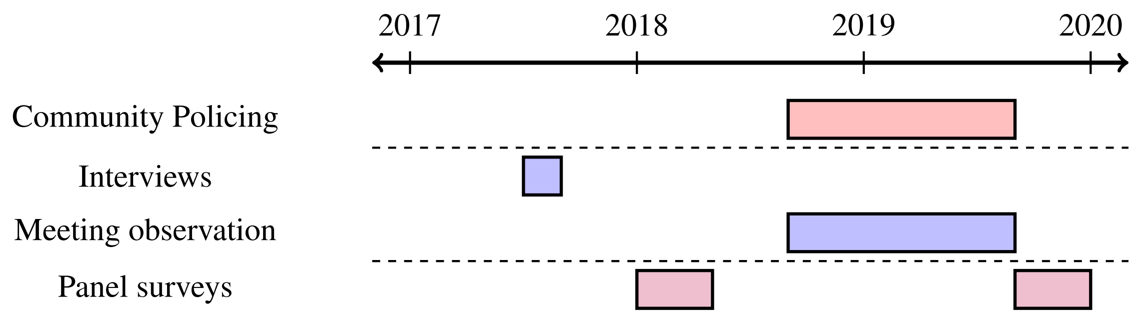


We also collected two types of qualitative data. First, we conducted one focus group prior to the start of the intervention. Second, we collected detailed notes on each of the 519 police–community meetings. One of the coauthors trained research assistants to take short-form field notes at the meetings 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 4 summarizes the timing of our data collection and intervention.

**Estimation.** Our primary pre-specified estimand is the intent-to-treat effect (ITT): the effect of assignment to police–community meetings on attitudes toward the police. We estimate the ITT effect using:

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

where $Y_{ijb}$ is the 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 $\boldsymbol{\gamma}_b$ is a vector of block fixed effects. Some specifications also include baseline measures of the survey response $Y_{ijb}$.

In addition to this pre-specified analysis, we run a series of simulations 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 ATT, for two reasons. First, censoring: if enough people express prior beliefs at the maximum of the measurement scale, then we cannot observe improvements in trust. Second, 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 average treatment effect on the treated. Censoring and mean reversion are fundamentally problems of measurement; in that sense, they pose challenges for the researcher but not for the policymaker.

Part of our discussion considers the average treatment effect on the treated (ATT), for which we use only data from treated neighborhoods to estimate:

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

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

# 4 Results

Using ITT estimates from (pre-specified) 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 5, and in Appendix Tables A6–A8. The point estimates are substantively small—generally less than 0.05 standard deviations, and always less than 0.1 standard deviations—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.

[12]Results are robust to the includsion of block fixed effects, $\gamma_b$. We exclude these fixed effects from the primary ATT specifications because there are some blocks in which there is no variation in $A_{jb}$ for one participation indicator. By omitting the fixed effects, we compare the same effective sample across specifications.

Figure 5: Intent-to-Treat Effects are Small and Indistinguishable from Zero

This figure plots estimates of (pre-specified) Equation 1 with 95% confidence intervals. Tables A3–A5 provide additional details.

