Class and the Development of Trust in Police in Latin America

Tara Slough Carolina Torreblanca

July 27, 2023

Abstract

Citizen trust in police is thought to facilitate the cooperation between communities and police. Large disparities in trust in police by race and socioeconomic status in the US are often assumed to translate to other contexts. We show that this generalization—made by regional and subject matter experts alike—is descriptively inconsistent with a large body of survey evidence in Latin America, world's most violent region. Drawing upon 147 unique cross-sectional surveys spanning 20 countries and two panel surveys, we show that measures of social class—thought to be strong predictors of treatment by police—have a weak negative correlation with individuals' trust in police. We measure and trace the origins of this variation in trust in police. In so doing, we propose a new framework for dynamic updating on police trustworthiness that organizes different explanations for variation in trust in police, producing implications for the design of policing policies to promote public safety.

A durable empirical regularity in the United States holds that citizens' race and class correlate with their beliefs about and attitudes toward the police. For example, in a 2021 Gallup survey, 56% of white respondents and 27% of Black respondents expressed "quite a lot" or "a great deal" of confidence in the police (Gallup, 2021). Such disparities in trust or confidence in the police, assessments of police service quality, and attitudes toward police have been observed for decades (Decker, 1981; Skogan, 2005; Macdonald and Stokes, 2006). Conventional wisdom holds that groups with worse experiences of the police—whether more abuse, worse quality of service, or poorer security outcomes—tend to trust the police less. In this paper, we document that this conventional wisdom about police service quality and trust from the United States does not travel to Latin America. In so doing, we propose a new account of the formation of trust in police.

Latin America and the Caribbean suffer the highest rates of crime and violence of any region (Muggah and Tobon, 2018; Vilalta, 2020). These high crime rates impose substantial welfare costs (Jaitman et al., 2015) and, in public opinion surveys, regularly register as important concerns of citizens (LAPOP, 2022). The primary institutional responses of the state to problems of crime and violence are the police and the justice system. But these state agents do not operate in a vacuum, in isolation from the citizens they police and serve. In this paper, we seek to more accurately characterize one dimension of the relationships between citizens and police in order to derive new implications for states' abilities to address crime and violence.

We focus on citizen trust in the police. Refining general definitions by Hardin (2003) and Bhattacharya, Devinney, and Pollutla (1998), we conceptualize trust in the police as a citizen's *belief* that a representative police agent will take an action that produces a beneficial outcome for the citizen. Citizen trust in police is important because many security outcomes are co-produced between citizens and police officers. Indeed, Blair et al. (2021: p. 1) assert that by generating more trust, interventions like community policing can "build more effective police agencies in environments of low trust." In these environments, trust is posited to promote information sharing (e.g., crime tips), which should increase the ability of police to locate, stop, remedy, investigate, or preempt crimes or misdemeanors. To the extent that community-policing interventions that seek

to build trust in police have occupied a central role in policing in Latin America (as elsewhere), it is important to understand this key outcome—trust in police—across the region.

Drawing on a growing body of literature on policing in Latin America, we evaluate a widely-held assumption that trust in police is increasing in social class or socioeconomic status. Latin America is the world's most unequal region (Hoffman and Centeno, 2003; Gasparini and Lustig, 2011) where social class has long been considered a salient social cleavage.[1] Work on police abuse and repression suggests that these behaviors target poor and marginalized communities (González, 2020; González and Mayka, 2022). Moreover, anti-poor bias is believed to be pervasive in many Latin American justice systems (O'Donnell, 1999; Brinks, 2007, 2019). Following the logic from the United States, that poor treatment and bad security outcomes reduce trust in police, we would expect trust in police to be *increasing* in socioeconomic status.

Leveraging 147 surveys from 20 countries in the region, we do not find support for this hypothesis. Indeed, pooling the 235,230 individual responses from all countries, we estimate that the correlation between income and trust in police is -0.053, and the correlation between education, a proxy for social class with less missingness, is -0.084. The negative correlation suggests that low-socioeconomic status individuals trust the police slightly *more* than their high-socioeconomic counterparts. Both correlations are statistically distinguishable from zero. Disaggregating across countries, we do not find a statistically significant positive correlation between either measure of socioeconomic status and trust in police in *any* of the 17 Latin American/Caribbean countries for which we have data. This represents a strong departure from decades of findings from the US, where income and the same trust question (in English) exhibit a *positive* correlation of 0.13.[2]

These results are surprising in light of conventional wisdom about the relationship between policing outcomes and trust in police. We measure the degree to which two groups of experts anticipated these findings through a forecasting exercise (DellaVigna, Pope, and Vivalt, 2019).

[1] Other identity-based cleavages (e.g., ethnicity or race) in Latin America vary more substantially across countries.

[2] In the US, the *racial* disparity between Black and White respondents is approximately twice the magnitude of the difference between respondents in the top and bottom income decile.

Our two groups of experts consisted of a sample of scholars of politics and public administration in Latin America and a sample of activists working on issues related to justice in Mexico. We show that the modal expert and the average forecast anticipated a positive correlation between socioeconomic status and trust in the police. Interestingly, inaccuracies in forecasts stem from underestimates of trust by *poor* respondents (those at the 10^{th} percentile) only. Respondents are quite accurate in their assessment of trust by *median* and *rich* respondents (those at the 50^{th} and 90^{th} percentiles, respectively). This exercise affirms that our findings challenge conventional wisdom, not only from the US but also from subject-matter and regional experts.

Why do our results depart so substantially from existing understandings of the correlates of trust in police derived from decades of research in the US? We explore but find no evidence for several explanations premised on the measurement of trust in police and conceptualization of trust. Using a variety of bounding approaches and ancillary analyses related to questions of institutional trust, we do not recover evidence that systematic measurement error drives our findings. We then consider the possibility that our concept of trust is mischaracterized. We examine whether trust in institutions is a fixed *trait* rather than a belief by examining the intra-cluster correlation in trust across multiple institutional trust measures (e.g., Ojeda, 2016; Mondak et al., 2011). Further, we ask whether trust might be picking up *preferences* for the role of police by considering respondent preferences over *mano dura*, which may be correlated with income. Here, we find that, if anything, preferences for *mano dura* are stronger among the rich and predict *lower* trust in police across the income distribution.

Given our conceptualization of trust as a *belief*, we consider multiple outcomes that citizens may use to learn about the trustworthiness of police. We show that crime victimization, bribe solicitation, and general feelings of insecurity are associated with lower trust in police across the income distribution. These outcomes—signals from which citizens might learn about the police—appear to yield patterns of learning that cohere with directional predictions produced by a model of Bayesian updating (Bhattacharya, Devinney, and Pollutla, 1998). Yet, two features of the data are inconsistent with a purely Bayesian account. First, the magnitude of updating (as a function

of socioeconomic status) is largely invariant to (average) citizen priors. Second, richer citizens report higher levels of victimization, in contrast to what standard characterizations of policing and violence in the region would anticipate. Less surprisingly, the rich report higher rates of bribe solicitation than the poor.

Comparing our findings to those from the well-studied US case where trust *increases* in socioeconomic status (and whiteness), we ask whether a similar or different process generates trust in the Latin American context. This is ultimately a question about the external validity of canonical findings from the US (Slough and Tyson, 2023). To organize our findings, we provide a simple model of the evolution of trust to understand the conditions under which these divergent patterns can be generated by a common set of theoretical mechanisms. We argue that the process by which policing outcomes or police behavior are perceived as informational signals about police quality varies by socioeconomic status in Latin America. We show how this mechanism—but not other more conventional alternatives like bias in police behavior or differential quality of police forces—is consonant with the survey evidence we marshal.

This paper contributes to the extensive literature on institutional trust (Levi and Stoker, 2000; Citrin and Stoker, 2018). We build on the insight that trust is cognitive (Bhattacharya, Devinney, and Pollutla, 1998; Hardin, 2003) by using data to isolate different attributes of citizens' updating on police agent trustworthiness. We show strong evidence that citizens of all socioeconomic statuses update in a roughly Bayesian manner (on average). By proposing a theoretical framework for thinking about the development of trust as a dynamic process, we offer a richer set of explanations for variation in trust in police that incorporate bias, police quality, experience, and behavioral features of citizens. We argue that this framework provides analytical traction on between-citizen and between-context variation in institutional trust.

Our paper is primarily descriptive, responding to recent calls for larger-scale quantitative description (Gerring, 2012; Munger, Guess, and Hargittai, 2021). While recent randomized interventions have sought to *increase* citizen trust in police (Blair et al., 2021; Karim, 2020; Peyton, Sierra-Arévalo, and Rand, 2019), our goal here is to try to understand baseline levels of trust in po-

lice in Latin America. Because causal effects are ultimately differences from these baseline levels, a richer descriptive understanding of this important outcome has two benefits. First, we can better interpret causal estimands and associated welfare effects of these interventions. Second, for policymakers and police forces intent on increasing trust in police, we show that these interventions are less easily targetable (on the basis of income) than in the US context.

We make two further innovations that serve to advance our understanding of trust in police. First, we integrate expert forecasting to assess the state of knowledge of these outcomes. While recent literature advances the use of forecasting in experiments (DellaVigna, Pope, and Vivalt, 2019; DellaVigna, Otis, and Vivalt, 2020), we show how these tools can be used in service of description. We argue that the forecasting data for description disciplines a reliance on heuristics and allows us to identify blindspots among scholars. Second, we show how a model can be used to discipline discussions of external validity when empirical patterns differ across different settings.

1 Trust in Police: Concept and Context

1.1 Concept of institutional trust

Social scientists have devoted substantial attention to the measurement of citizens' trust in government or governmental institutions (Levi and Stoker, 2000; Citrin and Stoker, 2018). Nevertheless, the definition of "trust"—and, indeed, the possibility of institutional trust—remains contested (PytlikZillig and Kimbrough, 2016). Building upon the conceptualizations of Hardin (2003) and Gerbasi and Cook (2009), we conceive of institutional trust as being *cognitive* and *relational*.

By cognitive, we mean that trust is fundamentally a *belief*. Bhattacharya, Devinney, and Pollutla (1998: p. 465) characterize trust as "an expectancy of positive (or nonnegative) outcomes one can receive based on the expected action of another party in an interaction characterized by uncertainty." In other words, individuals hold a belief about how the other party is likely to act or behave. In contexts of policing, this could be an individual's expectation about how police will treat them or an expectation of whether and how police will respond to a tip about a crime or misdemeanor.

By relational, we emphasize interactions between two parties, citizens and police agents. This is implicit in the formulation of trust by Bhattacharya, Devinney, and Pollutla (1998). We contend that beliefs are formed—and thus can be changed—by observation of the other party's behavior. In the context of citizen trust in police, thus, observed behavior by the police serves as a signal that provides some information about police trustworthiness. This signal can be used to update a citizen's belief about how the police might behave toward them in subsequent encounters. An unsavory encounter with a police agent, for example, can lead citizens to negatively update about the trustworthiness of police, in general, thereby reducing trust in police.

Hardin (2003) was skeptical of whether institutional trust is possible, largely because of limits to citizens' ability to form relationships with an institution. Whereas citizens may be able to interact with individual police agents, Hardin (2003) argues, it is nonstandard to think of an institution as an actor with whom these interactions might take place. We argue that interactions/relationships between individual citizens and individual police agents shape trust in police. In this context, citizens can hold beliefs about whether an individual officer is trustworthy. Moreover, they can make assessments about the share of trustworthy officers on a police force (or in a given police unit). Institutional trust is, therefore, both cognitive and relational.

It is useful to clarify two alternatives distinct from our institutional trust concept. First, some authors view or evaluate political trust as a trait (e.g., Ojeda, 2016; Mondak et al., 2011). Individuals from different groups may have different baseline propensities to trust other individuals or agents of institutions. If this were the case, environmental or genetic traits could confound the relationship between social class and trust. Alternatively, social/political trust may facilitate economic advancement, thereby increasing an individual's social class (Putnam, Leonardi, and Nanetti, 1993). Second, individuals undoubtedly hold varying preferences over what police *should* do or how the institution should function. The accounts of motivated reasoning or inference forwarded by Kunda (1987) and Taber and Lodge (2006) suggest that these preferences may affect how citizens form beliefs about police trustworthiness. While it is, of course, possible that preferences condition updating processes (Little, Schnakenberg, and Turner, 2022), we contend that

experiences with the police—good or bad—shape future expectations in a similar direction.

1.2 Policing and class in Latin America

Despite Latin America's regional turn towards democracy, police forces routinely engage in corrupt and abusive behavior (Macaulay, 2012; Magaloni, Franco-Vivanco, and Melo, 2020; Johnson, Mendelson Forman, and Bliss, 2012). Yet, experiences with police are far from uniform: research indicates that police forces can behave repressively toward lower-income individuals and individuals from marginalized groups while being responsive to the demands of privileged community members (González, 2020; González and Mayka, 2022). Additionally, regional scholars have pointed out bias against poor, indigenous, and other marginalized communities in the region's justice systems that lingers even after recent reforms (O'Donnell, 1999; Brinks, 2007, 2019).

While citizens' opinions of police in the region tend to be by and large negative (Malone and Dammert, 2021; Cao and Zhao, 2005), the disparate quality of policing could mask substantial heterogeneity in attitudes across social classes. However, the influence of socioeconomic status on trust in police is not straightforward to predict because class can affect exposure to police through different channels. Crime victimization, for instance, is negatively associated with trust (Cruz, 2015; Gaviria and Pagés, 2002), and positively with support for "*mano dura*" policing (Visconti, 2020), while feeling unsafe in your neighborhood correlates with more support for spending on the police (Altamirano, Berens, and Ley, 2022). To our knowledge, no study yet has systematically analyzed how support for police covaries with class across Latin American countries. Nevertheless, existing accounts provide support for a common premise: socioeconomic status predicts an individual's exposure to policing or the outcomes of policing. Our concept of trust suggests that different experiences with police agents or crime outcomes should provide different opportunities for learning about the trustworthiness of the police.

2 Research design

Our primary question is descriptive: how does trust in police vary in social class? Accordingly, we estimate the correlation between measures of socioeconomic status and reported trust in police. We

view these correlations as important in characterizing citizen-police relationships in Latin America and elsewhere.

Correlation speaks directly to our characterization of institutional trust as a belief. As a belief, trust is not an outwardly-observed characteristic, but individuals' level of trust in the police predicts at least some citizen behaviors toward the police. For example, Hanson, Kronick, and Slough (2022) show that citizen behavior toward police—in the context of community-police meetings—does vary in levels of trust in Medellín, Colombia. In the context of interactions with citizens, it is plausible that police officers may want to ascertain a citizen's level of trust when deciding how to engage. Because beliefs are unobserved, police may use observable characteristics to infer a citizen's trust and their likely behavior, a form of statistical discrimination (Phelps, 1972). In contrast to levels of trust, in Latin America, socioeconomic status is typically easily observable through an individual's dress, way of speaking, comportment, and surroundings (Britto Ruiz and Ordóñez Valverde, 2005; Sabatini, 2006; Villarreal, 2010).

2.1 Data

Our principal data source is LAPOP's AmericasBarometer (LAPOP, 2022). The sample incorporates the responses of 241,626 individuals collected from 147 unique surveys in 20 Latin American countries between 2004 and 2019. Each survey round was designed to be representative of the country's voting-age population that year. Table A1 in the Appendix lists the years and countries where the surveys were collected.

We use respondents' self-reported income bracket to measure socioeconomic status or class. We supplement the measure with individuals' educational attainment in years of schooling to assuage concerns that systematic misreporting of income might drive any results. In each round, respondents are asked about their trust in several institutions. We measure trust in the police with the question "To what extent do you trust the Police?" Responses range from 1 (not at all) to 7 (a lot).[3]

We complement the (repeated) cross-sectional data from LAPOP with two smaller panel sur-

[3]We report all survey questions and response scales in Table A4.

veys that measure trust in police: a five-wave nationally representative survey from Chile (COES, 2022) and a two-wave representative survey from Medellín, Colombia conducted in 2018 and 2019 (Hanson, Kronick, and Slough, 2022). While these surveys cover just two settings (Chile and Medellín, respectively), they allow us to examine within-individual variation over time.[4]

2.2 Estimation

We use ordinary least squares (OLS) to estimate the correlation between trust in police and socioeconomic status, as detailed in Equation 1. We regress individual i's self-reported trust in the police, Trust_i, on a self-reported measure of class (income or education), Class_i. Both trust and class measures are standardized within country-year to account for secular trends. We cluster standard errors at the level of each survey's primary sampling unit. Because the dependent and the independent variables are standardized, β is the estimator of the correlation coefficient.

$$\text{Trust}_i = \alpha + \beta \text{Class}_i + \epsilon_i \tag{1}$$

3 Baseline results

We present estimates of the correlation between socioeconomic status and trust in the police for the entire pooled sample of respondents in Figure 1. In addition, we also plot country-specific correlations. Contrary to conventional wisdom, the overall pooled correlation and all but one country-specific result are close to zero and slightly negative. The estimated correlation for the pooled sample is only -0.053 [95% CI: -0.059, -0.046] when class is operationalized as income and -.084 [95% CI: -0.089, -0.078] when we use education as a proxy. El Salvador has the most negative country-specific correlation, with an estimated correlation of -0.18 [95% CI: -0.20, -0.-0.16] when class is operationalized as education. For the rest of the countries, the correlation oscillates between -0.15 and .01. In the cases of Argentina, Chile, Costa Rica, Ecuador, Jamaica, and Peru, at least one correlation estimate is not statistically distinguishable from zero.

[4]While other existing longitudinal datasets cover other Latin American countries, we are unaware of any other that include repeated questions about policing.

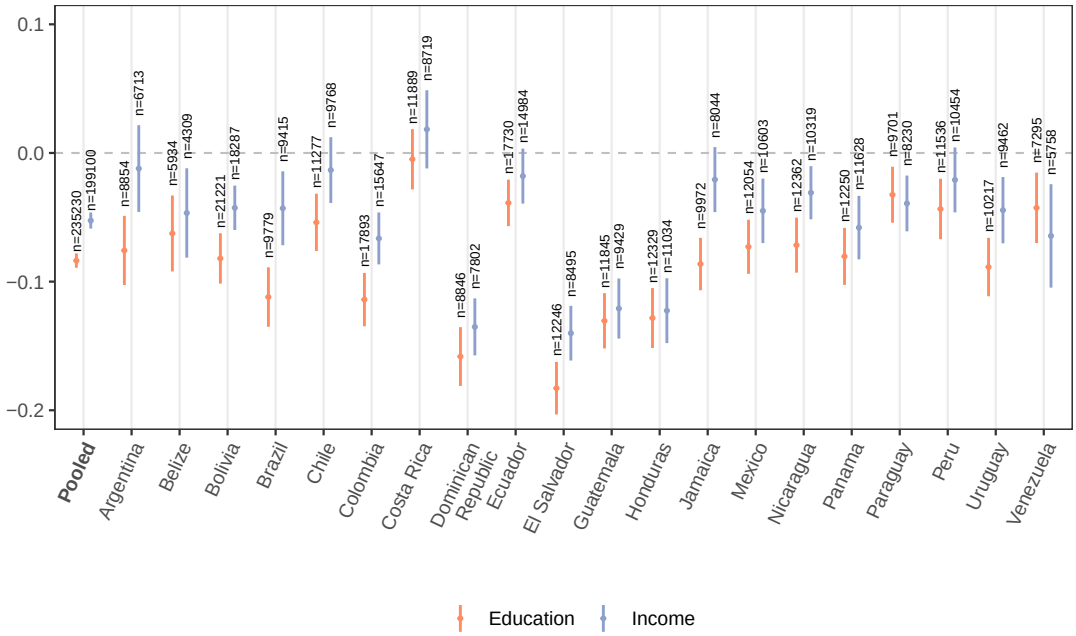

Figure 1: Correlation between LAPOP respondents' self-reported trust in police and two class measures: income (in blue) and education (in orange).

Correlation measures a linear relationship. Do the weak, negative correlations reported in Figure 1 mask a stronger, non-monotonic relationship between socioeconomic status and self-reported trust in the police? To explore the possibility, we divide respondents into class deciles and plot the mean level of trust in the police for the members of each decile. As before, we compute the means for the entire pooled sample and each country. Figure 2 reveals no evidence of a non-monotonic relationship between the two variables: the mean levels of trust are stable or decreasing only slightly in income or education in all countries and the pooled sample.