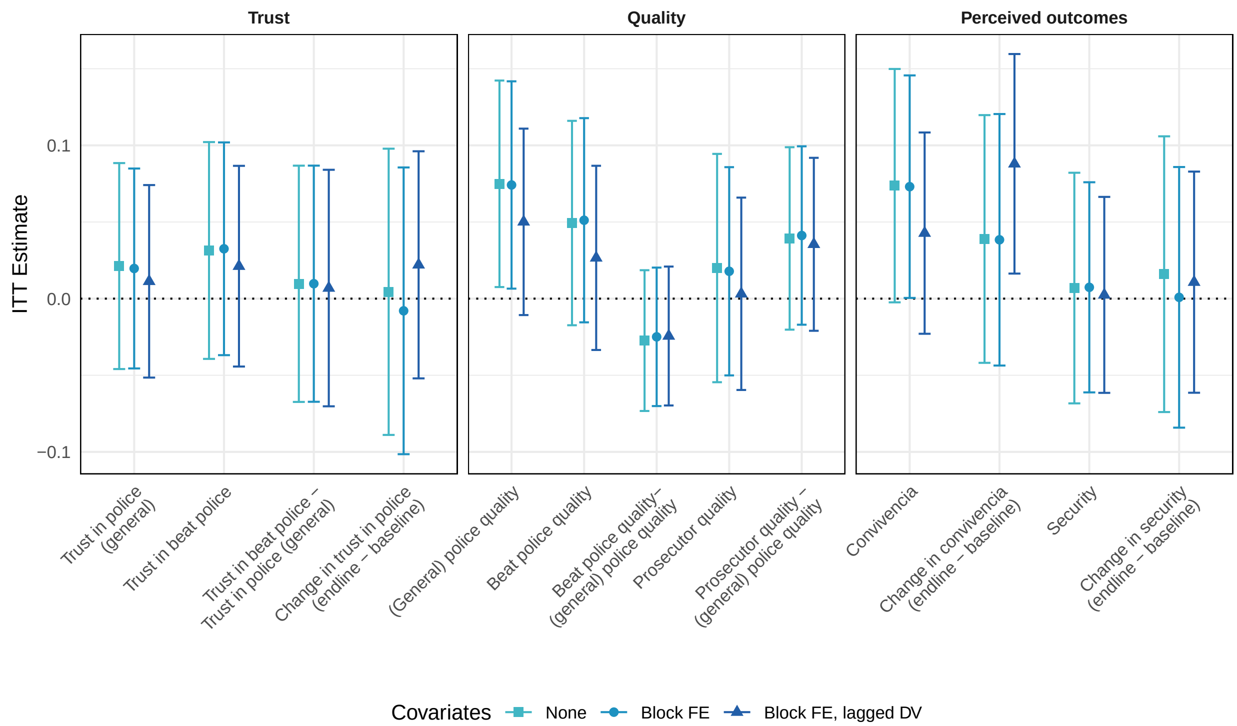


## 4.1 Imperfect reach

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 meetings) 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 of

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

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.

| | Baseline exposure | | Baseline attendance | | Endline exposure | | Endline Attendance | |
|---|---|---|---|---|---|---|---|---|
| | (1) | (2) | (3) | (4) | (5) | (6) | (7) | (8) |
| PANEL A | | | | | | | | |
| Meetings | $-0.003$ | $0.005$ | $-0.003$ | $0.003$ | $0.163^{***}$ | $0.216^{***}$ | $0.041^{***}$ | $0.052^{***}$ |
| | $(0.007)$ | $(0.010)$ | $(0.007)$ | $(0.010)$ | $(0.018)$ | $(0.022)$ | $(0.009)$ | $(0.011)$ |
| Block FE | ✓ | ✓ | ✓ | ✓ | ✓ | ✓ | ✓ | ✓ |
| Observations | 5205 | 2434 | 5205 | 2434 | 3644 | 2434 | 3644 | 2434 |
| Clusters | 347 | 347 | 347 | 347 | 347 | 347 | 347 | 347 |
| Control Mean | 0.054 | 0.054 | 0.051 | 0.052 | 0.264 | 0.291 | 0.059 | 0.063 |
| Sample | Baseline | Panel | Baseline | Panel | Endline | Panel | Endline | Panel |

$^{***}p < 0.001$; $^{**}p < 0.01$; $^{*}p < 0.05$

"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 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, i.e., prior to the start of the intervention. In other words, the treatment and control groups are balanced on pre-treatment exposure to police–community meetings. Moreover, Table 1 reveals that, in both groups, just 5% of respondents reported hearing about or attending police–community meetings during the year prior to the baseline survey. 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 Appendix Figure A3 also suggest that there were spillovers in *exposure* to police–community meetings, but not in attendance.

### 4.2 Varied meeting quality

A second source of attenuation in the effect of the intervention was heterogeneity in the signal provided by the meetings themselves. In other words, police officers looked good in some meetings and terrible in others. In this section, 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; second, by using the NRC emotion lexicon (in Spanish) and the improved SOL (iSOL) 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 alleged 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

Figure 6: 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.