One limitation of our analysis is that we measure class using self-reported education and income measures, which are indirect measures of socioeconomic status. Two alternative measures may be preferable for different purposes. First, in some countries, there exist administrative classifications of class. For example, in Colombia, dwellings are categorized by *estrato* (strata) to prorate public utility charges. Drawing upon the original panel survey in Medellín, we estimate a correlation between dwelling *estrato* and trust in the police of 0.068 (95% CI: [0.029, 0.107]). This correlation is somewhat higher than estimates based on self-reported education and income

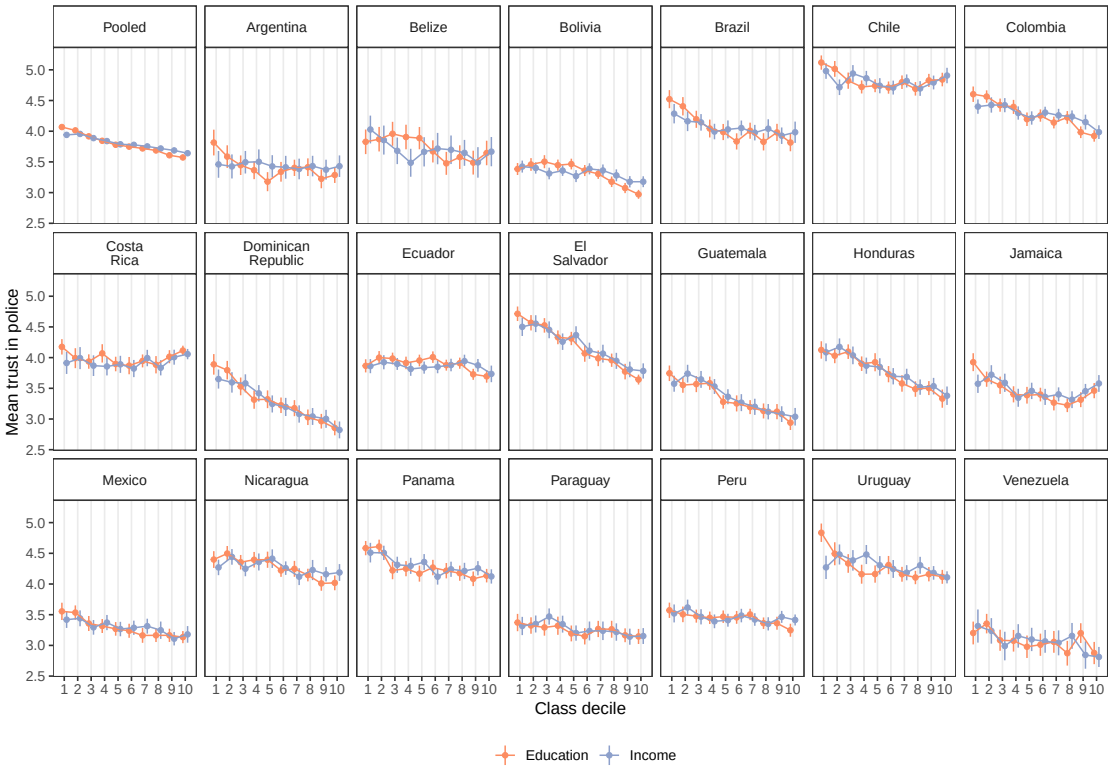

Figure 2: Mean trust in the police by decile of income (in blue) and education (in orange).

of -0.068 (95% CI: [-0.108, -0.028]) and 0.006 (95% CI: [-0.035, 0.046]), respectively. While the correlations from Medellín are somewhat higher than the Colombian averages reported in Figure 1, rescaling the national correlation estimates by a magnitude similar to the difference between estimates using self-reported class and the administrative *estrato* classification from the Medellín survey would preserve a weak negative correlation with trust in police. Second, citizens may identify with a higher or lower class than their income or education would suggest. From panel data in Chile, in which class identification is measured subjectively, we estimate a correlation of 0.038 (95% CI: [0.003, 0.073]) between self-identified class and trust in police. While this is the only positive correlation that we find which is (statistically) distinguishable from zero, it remains substantially weaker than correlations between class and trust in police observed in the US. These ancillary surveys suggest that our findings from the widely-available LAPOP proxies of class do not substantially mislead relative to plausible alternative measures.

4 Expert Forecasts

The finding that trust in the police covaries weakly and, in general, negatively in socioeconomic status in Latin America was surprising to us. We conducted an expert forecast elicitation with two samples to assess whether our findings were similarly surprising to other experts. Expert forecasts are increasingly used to measure experts' prior beliefs about quantities of interest in social science research (DellaVigna, Pope, and Vivalt, 2019).

Our two expert samples are (i) scholars of Latin American politics and (ii) activists working on issues of human rights and policing. We used a recent program of a Latin American politics conference known for the participation of scholars from all regions as our sample frame for the academic sample. One of the authors identified a network of activists through past non-academic employment related to policing in Mexico. All participants were invited to share the survey with other interested colleagues and students. In sum, we collected 121 country-level forecasts from 101 unique experts. Table A6 in the Appendix shows the count of forecasts by country and respondent type.

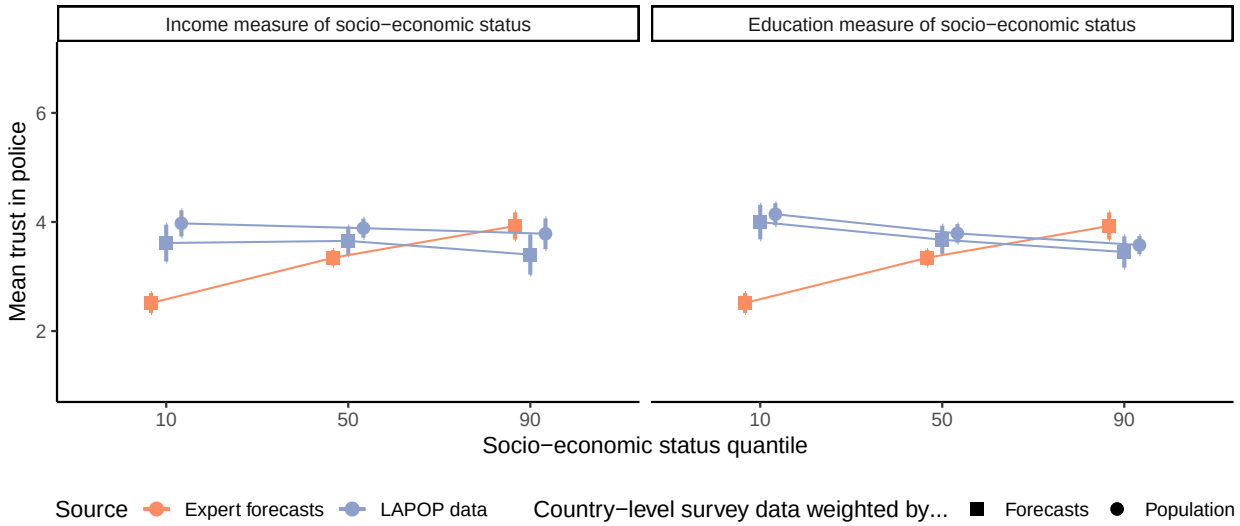

Figure 3: Divergence between average forecasts (in orange) and corresponding survey-based measures (in blue).

We asked experts to provide at least one *forecast* for one country in Latin America (or the region as a whole). A forecast consists of three quantities: mean levels of trust—per the LAPOP survey question—at the 10th, 50th, and 90th percentiles of household income. As depicted in Appendix A2, our forecasting instrument contextualized the income range by reporting average income at each level. Respondents had access to the text of relevant LAPOP questions in English and Spanish.

Figure 3 shows that, on average, expert respondents expected a positive correlation between socioeconomic status and trust in police. The average expert forecast is monotonically increasing from 2.51 [95% CI: 2.32, 2.69] for citizens in the first decile of income to 3.34 [95% CI: 3.20, 3.49] for citizens of median income, to 3.93 [95% CI: 3.68, 4.18] for citizens in the tenth decile of income. Because the forecast is an average over country-level forecasts, we report two measures from the survey data, drawn in blue. First, we weight country-level surveys by the prevalence of each country in the forecasts to ensure that frequency does not drive the divergence between forecasts and actual data with which countries are represented in the forecasts. We also weigh the survey data by each country's population, offering a more interpretable regional average. Both weighting schemes yield similar flat or slightly negative relationships between socioeconomic sta-

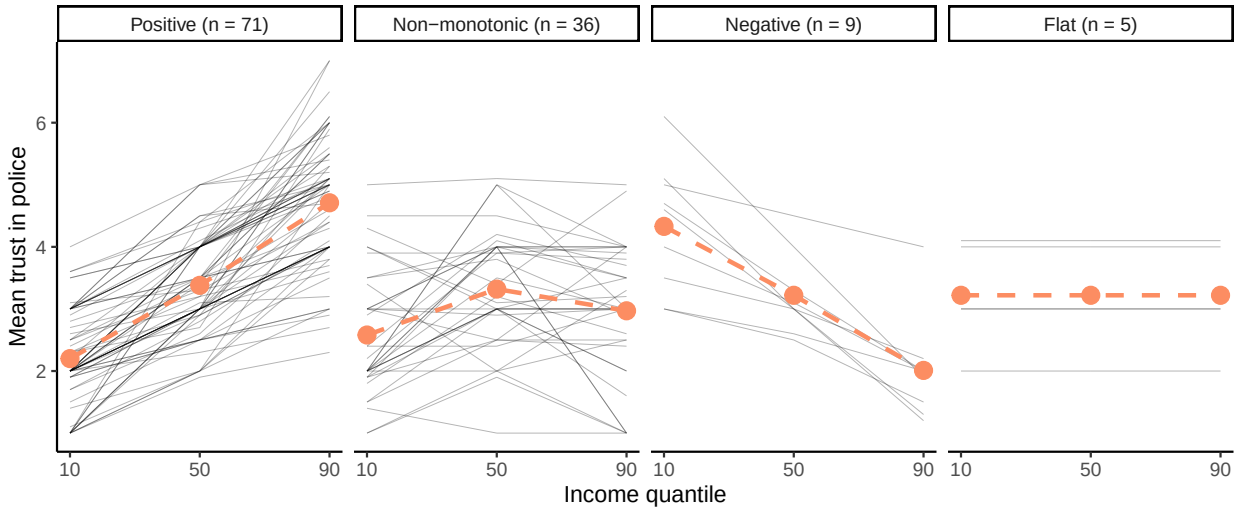

Figure 4: Classification of individual forecasts. Thin lines each represent individual forecasts. The points and dashed lines represent averages for each type of forecast.

tus and trust in police, in stark contrast to the forecasts. The right panel of the graph shows that using education rather than self-reported income to measure class in the LAPOP survey data does not change our qualitative finding that the positive relationship between income and crime anticipated by experts—including us—is not borne out in the data.