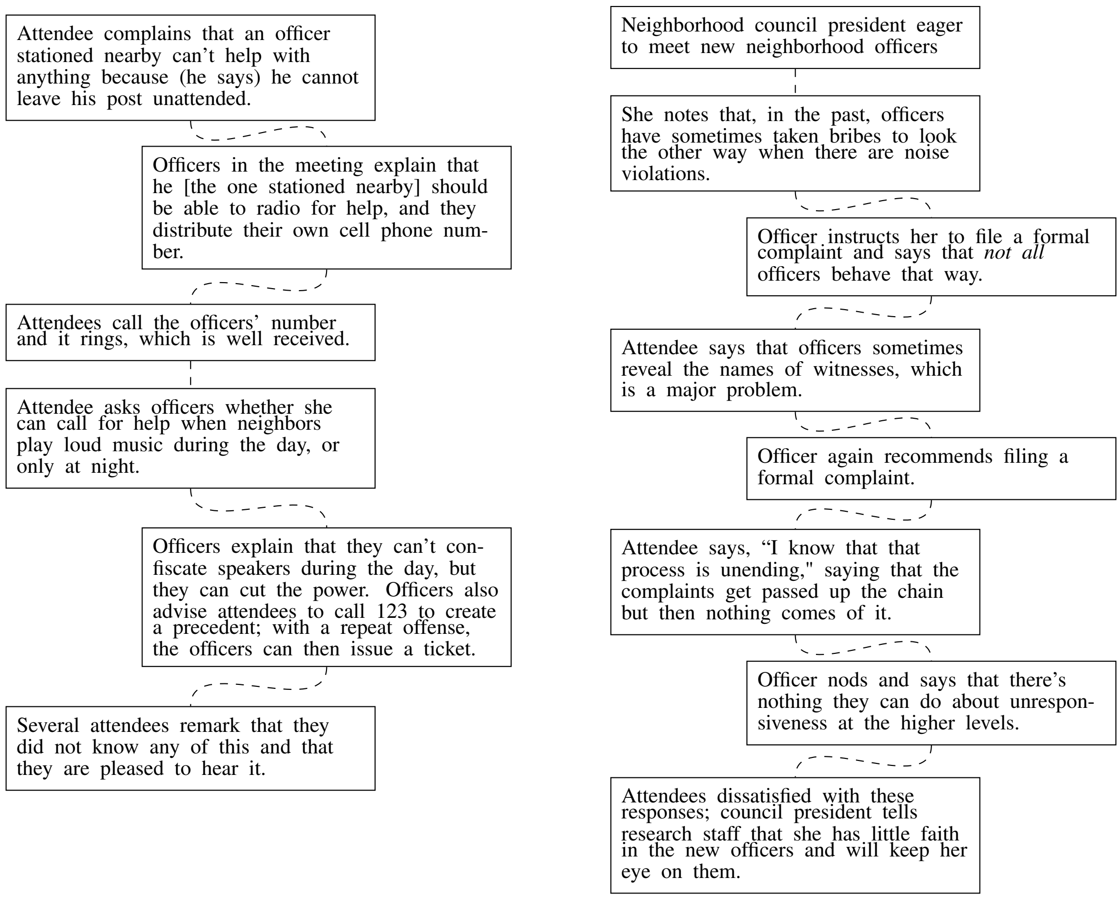


council (JAC) berated the officers, accusing them over and over of corruption. The president of the JAC accused the officers of looking the other way when local businesses violated noise ordinances, "as long as they pay you off." One of the two officers responded by invoking a bad-apple narrative, noting that "that is not how all police officers behave," but the council president wasn't 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. "He thinks that I don't know him, but I know exactly who he is."

Figure 6 summarizes these interactions. Our objective here is both to succinctly illustrate the process by which police–community meetings might have changed beliefs, and to show 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 (Appendix Figure A8). In other words, it is not the case that the most-positive meetings generally took place in the most pro-police neighborhoods. Rather, the tone of the meetings was highly idiosyncratic. As we show below, this pattern of heterogeneity in the signal from the meetings likely attenuated their effects.

### 4.3 Positive selection into police–community meetings

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. In other words, there was positive selection into meetings.

For Colombia, the top left panel of Figure 7 plots the rate at which respondents (in treated neighborhoods) 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 (top right panel of Figure 7). Nor are these differences driven by small numbers of people at the extremes of the distribution of trust in police. The bottom two panels of Figure 7 show that substantial fractions of respondents say that they trust the police "not at all" (1) or "very much" (4). Indeed, among meeting attendees, trusting the police "very much" (at baseline) is the modal response (bottom right panel).

The five other sites in the community-policing meta-study were Liberia, Uganda, the Philippines, Pakistan, and Brazil. In Liberia, Uganda, and the Philippines, as in Colombia, we observe

Figure 7: Baseline Trust in Police Strongly Predicts Engagement

The top panels plot 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. The bottom panel plots the distribution of baseline trust in police among those who *heard about* or *did not hear about* meetings (left), and among those who *attended* or *did not attend* (right).