Two further disaggregations of the forecast data help to clarify experts' expectations. First, Figure 4 disaggregates four "types" of forecasts. As in Figure 3, the modal prediction suggested a *positive* relationship between socioeconomic status and trust in police. Additionally, 30% of forecasts posited a non-monotonic relationship. On balance, these forecasts suggested that a citizen with a median household income should trust the police more than the rich or the poor. Just 7.5% and 4% of forecasts posited a negative or flat relationship, respectively. Second, Figure A4 in the Appendix shows that the prediction of a positive relationship between income and trust in police is not specific to any country: we observe a similar pattern in all countries for which we have ≥ 8 predictions: Argentina, Brazil, Chile, Colombia, and Mexico.

5 What accounts for these patterns of trust?

We have documented that empirical patterns of trust in police depart from our expectations and those of other experts. We now seek to understand why this is the case. We consider four classes of explanations: measurement concerns; two alternative conceptualizations of institutional trust; and self-reported experiences with police.

5.1 Artifacts of measurement

One possible cause of the divergence between experts' beliefs and what LAPOP survey data shows could be how class and trust are measured. Missingness may be correlated with socioeconomic status and trust in police. For example, if low-socioeconomic-status individuals who do not trust the police are less likely to respond or high-socioeconomic-status individuals who trust the police are more likely to respond to the relevant questions, we would underestimate the correlation between socioeconomic status and trust in police. While 16.8% of the respondents provided no income information, only 2.2% refused to answer an institutional trust question. However, the average percentage of missingness in self-reported education for survey rounds is only around 1%. The similarity of correlations across both measures of socioeconomic status in Figure 1 suggests that missingness is unlikely to bias our inferences about trust in police systematically.

Additionally, we conduct a bounding exercise reported in Appendix A6 to assess the maximum extent to which missingness could bias our conclusions. To generate worst-case bounds, note that correlation is bounded between -1 and 1. Since both trust in police and our measures of socioeconomic status are coded as Z-scores in (1), we can calculate these worst-case bounds for any respondent that reports at least one of the two measures. For this subset of missing observations (99.8% of all missing responses), we impute the observed Z-score for the missing Z-score (such that the imputed observation lies on the 45° line) to generate the worst-case upper bound. We then impute the negative of the observed Z-score (such that the imputed observation lies on the -45° line) to generate the worst-case lower bound. We estimate the correlation with all observed and imputed observations to generate worst-case bounds for the correlations reported in Figure

1. We estimate a worst-case bound using the education measure of [-0.11, -0.05], suggesting that missingness can have only a minimal effect on our conclusions. The worst-case bounds for the income measure of social class are [-0.22, 0.13]. The width of these bounds is unsurprising, given the different degrees of missingness. Nevertheless, it is useful to note that this worst-case upper bound on trust in the Latin American case is equivalent to the estimated correlation between social class and trust in the United States (reported above), showing how different these empirical patterns are in practice.

Alternatively, one may be worried about bias in the semantic content of respondents' opinions. Suppose LAPOP survey respondents have no prior opinion about how much they trust the police or other institutions. In that case, they might offer meaningless answers to questions about institutional trust when prompted. To test this possibility, we leverage changes in the ideological orientation of a country's subsequent government and all the questions that ask after trust in different political institutions. If respondents answer questions on institutional trust randomly, on average, their answers should not vary systematically between authorities within a single year, nor should they respond to changes in the political environment.

We first identify nine spells of right-wing presidencies experienced by countries in the LAPOP data. Next, using a difference-in-difference-inspired design (with repeated cross-sectional data), we compare within-country the correlation between trust in each authority and socioeconomic status in the survey immediately before the right-wing spell and in the subsequent two survey waves. Figure A7 in the Appendix shows the results. Having a right-wing president increases the correlation between trust in the president and income, as we would expect if respondents' self-reported evaluation of trust did respond to changes in who political actors *are*. Further, the correlation between trust in other political institutions, like parties, congress, the police, and the army, is also positively affected by the incoming of a right-wing president, but to a lesser degree than trust in the presidency itself. In contrast, trust in the Supreme Court, an institution less directly dependent on the president, remains more stable. The evidence points to respondents' assessment of trust in authorities being meaningful.

Finally, we draw upon a finding from the Colombian panel survey and associated field experiment to suggest that survey-measured trust in police predicts subsequent engagement with police. Hanson, Kronick, and Slough (2022) show that respondents that report the top category of trust in the police (out of four categories) in a baseline survey are more than twice as likely to attend community-police meetings in beats assigned to treatment (18.4% vs. 8.4% of respondents). If survey measures of trust in the police were entirely random noise, we would not expect this alignment between reported trust and observed behavior.

5.2 Institutional trust as a trait

So far, we have shown that in Latin America, trust in the police does not vary in class in the way most experts predicted. The observed correlation between trust and class is generally negative but near zero. The alternative concepts of trust we described above may imply different predictions for the correlation between socioeconomic status and trust in police. If trust or propensity for trust were a trait rather than a relational expectation of police behavior, our findings could be explained by a weakly negative correlation between this predisposition and class.

To examine the possibility that our results are driven by stable individual differences in LAPOP respondents' *trust propensities*, we estimate the intra-class correlation between each respondent's trust in multiple institutions: the army, political parties, the sitting president, the supreme court, the national legislature, and the police. The intra-class correlation gives the ratio of between-respondent variance to the total variance in trust in these institutions. If the ICC were close to 1, it would suggest limited variance in an individual's assessment of multiple institutions, suggesting that institutional trust functions as a stable trait or predisposition. However, we do not observe a high ICC. Across all subjects in the LAPOP surveys, we estimate an ICC of .05 (95% CI: [.02,.23]). Disaggregating by country in figure A8, we show similarly small estimates in all countries. It is, therefore unlikely, that a stable individual-level predisposition to trust can explain away our results.

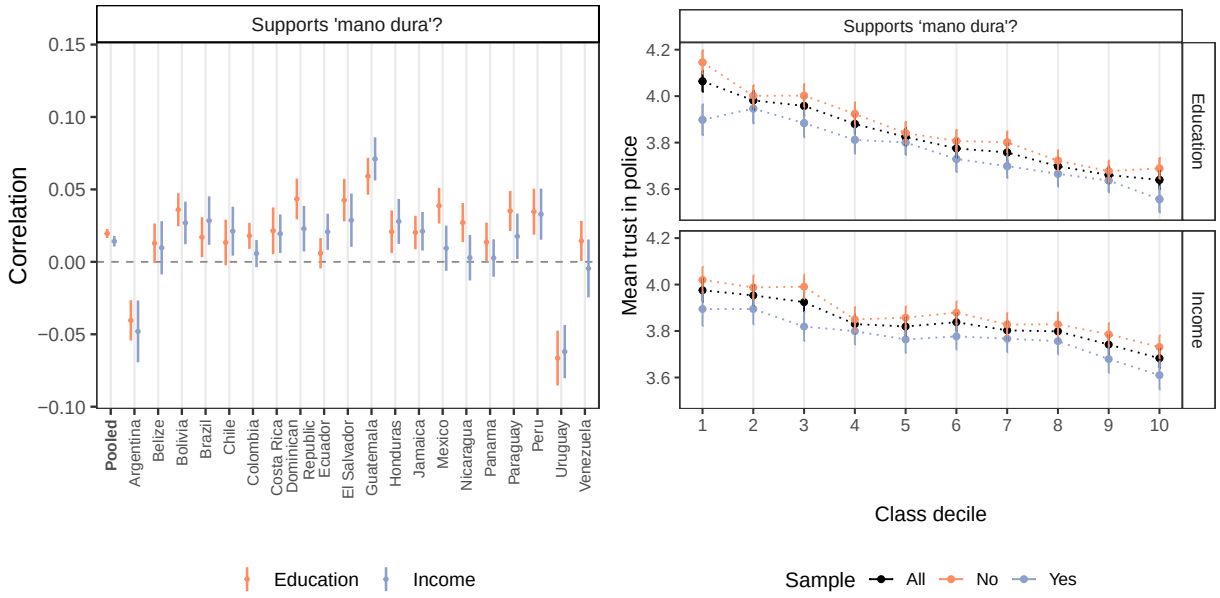

Figure 5: The left panel shows the estimated correlation between income (orange), education (blue), and support for tough-on-crime or *mano dura* policing. The right panel shows the predicted level of trust in police for the pooled sample, by class decile, as a function of support for *mano dura* (yes in blue/no in orange).

5.3 Beliefs vs. Preferences

We have argued that trust should be characterized as a belief, but skeptical readers may argue that it is, instead, a manifestation of a preference with regard to policing. Moreover, literature in psychology and political psychology argues for the plausibility of an interaction between (prior) beliefs and preferences in the updating process (Kunda, 1987; Taber and Lodge, 2006; Little, Schnakenberg, and Turner, 2022). Specifically, the evolution of trust could be subject to motivated reasoning or directional motives, whereby trust becomes a function of an individual's prior preferences regarding policing (Ibid.). If this were the case, a respondent who prefers a policy that necessitates active police involvement may be motivated to hold more positive views of the police, thereby generating higher levels of trust in police (all else equal).

To gauge if respondents' beliefs about police trustworthiness may be shaped by their prefer-

ences over policing practices or policy, we characterize the relationship between socioeconomic status, self-described support for tough-on-crime or *mano dura* policing, and trust in police. Two expectations are worth clarifying. First, a motivated-reasoning or inference account of updating on police trustworthiness holds that pro-*mano dura* individuals have higher trust in police. Second, given the generally negative (if weak) correlations between socioeconomic status and trust in police reported in Figure 1, we would expect that the poor hold more favorable views of *mano dura* policies. The figure counters both expectations.

First, the left panel in Figure 5 shows a close-to-zero and *positive* correlation between income and support for tough-on-crime policing across all countries except Argentina and Uruguay. This finding is in line with recent research showing there is a positive relationship between (i) crime victimization and support for 'mano-dura' (Visconti, 2020) and (ii) income and urban property crime victimization (Gaviria and Pagés, 2002). The right panel in Figure 5 shows the predicted level of trust in police by class decile as a function of respondents' self-reported support for 'mano dura.' The black line plots the expected level of trust in police for respondents in each decile, and the blue line plots the conditional expectation for respondents in that decile who are supportive of 'mano dura.' In contrast, the orange line plots the conditional expectation for respondents in that decile who are *unsupportive* of 'mano dura.' As we can see, the expected level of trust for individuals supportive of *mano dura* is lower than for individuals unsupportive of the measure across all income levels. Additionally, trust for both groups decreases at a similar rate. The results show the opposite empirical pattern we would expect to find if trust largely driven by individuals' preferences, discounting the possibility of a motivated-reasoning explanation of our results.

5.4 Experiences of police

Our concept of trust in police centers on citizens updating their beliefs about police trustworthiness based on their experience of police behavior or observation of policing outcomes. To this end, our survey data identifies one such experience—police solicitation of bribes (corruption)—and two policing outcomes—crime victimization and perceptions of safety—that allow for empirical investigation of these dynamics. While the experiences of crime victimization and feeling "un-

safe" do not necessarily require active interaction with police, they at least imply that police failed to prevent crime or inspire a feeling of security. To the extent that these observations or experiences serve as signals of police trustworthiness, we are interested in how such signals vary in socioeconomic class. We are first interested in the frequency or degree to which citizens of different socioeconomic classes experience different signals of police trustworthiness. We are then interested in whether these experiences differentially predict trust in the police for respondents of different socioeconomic statuses.

In Figure 6, we report the correlation between respondent socioeconomic status and the three (possible) binary signals of police trustworthiness. The left panel of the figure shows that socioeconomic status is correlated with higher rates of self-reported crime victimization in every country in the region. These correlations are statistically distinguishable from zero at the $\alpha = 0.05$ level for all countries and measures of socioeconomic status except for the income measure in Panama and Venezuela. Whether this is a surprising finding is unclear. Our evidence aligns with findings on urban property crime by Gaviria and Pagés (2002), albeit in a larger sample of urban/rural municipalities and with a broader range of crimes. On the other hand, poorer neighborhoods and municipalities are often distinguished by high rates of violent crime and insecurity. The panel survey from Medellín elicited victimization experience by type of crime, thus speaking to these questions. Figure A9 shows that exposure to violent crimes (especially homicide) is more common among the poor, whereas property crimes—which happen more frequently—disproportionately target the rich.

The second panel of Figure 6 suggests that lower socioeconomic-status respondents report feeling unsafe in their neighborhoods at higher rates. These estimates are noisier, and many country-level correlations are indistinguishable from zero. As such, the rate at which the poor and rich observe poor security outcomes is ambiguous in the survey data. The disparities between survey and administrative data leave open the possibility that what is perceived as an informational signal varies as a function of socioeconomic status.

The right panel of Figure 6 reports the correlation between socioeconomic status and self-

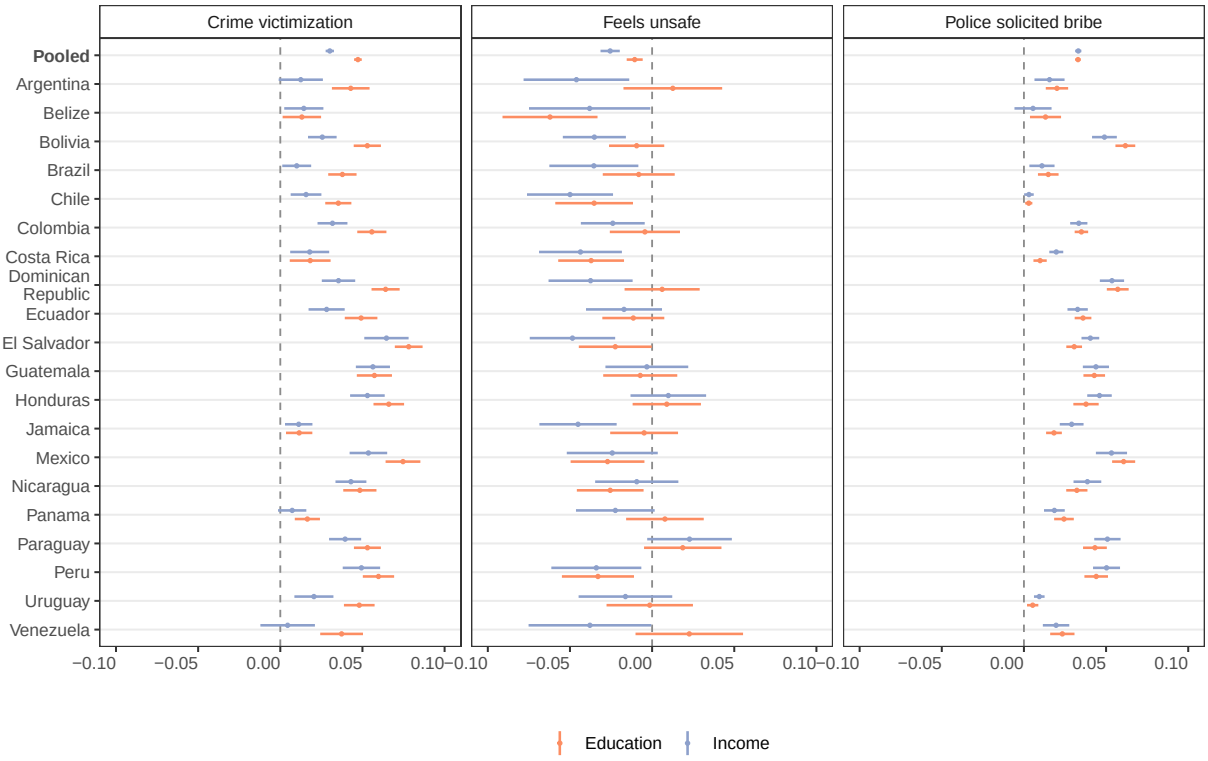

Figure 6: Pooled and country-specific correlations between income (in blue) and education (in orange) and each of the following three variables (from left to right): perceiving yourself as having been a victim of a crime during the past 12 months, perceiving your neighborhood as unsafe, and a police officer soliciting a bribe from you.

reports that a police agent asked a citizen for a bribe. The positive correlations suggest that across countries and class measures, the rich are asked for bribes more frequently. From the perspective of rent maximization, these are the citizens from whom police may be able to extract larger bribes. Collectively, the three panels of Figure 6 suggest that experiences with policing vary in social class, albeit higher self-reported crime victimization among the rich may be surprising relative to conventional wisdom.

Second, we show that exposure to these signals about the police is associated with the level of trust reported. Figure 7 plots the predicted level of trust in the police as a function of our three signals (denoted S_i) of police trustworthiness: (i) crime victimization in the past 12 months (left), (ii) feeling "unsafe" in their neighborhood (center), and (iii) whether a police officer asked for a bribe during the past 12 months (right). In each panel, the black line plots the mean level of trust, by decile, across the full sample. This line is very similar across all three vertical panels: the only (slight) differences come from variation in the presence of questions measuring the aforementioned signals across country-year survey waves. For each measure of socioeconomic status (the horizontal panels), trust decreases slightly and monotonically as income increases. We note that these means can be additively decomposed as follows:

$$E[\text{Trust in police}_i] = E[S_i = 1]E[\text{Trust in police}_i|S_i = 1] + E[S_i = 0]E[\text{Trust in police}_i|S_i = 0] \tag{2}$$

The orange and blue points and lines report our estimates of the conditional expectations in (2). The blue line ($E[\text{Trust in police}|S_i = 1]$) consistently falls below the orange line ($E[\text{Trust in police}|S_i = 0]$). This is consistent with our expectations—and the conventional wisdom—that poor security outcomes or treatment by police reduce trust in police. Importantly we see evidence consistent with updating—the difference in the orange and blue lines—for all deciles of socioeconomic status. The idea that poor treatment or outcomes reduce trust is consistent with many existing theoretical and empirical accounts of trust in government or government institutions (Hardin, 2003;

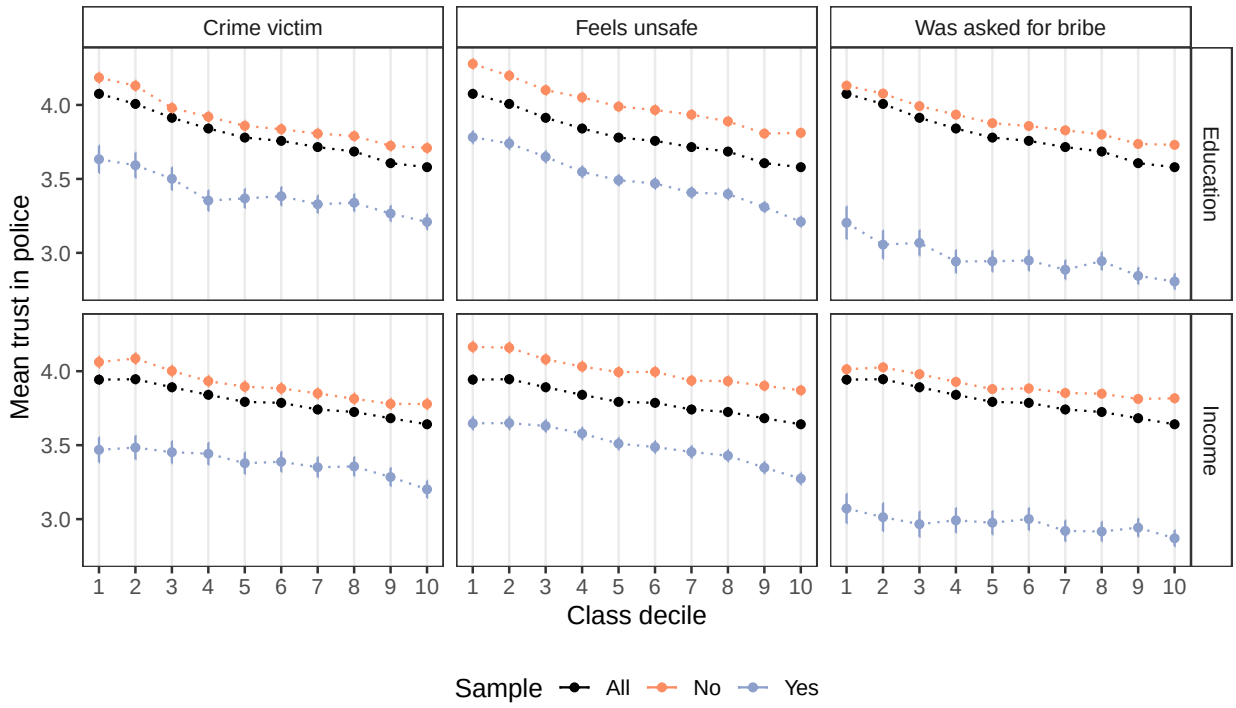

Figure 7: Predicted level of trust in police, by class decile, as a function of the following binary signals: crime victimization in the past 12 months (first panel), feeling 'unsafe' in their neighborhood (second panel), and a police officer asked for a bribe during the past 12 months (third panel). Model was fit on pooled data from all country-waves.