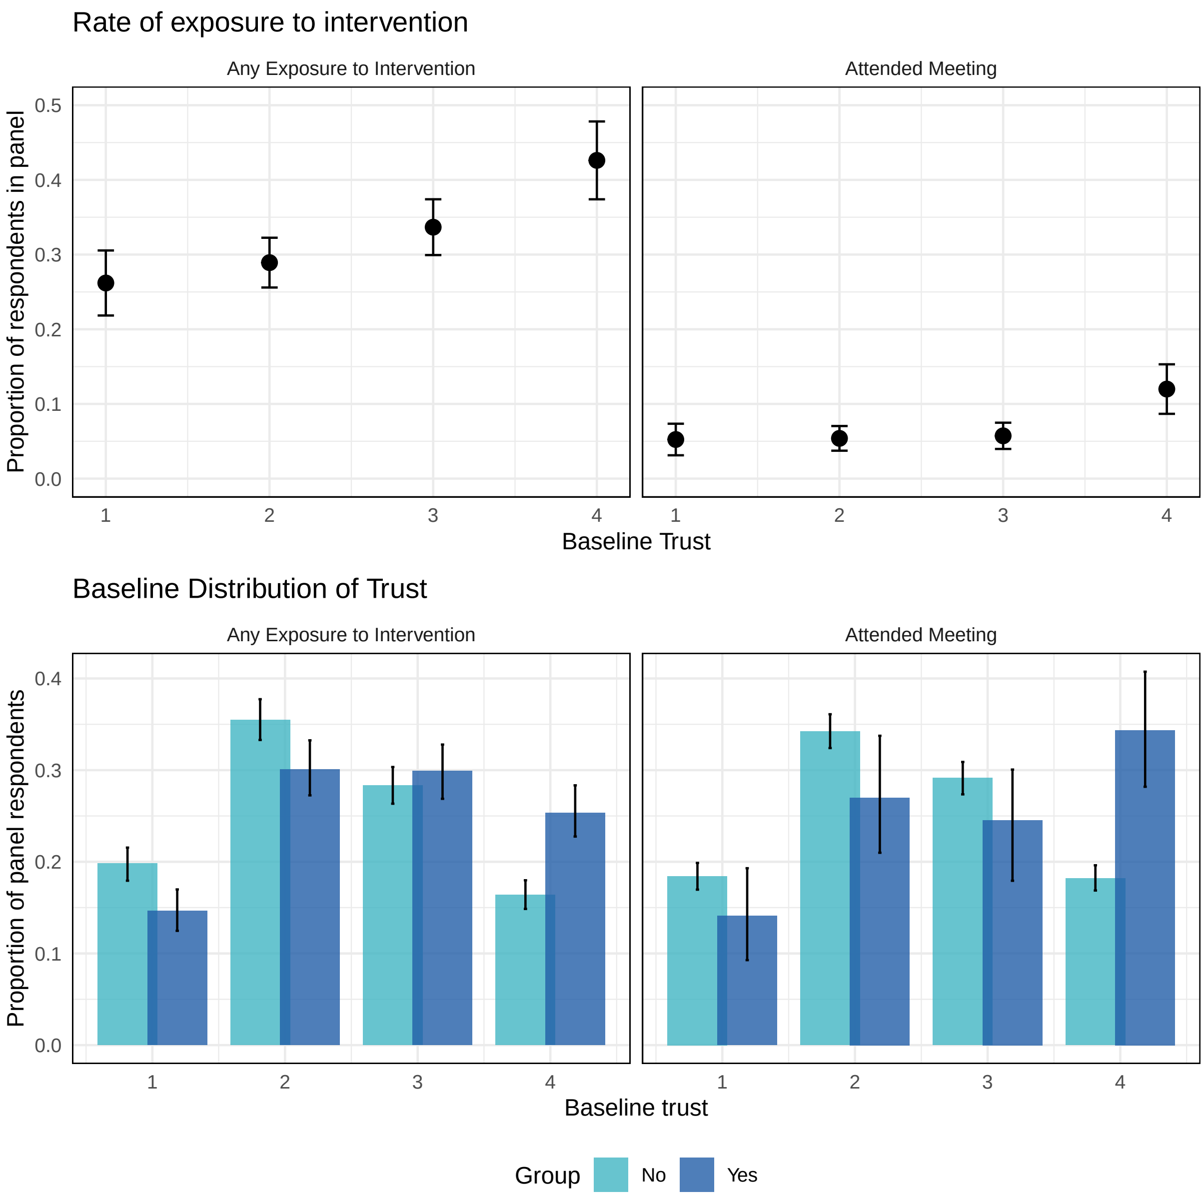


substantively large differences in attendance rates between people with the lowest and highest levels of trust in police. In Pakistan, attendance was flat across baseline trust categories; one possible reason is that self-reported trust appears more fluid: the correlation between baseline and endline

Table 2: Meeting Attendance is Positively Correlated with Trust

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.

| | P(Attend ∣ Lowest Trust) | P(Attend ∣ Highest Trust) | Difference |
|---|---|---|---|
| Liberia | 0.255 | 0.376 | 0.121 |
| | (0.058) | (0.064) | (0.108) |
| Uganda† | 0.286 | 0.386 | 0.100 |
| | (0.037) | (0.032) | (0.045) |
| Colombia | 0.084 | 0.184 | 0.100 |
| | (0.019) | (0.029) | (0.034) |
| Philippines† | 0.097 | 0.179 | 0.083 |
| | (0.053) | (0.011) | (0.052) |
| Pakistan | 0.059 | 0.038 | -0.020 |
| | (0.019) | (0.037) | (0.041) |
| Brazil | 0.167 | 0.068 | -0.098 |
| | (0.081) | (0.019) | (0.079) |

† 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.

beliefs is zero (compared to 0.48 in Colombia or 0.38 in Uganda). In Brazil, similarly, we do not observe evidence of positive selection into meetings, though the low recontact rate in the panel survey adds considerable noise to our estimates.

To see this, consider the estimates in Table 2. 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. (In Liberia, the survey was a repeated cross-section rather than a panel, meaning that we cannot study selection at the individual level.) 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.[13] If we instead compare the top *two* trust categories in the Philippines to the bottom two, the difference in attendance rates is 6.4 percentage points (s.e. 3.2%). In Pakistan, in contrast, attendance was nearly flat across baseline and endline trust categories; one possibility is that *trust in police* is simply more fluid in this context, given that the correlation between baseline and endline beliefs is zero. In Brazil, we see no evidence of positive selection into meetings, though the low recontact rate makes it somewhat difficult to assess. Of the 493 recontacted respondents in treated units, just 24 reported the lowest level of trust at baseline, of whom 4 (16%) attended, compared to 19 of the 278 respondents (7%) with the highest level of trust. (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.) 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).

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 ten categories of predictors (treatment assignment, age, gender, and various measures of socio-economic status), we find that baseline trust in police is second only to treatment assignment in predicting whether respondents *hear about* meetings—and that baseline trust is the *strongest* predictor of whether respondents attend meetings, even stronger than treatment assignment (Appendix A6.2).

For this reason, we investigate censoring and mean reversion in Appendices A8.1 and A8.2, respectively, focusing our discussion on the problem of preaching to the choir. One way to get a sense of the magnitude of this problem is to consider the ATT (Equation 2). Columns (2)–(3) and (5)–(6) of Table 3 (Panel A) reveal that, if we define the endogenously treated population as

[13]In treated neighborhoods in the Philippines, 31 respondents (1%) report the lowest level of trust, 62 (2%) report the second-lowest level, 313 (11%) report the middle category, while 1,044 (35%) and 1,495 (51%) report the second-highest and highest levels, respectively.

*attendees*, 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.