Levi and Stoker, 2000). We do not find evidence that poor, middle-class, and rich Latin Americans update according to fundamentally different cognitive processes. This analysis does not, for example, support claims that citizen rationality (in this domain) varies in education or socioeconomic status.

The distance between the black lines and the orange/blue lines reflects the share of respondents that experience a given signal (e.g., $E[S_i = 1]$ from (2)). We see that crime victimization and especially bribe solicitation are *rare* at all levels of socioeconomic status. This is evident because the black line is much closer to the orange line, the conditional means for citizens who did not observe the signal in the last year. Indeed, in the full sample, only 21.0% and 10.6% of respondents reported crime victimization or bribe solicitation in the last year, respectively. The correlations in Figure 6 suggest that these outcomes of policing are increasing in socioeconomic status, which

is evident from the growing distance between the black and orange lines as socioeconomic status increases. For example, moving from the first to the tenth decile of education corresponds to (reported) increases from 13.5% to 28.1% in crime victimization and from 5.7% to 16.3% in bribe solicitation. While these differences in exposure to poor police behavior do increase the (negative) gradient of socioeconomic status and trust in police by pulling the black line toward the orange line, we note that these differences in isolation do not account for the negative gradient of the orange and blue lines.

Our language in this section has veered closer to causal language. Ideally, we would describe the updating in Figure 7 as the *effect* of different signals of police trustworthiness on trust in police at different income levels. However, we lack a research design capable of supporting such an inference with the LAPOP data. However, the panel surveys from Medellín and Chile permit estimation of average treatment effects on the treated (ATTs) for several closely related signals. We employ a two-way fixed-effect estimator and the fixed-effect counterfactual estimator proposed by Liu, Wang, and Xu (2022). From the latter, we report the ATT that weights units equally. As in Figure 7, these signals are self-reported, though the questions vary slightly (as we report in Table A4). In Figure 8, we compare the estimated ATTs to associations (analogous to Figure 7) for the full sample of respondents. We show that all of the estimated ATTs are significant at the $\alpha = 0.05$ level and signed in the same direction as in the cross-sectional analyses from the Medellín, Chile, and LAPOP samples: feeling unsafe, crime victimization, and viewing the police as corrupt reduces trust in police. However, the magnitude of the ATTs relative to the pooled cross-sectional association is reduced by 49-81%. This suggests that the magnitude of findings in Figure 7 is likely overstated, but that updating on the basis of security outcomes or interactions with police leads to lower levels of trust in police.

6 A Framework for the Development of Institutional Trust

To summarize our empirical results to this point, we show that across Latin American countries, the correlation between social class and trust in police is weak and, in most countries, negative,

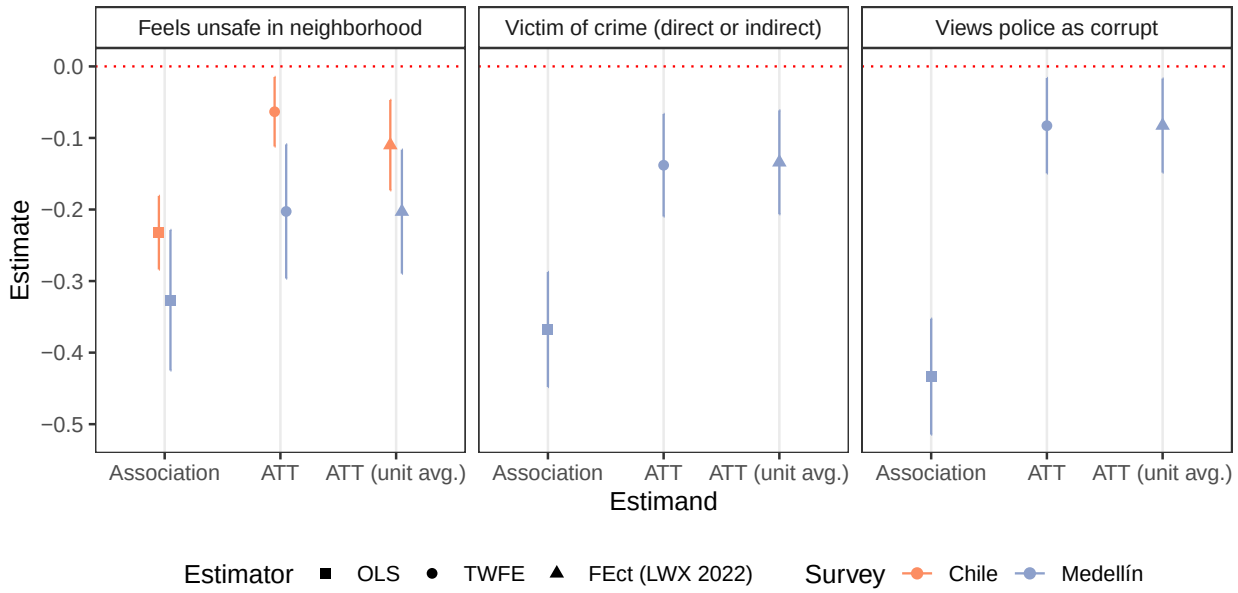

Figure 8: Estimates of pooled associations (across waves) to estimates of the average treatment effect (ATT) on the treated of signals analogous to those in Figure 7. LWX (2022) indicates the fixed effects counterfactual estimator proposed by Liu, Wang, and Xu (2022). 95% confidence intervals are calculated on standard errors clustered at the primary sampling unit.

meaning that poor (or less-educated) respondents report higher trust in police than rich (or more-educated) respondents. This finding counters conventional wisdom from the US as well as the expectations of regional experts. Despite this surprising finding, we show that respondents update in response to signals of security outcomes or police behavior as predicted by conventional wisdom: poor service or abuse by police seems to beget poorer assessments of police trustworthiness. More surprising, perhaps, the rich report higher levels of crime victimization, higher levels of police bribe solicitation, and similar assessments of neighborhood security. This means that the rich (self)-report a higher frequency of adverse signals of police quality than the poor.

We propose a simple theoretical framework to organize and synthesize these empirical results. Our framework builds on our conception of trust as cognitive. It seeks to illuminate how disparities in trust might emerge (or fail to emerge) within the course of citizen-police agent interactions. Ultimately, we are interested in characterizing citizen learning about police trustworthiness. Our measure of trust in police operationalizes citizen beliefs about the trustworthiness of the police, which we will refer to as μ_t, where t indexes time. Police agents are of type $\theta \in \{0, 1\}$, where

where $\theta = 1$ indicates trustworthy and $\theta = 0$ indicates non-trustworthy. We think of citizen beliefs as a subjective assessment of the pool of police agents, e.g., $\mu_t = \Pr(\theta = 1)$. As the share of trustworthy agents in the police increases, a citizen trusts the police more.

Citizens learn about police trustworthiness by observing police behavior or security outcomes. We denote a citizen's prior belief about the trustworthiness of the police as μ_{t-1}, which has an analogous interpretation to μ_t. At a given time, t, police choose an action $S_t \in \{0, 1\}$. Without loss of generality, we think of $S_t = 1$ as an action that harms the citizen's welfare (e.g., solicitation of a bribe or failure to prevent a crime). A trustworthy police agent chooses the action that harms the citizen, $S_t = 1$, with probability $\Pr(S_t = 1|\theta = 1) = \sigma_1 \in (0, 1)$, whereas the non-trustworthy police agent chooses $S_t = 1$ with probability, $\Pr(S_t = 1|\theta = 0) = \sigma_0 \in (\sigma_1, 1)$. We will assume that non-trustworthy police agents are more likely to take the action that harms the citizen, $\sigma_0 > \sigma_1$. A citizen who observes S_t can update according to Bayes' rule as follows:

$$\mu_t = \begin{cases} \frac{\mu_{t-1}\sigma_1}{\mu_{t-1}\sigma_1 + (1-\mu_{t-1})\sigma_0} & \text{if } S_t = 1 \\ \frac{\mu_{t-1}(1-\sigma_1)}{\mu_{t-1}(1-\sigma_1) + (1-\mu_{t-1})(1-\sigma_0)} & \text{if } S_t = 0. \end{cases}$$

The assumption of Bayesian updating is (directionally) consistent with our findings in Figure 7 for all socioeconomic strata when we examine how levels of trust vary as a function of crime victimization, bribe solicitation, and feelings of safety.

To think about citizen beliefs in the aggregate, we consider two additional features. The first, institutional quality, π, is the share of trustworthy agents among the police. Since citizens will generally hold heterogeneous beliefs about police trustworthiness—as in our data—given the updating process above, it is useful to parameterize the (objective) pool of agents. We will define $\pi = \Pr(\theta = 1)$. One direct consequence of this parameterization holds that:

$$\Pr(S_t = 1) = \pi\sigma_1 + (1 - \pi)\sigma_0 \qquad \text{and} \qquad \Pr(S_t = 0) = \pi(1 - \sigma_1) + (1 - \pi)(1 - \sigma_0)$$

The second additional feature, citizen attention, introduces a behavioral element to our model.

With probability $\gamma \in (0, 1]$, citizens observe police behavior or security outcomes. With complementary probability, they remain oblivious and do not update, such that $\mu_{t-1} = \mu_t$. Collectively, we can then express the expectation of police trusworthiness in time t, μ_t, as follows:

$$E[\mu_t] = \gamma\Bigg(\overbrace{(\sigma_0 + \pi(\sigma_1 - \sigma_0))}^{\Pr(S_t=1)} \overbrace{\frac{\mu_{t-1}\sigma_1}{\mu_{t-1}\sigma_1 + (1-\mu_{t-1})\sigma_0}}^{\mu_t|S_t=1} + \underbrace{(1 - \sigma_0 - \pi(\sigma_1 - \sigma_0))}_{\Pr(S_t=0)} \underbrace{\frac{\mu_{t-1}(1-\sigma_1)}{\mu_{t-1}(1-\sigma_1) + (1-\mu_{t-1})(1-\sigma_0)}}_{\mu_t|S_t=0}\Bigg) + (1-\gamma)\mu_{t-1} \tag{3}$$

To the extent that our goal is to understand systematic *variation* in trust in police, this framework broadens the set of explanations for this variation, as discussed in Table 1. To understand how trust varies in social class, suppose that one or more of these parameters varies in social class. Placing the conventional wisdom—poor service/outcomes begets distrust—in the context of our framework shows where this account is underspecified. First, we do not know whether variation in service quality comes from variation in the likelihood of police treatment and security outcomes (σ_1 or σ_0), contact with different pools of police agents (π), or both. Second, we identify other sources of variation in trust that could covary with socioeconomic status and drive the weak, negative correlation between socioeconomic status and trust in police that we have documented. Importantly, these features fall within our conception of institutional trust rather than the alternative accounts we have provided evidence against.