Positive selection into meeting awareness and attendance is unsurprising if we believe that the meetings provide (at least some) consumption value, as we posited above. To the extent that the opportunity for non-enforcement contact with officers is itself a police service that residents value, we might expect those who trust the police to value it more. People who dislike or distrust the police are unlikely to choose to spend an evening with them.

Positive selection should attenuate our *estimates* of treatment effects via *censoring* and *mean reversion*, and it should attenuate the treatment effects themselves because of the preaching-to-the-choir problem. Censoring and mean reversion are issues of measurement: the former arises because of the scale of the survey question, 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. Preaching to the choir, in contrast, is an issue of substance: it attenuates the actual effect of the intervention, not just our estimate of the effect.

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 | $-0.088$ | | $-0.057$ | $0.020$ | | $0.058$ |
| | $(0.064)$ | | $(0.071)$ | $(0.059)$ | | $(0.065)$ |
| Attended meeting | | $-0.179$ | $-0.143$ | | $-0.141$ | $-0.177$ |
| | | $(0.110)$ | $(0.121)$ | | $(0.094)$ | $(0.102)$ |
| PANEL B | | | | | | |
| Any exposure | $0.088$ | | $0.094$ | $0.121^{*}$ | | $0.131^{*}$ |
| | $(0.052)$ | | $(0.057)$ | $(0.051)$ | | $(0.056)$ |
| Attended meeting | | $0.029$ | $-0.028$ | | $0.032$ | $-0.049$ |
| | | $(0.095)$ | $(0.103)$ | | $(0.074)$ | $(0.082)$ |
| Baseline as covariate | ✓ | ✓ | ✓ | ✓ | ✓ | ✓ |
| Any exposure: Panel A = Panel B, $p$-value | $0.035$ | | $0.099$ | $0.197$ | | $0.393$ |
| Attended mtg: Panel A = Panel B, $p$-value | | $0.157$ | $0.206$ | | $0.153$ | $0.136$ |
| Untreated mean (sd) | $-0.121$ | $-0.143$ | $-0.121$ | $0.034$ | $0.055$ | $0.034$ |
| Untreated sd | $1.175$ | $1.171$ | $1.175$ | $0.961$ | $0.951$ | $0.961$ |
| Observations | 1191 | 1191 | 1191 | 1206 | 1206 | 1206 |
| Clusters | 173 | 173 | 173 | 173 | 173 | 173 |

$^{***}p < 0.001$; $^{**}p < 0.01$; $^{*}p < 0.05$

### 4.4 Comparing the influence of reach, signal, and selection

To gauge how these challenges influence the effect of police–community meetings, we run a series of simulations. The simulation starts with a population whose prior trust in police, $\pi$, is drawn from a discrete uniform distribution between $1$ and $4$ (mimicing the four-point *trust in police* measure from our survey). We then assign half of the population to be invited to police–community meetings, and we vary both (a) the overall rate of meeting attendance and (b) the extent to which attendance covaries with prior beliefs. For simplicity, we consider only linear relationships between prior beliefs and attendance, so that the change in attendance rates with a one-point change in prior beliefs is $\Delta$. We then define *positive selection* as:

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

In other words, $\theta$ 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 $\Delta$ 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 expose attendees (and only attendees) to a signal about police trustworthiness, allow them to update in a Bayesian manner, and then calculate the ITT and the ATT.

We begin with a version of the simulation in which the meetings provide maximally positive information about the police, i.e., in which 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 8 presents the results. Consider first the simulation in which the overall attendance rate (compliance) is 50%. For an overall compliance rate of 50%, the purple line in the left panel of Figure 8 shows how our estimate of the ITT changes with positive selection. 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. In our data, this is 0.36 standard deviations. But positive selection quickly shrinks this estimate: 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 ($\theta = 0$). With lower rates of compliance, such as 30% (blue line) or 10% (teal line), positive selection similarly attenuates the ITT.

Positive selection also quickly attenuates the ATT. The right panel of Figure 8 shows that, as selection increases from the theoretical minimum ($\theta = 0$) to the theoretical maximum ($\theta = 2$), the ATT shrinks from 0.75 to 0.5 on a 1–4 scale. Even at $\theta = 1$, the value that we observe in our data, the ATT is just 82% of its value in the absence of selection.

This simulation emphasizes that positive selection into meetings significantly shrinks the estimated treatment effects *even when* the intervention (a) provides maximally positive information

Figure 8: The Problem of Preaching to the Choir

Using a series of simulations, these figures plot how estimates of the intent-to-treat effect (ITT) and the average treatment effect on the treated (ATT) change with the degree of positive selection into meetings. The x-axes span $\theta$, our parameterization 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.