The updating process we describe in (3) has some notable features. First, it is useful to note that $\lim_{t\to\infty} E[\mu_t] = \pi$. Yet, individual beliefs (μ_t) will not generally converge to π. Second, it is also useful to examine $\text{Var}[\mu_t]$, but whether variance increases or decreases in t is ambiguous. Within the LAPOP repeated cross-sectional data or the (relatively) short panels from Medellín and Chile, our ability to track individual updating over long histories (sufficiently large $|t|$) is limited. Given that we only have surveys of adults, these survey data do not offer a persuasive measure of μ_0. Nevertheless, as long as citizens (sometimes) see signals of police quality, μ_0 (so long as it is not 0 or 1) should play little role in characterizing trust in a cross-section of citizens.

Parameter(s)	Source of variation	Description
μ_0	Differences in (initial) beliefs	*Ex-ante* beliefs about the police emerging from early socialization.
$\|t\|$	Different histories of observation	Number of periods of updating on police quality, likely proxied by age.
σ_0, σ_1	Different likelihoods of poor police treatment of citizens or poor security outcomes.	Police may be biased toward some citizens over others. Bias may be a result of police tastes or institutional incentives for abuse/poor service in some communities but not others.
π	Different pools of police agents.	Better (more trustworthy) officers may select into or be assigned to some jurisdictions than others. Alternatively, changes in police compensation or selection may lead to different pools of police over time.
γ	Different levels of attentiveness to police behavior	Some citizens may be more observant than others due to cognitive load

Table 1: Sources of variation in trust in police.

To examine variation in $|t|$, we compare trust in police as a function of respondent age. Trust in police is increasing in age (Figure A11). This observation posits an alternate explanation for the main finding, that trust is decreasing (weakly) in social class, given compositional differences between rich and poor respondents. In our sample, the young are both better educated, slightly more affluent (income is maximized at middle age), and less trusting of the police. Figure A11 suggests that while these compositional differences may strengthen our finding of a negative correlation between socioeconomic status and income, they alone cannot account for the negative correlation we observe. This is evident from the very similar negative gradients of income on trust within each age bracket.

Given the observation of a signal—a poor security outcome or mistreatment by police—different citizens could emerge with different beliefs because (a) they had different priors; or (b) they update via different likelihoods (σ_1 and σ_0). To the extent that a posterior today (in t) serves as a prior tomorrow (in $t+1$), our main finding is that average beliefs in police trustworthiness decrease slightly as income increases. Yet, when we compare the beliefs of citizens exposed to the adverse signal ($S_t = 1$) vs. those that were not exposed ($S_t = 0$) in Figure 7, there exists a high degree of stability in the extent of updating (the difference between the orange and blue lines) across all deciles of socioeconomic status. Nevertheless, because the rich have (slightly) lower priors about police trustworthiness, we should expect them to update slightly less on the basis of a "bad" signal

when all citizens share the same likelihoods. Table A8 reveals that, for both measures of social class, this is the case for crime victimization and bribe solicitation, though these differences are small in magnitude as we might expect for such small differences in prior beliefs.[5] The data we observe are broadly consistent with Bayesian updating in the model and do not point to different likelihoods or differential ability to update as a source of the pattern of beliefs that we observe. This is an important observation in light of debates about education, cognitive sophistication, and the ability to update.

The final two sources of variation in trust in police deal with the probability that citizens encounter or observe a given signal of police trustworthiness. Table 1 suggests two possible sources of this variation: the share of trustworthy vs. non-trustworthy police agents that citizens encounter (π) or the probability that citizens detect signals of police behavior (γ). In our conceptualization, the former is institutional, and the latter is behavioral. On the one hand, the sorting of police officers into police precincts, either by police commanders or officers themselves, could lead some populations to greater exposure to untrustworthy or abusive officers. If this varied in citizen social class, these differences in rates at which bad signals ($S_t = 1$) emerge could lead to variation in the frequency of downward updating. Importantly, this can occur even if all police agents (of a given type) treat citizens equally, e.g., without bias.[6]

It may alternatively (or also) be the case that poor security outcomes emerge at similar rates across rich and poor communities—e.g., pools of unbiased officers are similar—but citizens observe these signals and thus update at different rates. Various mechanisms could explain this differential attentiveness, including different sets of concerns, different patterns of time use, or different mobility patterns (Pereira et al., 2019; Bick, Fuchs-Schündeln, and Lagakos, 2018).

Our evidence in Figure 6 suggests that contrary to conventional wisdom, higher socioeconomic status citizens across Latin America report encountering more police abuse (e.g., corruption), more crime victimization, and similar (if variable) perceptions that their neighborhoods are unsafe. We

[5]For the "feeling unsafe" signal, point estimates on the interaction between class and the signal are negative and significant, but small in magnitude.

[6]In our model, equal treatment is equivalent to non-class specific likelihoods (σ_1 and σ_0).

do not have definitive evidence to distinguish between the two former possibilities. To observe the negative correlation between trust in the police as well as the self-reported frequency of signals in the LAPOP data, the poor would need to be exposed to a more trustworthy pool of officers than the rich (higher π for the poor than the rich). This is at odds with most accounts of police hiring and assignment in the region (González, 2017; Brinks, 2007).[7]

Instead, we suggest that the way observations of security outcomes are translated into signals of police trustworthiness may be a behavioral source of the patterns we document. The strongest piece of evidence supporting this observation comes from the analysis of variation in trust of police by social class. Specifically, Figure 9 shows that the standard deviation of trust in police is *decreasing* in decile of education and income across the region. This is consistent with the idea that γ is *increasing* in socioeconomic status. In other words, higher socioeconomic status citizens observe signals of police behavior/outputs at higher rates, holding fixed the rate at which these abuses occur. It could also be consistent with the rich having a lower threshold for interpreting police behavior as a bad signal of police trustworthiness (and thus seeing more bad signals, even fixing police behavior). These possibilities are observationally equivalent in our data and merit further investigation.

Our framework provides analytic leverage to hypothesize the source of differences between trust in police in the US and in Latin America. First, there exists a much wider literature documenting discriminatory policing in the US context. The origins of this discrimination can include bias in the behavior of police toward certain citizens (the likelihoods σ_1 and σ_0) or differences in the pools of officers that serve different populations (π). Indeed, police hiring and funding in the US is far more decentralized than in most (if all) of Latin America, which may lead to a higher potential for sorting. It need not be the case that similar processes of processing police signals are *not* present in the US. Because both the bias and police quality explanations cut against this process in the US context, it may be the case that we simply cannot disentangle this process (mechanism)

[7]It would, of course, be possible to design a public servant allocation scheme that assigns higher-rated or more experienced public servants to *harder* or *less desirable* postings or jurisdictions. We are not aware of such a scheme in Latin America with respect to police assignment.

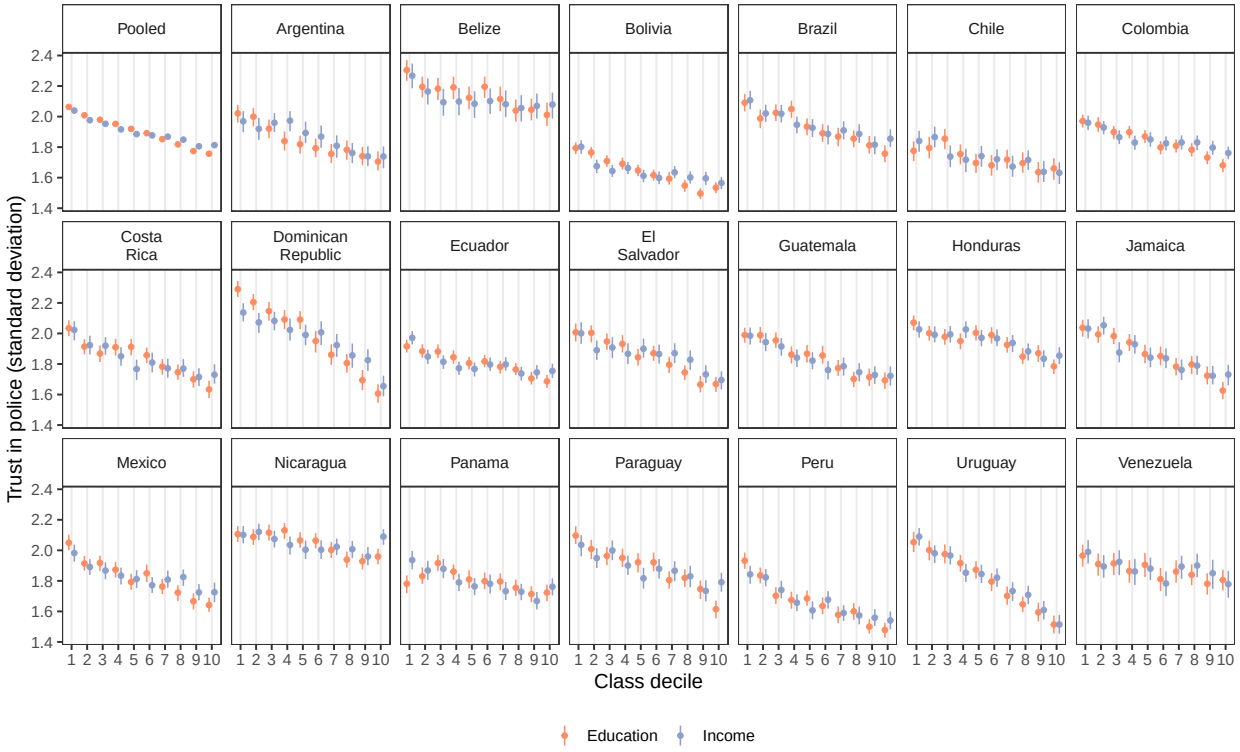

Figure 9: Figure shows the estimated standard deviation of trust in police by class decile for the country-specific and pooled samples its the 95% bootstrapped confidence intervals.

from aggregate survey data on trust in police.