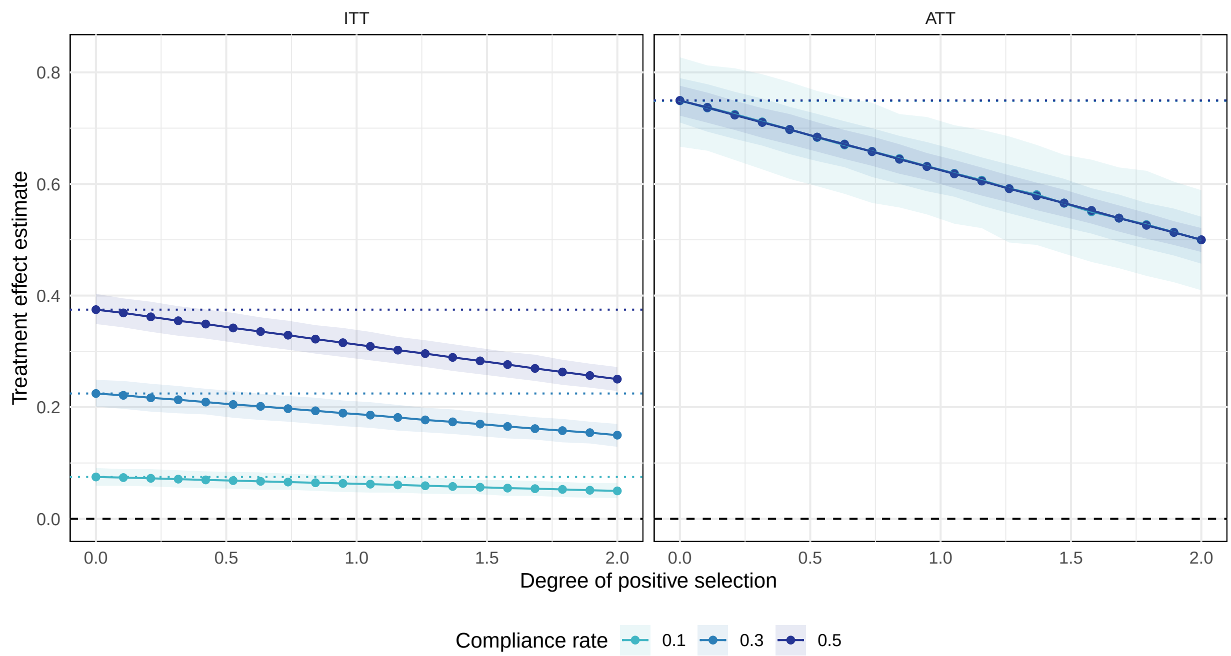


about police and (b) reaches a sizable fraction of the population. The consequences of positive selection are at least as perverse in an environment less favorable to the intervention. If, for example, prior beliefs are *not* biased—if, for example, prior beliefs are uniformly distributed on [1,4] and the meetings provide a signal of $2.5$—then the ITT and ATT would be zero in the absence of positive selection but quickly become *negative* as $\theta$ increases.

It is, of course, impossible to quantify precisely how positive selection affected our actual results. But a back-of-the-envelope exercise (see Appendix A8.3) suggests that, in our case, positive selection attenuated the ITT estimates by about $33\%$ and the ATT estimates by approximately $50\%$. We conclude that, while imperfect reach and heterogeneous meeting quality were indeed impediments to an effective intervention, positive selection into police–community meetings also posed a significant challenge.

## 5 Discussion

Our analysis produces a counterintuitive policy recommendation for policymakers who seek to build trust in government: hold *more* public meetings where they are likely to attract critics and perhaps be contentious; invest less in such meetings where they are likely to draw fans and proceed harmoniously. In this section, we suggest that implementing this recommendation is more tractable than it might seem.

In particular, we propose an empirical approach to *anticipating* the extent of the preaching-to-the-choir problem. We find, using our data from Medellín, that the relationship between prior trust and participation in pre-intervention meetings—which is to say, other types of meetings that the police had held in the past[14]—looks almost exactly like the relationship between prior trust and participation in the police–community meetings of the intervention. Using our baseline measure of trust in police and the baseline measure of attendance at *previous* meetings, we estimate positive selection of $\theta = 1.10$ (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 $\theta = 1.01$ (95% CI: $[0.43, 1.33]$). In other words, we could have anticipated the degree of positive selection into police–community meetings simply by studying the relationship between trust and attendance in the baseline survey alone. Had we done so, we might have advised against investment in public meetings.

Policymakers elsewhere can anticipate the extent of the preaching-to-the-choir problem by analyzing widely available survey data. The Latin American Public Opinion Project (LAPOP), for example, routinely asks respondents about trust in public institutions and about attendance at public meetings (such as those convened by city councils or neighborhood associations). Where the relationship between trust and attendance is strong and positive, policymakers might expect participatory interventions to have more muted effects on trust (because of the preaching-to-the-choir problem). Where it is flat or negative—i.e., where meetings appear to attract more critics than fans—policymakers might expect more salutary effects of participatory interventions on trust.

[14] We use the same measure of baseline attendance that is reported in Columns (1)–(4) of Table 1.

Figure 9: 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) the difference between city-council-meeting attendance rates among respondents with the most trust in city government and city-council-meeting attendance rates among respondents with the least trust. On the right, we report analogous differences for attendance at neighborhood association meetings, given trust in community members. The segments depict 95% confidence intervals; the estimates in black are distinguishable from zero at the $\alpha = 0.05$ level.

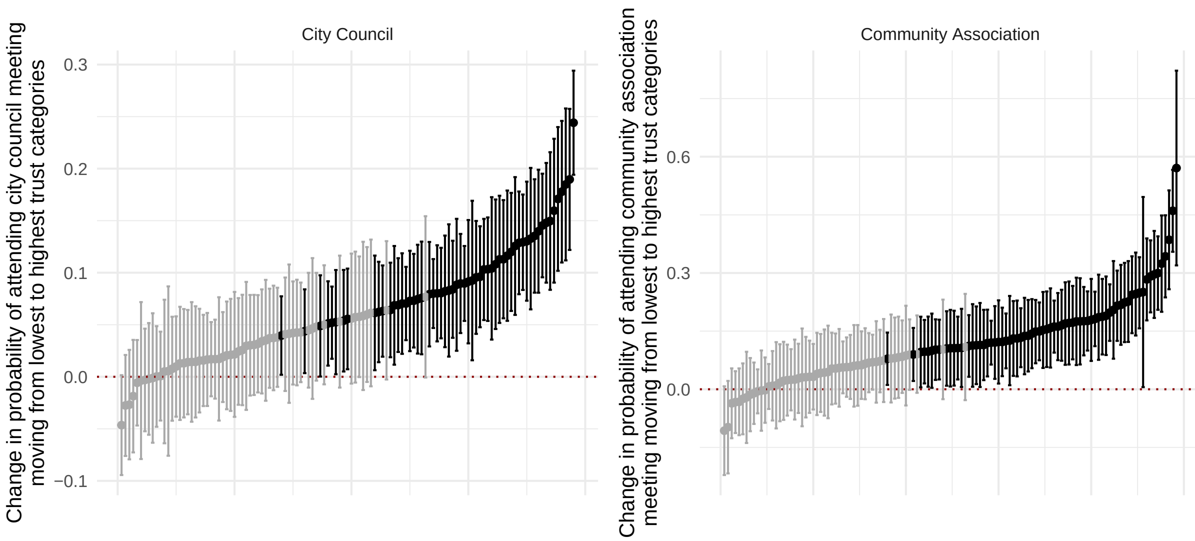


Simple descriptive analysis of existing data can therefore inform decisions over investment in public meetings.

Conducting this exercise ourselves, we find both that (1) the preaching-to-the-choir problem is not specific to our meta-study or to police–community meetings, and that (2) the magnitude of the problem varies significantly across contexts. Figure 9 plots estimates of the difference in meeting attendance rates between the highest and lowest categories of trust in two other groups: city government and members of the community. 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 five 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 in 106 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%.

These positive correlations are not obvious ex-ante: one might think that public meetings would attract more critics than fans, serving as fora for expressing frustration, airing grievances, or making demands. The consistency of the positive correlation 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. At the same time, the tremendous variance in the magnitude of the relationship imply that some contexts do present an opportunity for government officials to win hearts and minds at public meetings.

This is not to say, of course, that building trust is the only or even the primary objective of participatory interventions (though it was the primary pre-registered objective of our intervention). Many participatory fora seek primarily to reshape public budgets or reallocate public health resources, for example. Our finding is most relevant for the many participatory interventions that do focus on winning hearts and minds. But it is also relevant, if less directly, for other participatory governance initiatives. Our analysis of the consequence of selection adds structure to the oft-cited concern that participatory governance may simply serve the interests of privileged citizens with the leisure time to participate Our application is specific, but two takeaways are general: (1) we cannot understand the consequences of participatory fora without considering selection, and (2) policymakers can often anticipate selection by analyzing widely available data.

A different approach to avoiding the preaching-to-the-choir problem suggests imposing rather than inviting non-enforcement contact with officers. Peyton, Sierra-Arévalo, and Rand (2019) and Karim (2020) describe community policing interventions in which police knock on residents' doors. We have normative and empirical concerns with this approach. Our measures of uptake suggest that many citizens prefer not to engage with police, and that they may even be harmed by this type of contact. Imposition of police presence entails risks for residents and may also entail risks for police officers. Our proposed solution instead preserves the benefits of opt-in, voluntary non-enforcement contact while minimizing the costs of ineffective interaction with constituents who already trust the police.

## 6 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 more than 500 meetings were held and that 8% of the adult population attended (in treated neighborhoods). We attribute this failure in part to positive selection into meetings: those who liked the police to begin were more likely to attend and less likely to be impressed by what they saw. Officers ended up preaching to the choir.

Our findings suggest an additional cause for the null results in the larger meta-study (Metaketa) that includes our experiment. Blair et al. (2021) pin the null on implementation failures, citing limited commitment to community policing in the Global South. But rational politicians or rational police leaderships should only prioritize community policing to the extent that it generates positive outcomes.

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 policy-making and implementation, or "flattening access" (Grossman, Humphreys, and Sacramone-Lutz, 2020). But when selection into participation is non-uniform, participatory institutions may be highly inefficient (as in the present study) or even detrimental (Hanson, 2018; Slough, 2022*a*). Better understanding of this citizen selection into participation can inform the design of more effective public institutions.

## 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." *Amnesty International* .

Ba, Bocar A, and Roman Rivera. 2019. "The Effect of Police Oversight on Crime and Allegations of Misconduct: Evidence from Chicago." Faculty Scholarship at Penn Law, `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." *The Academy of Management Review* 23 (3): 459–472.

Blair, Graeme, Jeremy M Weinstein, Fotini Christia, Eric Arias, Emile Badran, Robert A Blair, Ali Cheema, Ahsan Farooqui, Thiemo Fetzer, Guy Grossman et al. 2021. "Community policing does not build citizen trust in police or reduce crime in the Global South." *Science* 374 (6571): eabd3446.

Blattman, Christopher, Gustavo Duncan, Benjamin Lessing, and Santiago Tobon. 2022. "State-building on the Margin: An Urban Experiment in Medellín." NBER Working Paper 29692 `10.3386/w29692`.

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): 1455–1485.

Casey, Katherine, Rachel Glennerster, and Edward Miguel. 2012. "Reshaping institutions: Evidence on aid impacts using a preanalysis plan." *The Quarterly Journal of Economics* 127 (4): 1755–1812.

Castillo, Juan Camilo, and Dorothy Kronick. 2020. "The Logic of Violence in Drug War." *American Political Science Review* 114 (3): 874–887.

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* .

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): 902–916.

Dipoppa, Gemma. 2021. "How Criminal Organizations Expand to Strong States: Migrant Exploitation and Political Brokerage in Northern Italy." Working Paper, https://www.dropbox.com/s/4egwnanyfkpe11b/Dipoppa_JMP.pdf?dl=0.

Einstein, Katherine Levine, David M Glick, and Maxwell Palmer. 2019. *Neighborhood defenders: Participatory politics and America's housing crisis*. Cambridge University Press.

Esberg, Jane, and Jonathan Mummolo. 2018. "Explaining Misperceptions of Crime." Working Paper, https://scholar.princeton.edu/sites/default/files/jmummolo/files/esberg_mummolo_ms_share.pdf.

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. "A Dark Side of Civic Engagement." *Civic Engagement in American Democracy* pp. 395–425.

Garcia, Juan, Daniel Mejia, and Daniel Ortega. 2013. "Police Reform, Training and Crime: Experimental Evidence from Colombia's Plan Cuadrantes." Fundación Ideas para la Paz, Working Paper Series, https://ideas.repec.org/p/col/000089/010497.html.

Goldfrank, Benjamin. 2007. "The Politics of Deepening Local Democracy: Decentralization, Party Institutionalization, and Participation." *Comparative Politics* pp. 147–168.

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

Gonzalez, Yanilda, and Lindsay Mayka. 2022. "Policing, Democratic Participation, and the Reproduction of Asymmetric Citizenship." *American Political Science Review* pp. 1–17.

González, Yanilda María. 2022. "Reforming to Avoid Reform: Strategic Policy Substitution and the Reform Gap in Policing." *Perspectives on Politics* Firstview: 1–19.

Greene, Jack R. 2000. "Community policing in America: Changing the nature, structure, and function of the police." *Criminal justice* 3 (3): 299–370.

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." *The Journal of Politics* 82 (4): 1321–1336.

Hanson, Rebecca. 2018. "Deepening distrust: Why participatory experiments are not always good for democracy." *The Sociological Quarterly* 59 (1): 145–167.

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*. Russell Sage Foundation.

Holland, Alisha C. 2013. "Right on crime? conservative party politics and" mano dura" policies in el salvador." *Latin American Research Review* pp. 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): 203–210.

Humphreys, Macartan, Raul Sanchez de la Sierra, and Peter Van der Windt. 2015. "Social Engineering in the Tropics: A Grassroots Democratization Experiment in the Congo." Working Paper, `http://www.columbia.edu/~mh2245/papers1/20150519%20HSW.pdf`.

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): 536–551.

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): 1532–1546.

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 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): 552–572.

Mansuri, Ghazala, and Vijayendra Rao. 2013. "Can participation be induced? Some evidence from developing countries." *Critical Review of International Social and Political Philosophy* 16 (2): 284–304.

Mayka, Lindsay. 2019. *Building participatory institutions in Latin America: reform coalitions and institutional change*. 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 *International Conference on Applications of Natural Language to Data Bases/Information Systems*. Springer pp. 214–219.

Mummolo, Jonathan, and Erik Peterson. 2017. "How content preferences limit the reach of voting aids." *American Politics Research* 45 (2): 159–185.

Osse, Anneke, and Ignacio Cano. 2017. "Police deadly use of firearms: an international comparison." *The international journal of human rights* 21 (5): 629–649.

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): 19894–19898.

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

Slough, Tara. 2022*a*. "Oversight, Capacity, and Inequality." Working Paper `http://taraslough.com/assets/pdf/oci.pdf`.

Slough, Tara. 2022*b*. "Squeaky Wheels and Inequality in Bureaucratic Service Provision." *Working Paper,* *`http://taraslough.com/assets/pdf/colombia_audit.pdf`* .

Trejo, Guillermo, and Sandra Ley. 2020. *Votes, drugs, and violence: The political logic of criminal wars in Mexico*. Cambridge University Press.

Tyler, Tom, and Jonathan Jackson. 2013. "Future challenges in the study of legitimacy and criminal justice." *Yale Law School, Public Law Working Paper* (264).

Tyler, Tom R. 2006. "Psychological perspectives on legitimacy and legitimation." *Annu. Rev. Psychol.* 57: 375–400.

Ungar, Mark, and Enrique Desmond Arias. 2012. "Reassessing community-oriented policing in Latin America." *Policing & Society* .

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): 341–362.

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): 1213–1229.