7 Conclusion: Implications for the study of institutional trust

Conventional wisdom from the US suggests that racially and socioeconomically marginalized populations have lower trust in police than their advantaged counterparts. This apparent conventional wisdom has been extended to other contexts by police agencies that design interventions to increase trust (Blair et al., 2021) and, as our forecasting exercise shows, expert beliefs about trust in police. We show that descriptively, a distinct pattern obtains in Latin America, as a whole, and effectively all countries therein.

We argue that trust in the police should be viewed as a belief about the trustworthiness of the officer pool as a whole and that there is ample evidence that citizens update in a roughly Bayesian manner. We find suggestive evidence that low- and high-class citizens translate observations of police outputs or behavior into signals of trustworthiness in different ways. To our knowledge, this is a novel mechanism for explaining differences in trust in police. More research is needed to integrate this process for how citizens observe or interpret government outputs into an informational signal.

This mechanism has important implications for the design of policing interventions intended to increase trust or improve security outcomes. A central mechanism underpinning community policing, for example, is that greater exposure to local police officers increases citizen observation of these officers in a positive setting. But if trust is enhanced by *lack* of observation of, or translation of police outputs into, informational signals, community policing could well backfire. Moreover, if these meetings raise the expectations for police, it may have a similar effect of changing how observations are converted into informational signals, reducing trust in police. This mechanism also suggests that improvements in police quality or reductions in police misbehavior/bias should have effects on trust that are *heterogeneous* in magnitude. In the Latin American context, these interventions may yield the highest trust increases among high-socioeconomic status populations, introducing a heretofore undocumented set of distributive considerations.

References

Altamirano, Melina, Sarah Berens, and Sandra Ley. 2022. "Security or Social Spending? Perceptions of Insecurity, Victimization, and Policy Priorities in Mexico and Brazil." *Political Studies* 1 (24): 00323217221096559.

Bhattacharya, Rajeev, Timothy M. Devinney, and Madan M. Pollutla. 1998. "A Formal Model of Trust Based on Outcomes." *The Academy of Managment Review* 23 (3): 459–472.

Bick, Alexander, Nicola Fuchs-Schündeln, and David Lagakos. 2018. "How Do Hours Worked Vary with Income? Cross-Country Evidence and Implications." *American Economic Review* 108 (1): 170–199.

Blair, Graeme, Jeremy M. Weinstein, Fotini Christia, Eric Arias, Emile Badran, Robert A. Blair, Ali Cheema, Ahsan Farooqui, Thiemo Fetzer, Guy Grossman, Dotan Haim, Zulfiqar Hameed, Rebecca Hanson, Ali Hasanain, Dorothy Kronick, Benjamin S. Morse, Robert Muggah, Fatiq Nadeem, Lily L. Tsai, Matthew Nanes, Tara Slough, Nico Ravanilla, Jacob N. Shapiro, Barbara Silva, Pedro C. L. Souza, and Anna M. Wilke. 2021. "Community Policing Does Not Build Citizen Trust in Police or Reduce Crime in the Global South." *Science (New York, N.Y.)* 374 (6571): eabd3446.

Brinks, Daniel M. 2007. *The Judicial Response to Police Killings in Latin America: Inequality and the Rule of Law*. Cambridge: Cambridge University Press.

Brinks, Daniel M. 2019. "Access to What? Legal Agency and Access to Justice for Indigenous Peoples in Latin America." *The Journal of Development Studies* 55 (3): 348–365.

Britto Ruiz, Diana, and Jorge Ordóñez Valverde. 2005. "Las prácticas de distinción social." *Polis. Revista Latinoamericana* (11).

Cao, Liqun, and Jihong Solomon Zhao. 2005. "Confidence in the Police in Latin America." *Journal of Criminal Justice* 33 (5): 403–412.

Citrin, Jack, and Laura Stoker. 2018. "Political Trust in a Cynical Age." *Annual Review of Political Science* 21: 49–70.

COES, Centre for Social Conflict and Cohesion Studies. 2022. "Estudio Longitudinal Social de Chile 2016-2021.".

Cruz, José Miguel. 2015. "Police Misconduct and Political Legitimacy in Central America." *Journal of Latin American Studies* 47 (2): 251–283.

Decker, Scott H. 1981. "Citizen Attitudes Toward the Police: A Review of Past Findings and Suggestions for Future Policy." *Journal of Police Science and Administration* 9: 80–87.

DellaVigna, Stefano, Devin Pope, and Eva Vivalt. 2019. "Predict Science to Improve Science." *Science (New York, N.Y.)* 366 (6464): 428–429.

DellaVigna, Stefano, Nicholas Otis, and Eva Vivalt. 2020. "Forecasting the Results of Experiments: Piloting an Elicitation Strategy." *AEA Papers and Proceedings* 110: 75–79.

Gallup, Inc. 2021. "In U.S., Black Confidence in Police Recovers From 2020 Low." https://news.gallup.com/poll/352304/black-confidence-police-recovers-2020-low.aspx.

Gasparini, Leonardo, and Nora Lustig. 2011. "The Rise and Fall of Income Inequality in Latin America." In *The Oxford Handbook of Latin American Economics*, ed. José Antonio Ocampo, and Jaime Ros. Oxford University Press p. 0.

Gaviria, Alejandro, and Carmen Pagés. 2002. "Patterns of Crime Victimization in Latin American Cities." *Journal of Development Economics* 67 (1): 181–203.

Gerbasi, Alexandra, and Karen Cook. 2009. "Trust.". Oxford University Press pp. 218–241.

Gerring, John. 2012. "Mere Description." *British Journal of Political Science* 42 (4): 721–746.

González, Yanilda, and Lindsay Mayka. 2022. "Policing, Democratic Participation, and the Reproduction of Asymmetric Citizenship." *Forthcoming, American Political Science Review* .

González, Yanilda María. 2017. ""What Citizens Can See of the State": Police and the Construction of Democratic Citizenship in Latin America." *Theoretical Criminology* 21 (4): 494–511.

González, Yanilda María. 2020. *Authoritarian Police in Democracy: Contested Security in Latin America*. Cambridge University Press.

Hanson, Rebecca, Dorothy Kronick, and Tara Slough. 2022. "Preaching to the Choir: A Problem of Participatory Interventions.".

Hardin, Russel. 2003. "Trust in Government.". Russell Sage Foundation chapter 1, pp. 9–27.

Hoffman, Kelly, and Miguel Angel Centeno. 2003. "The Lopsided Continent: Inequality in Latin America." *Annual Review of Sociology* 29 (1): 363–390.

Jaitman, Laura, Rodrigo Soares, Mauricio Olavarría-Gambi, and Roberto Guerrero Compeán. 2015. "The Welfare Costs of Crime and Violence in Latin America and the Caribbean.".

Johnson, Stephen, Johanna Mendelson Forman, and Katherine Bliss. 2012. Police Reform in Latin America: Implications for U.S. Polic. Technical report Center for Strategic and International Studies.

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.

Kunda, Ziva. 1987. "Motivated Inference: Self-serving Generation and Evaluation of Causal Theories." *Journal of Personality and Social Psychology* 53 (4): 636–647.

LAPOP, Lab. 2022. "The AmericasBarometer." www.vanderbilt.edu/lapop.

Levi, Margaret, and Laura Stoker. 2000. "Political Trust and Trustworthiness." *Annual Review of Political Science* 3: 457–507.

Little, Andrew T., Keith E. Schnakenberg, and Ian R. Turner. 2022. "Motivated Reasoning and Democratic Accountability." *American Political Science Review* 116 (2): 751–767.

Liu, Licheng, Ye Wang, and Yiqing Xu. 2022. "A Practical Guide to Counterfactual Estimators for Causal Inference with Time-Series Cross-Sectional Data." *American Journal of Political Science* n/a (n/a).

Macaulay, Fiona. 2012. "Cycles of Police Reform in Latin America." In *Policing in Africa*, ed. David J. Francis. New York: Palgrave Macmillan US pp. 165–190.

Macdonald, John, and Robert J. Stokes. 2006. "Race, Social Capital, and Trust in the Police." *Urban Affairs Review* 41 (3): 358–375.

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.

Malone, Mary Fran T., and Lucia Dammert. 2021. "The Police and the Public: Policing Practices and Public Trust in Latin America." *Policing and Society* 31 (4): 418–433.

Mondak, Jeffery J., Damarys Canache, Mitchell A. Seligson, and Matthew V. Hibbing. 2011. "The Participatory Personality: Evidence from Latin America." *British Journal of Political Science* 41 (1): 211–221.

Muggah, Robert, and Katherine Tobon. 2018. "Citizen Security in Latin America: Facts and Figures." *Security Research Hub Reports* .

Munger, Kevin, Andrew M. Guess, and Eszter Hargittai. 2021. "Quantitative Description of Digital Media: A Modest Proposal to Disrupt Academic Publishing." *Journal of Quantitative Description: Digital Media* 1: 1–13.

O'Donnell, Guillermo. 1999. "The (Un)Rule of Law and the Underprivileged in Latin America.". Notre Dame Press chapter Polyarchies and the (un) rule of law in Latin America: A Partial Conclusion.

Ojeda, Christopher. 2016. "The Effect of 9/11 on the Heritability of Political Trust." *Political Psychology* 37 (1): 73–88.

Pereira, Rafael H. M., David Banister, Tim Schwanen, and Nate Wessel. 2019. "Distributional Effects of Transport Policies on Inequalities in Access to Opportunities in Rio de Janeiro." *Journal of Transport and Land Use* 12 (1): 741–764.

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.

Phelps, Edmund. 1972. "The Statistical Theory of Racism and Sexism." *American Economic Review* 62: 659–61.

Putnam, Robert D., Robert Leonardi, and Raffaella Y. Nanetti. 1993. *Making Democracy Work : Civic Traditions in Modern Italy*. Princeton, UNITED STATES: Princeton University Press.

PytlikZillig, Lisa M., and Christopher D. Kimbrough. 2016. "Interdisciplinary Perspectives on Trust: Towards Theoretical and Methodological Integration.". New York: Springer chapter Consensus on Conceptualizations and Definitions of Trust: Are We There Yet?, pp. 49–70.

Sabatini, F. 2006. The Social Spatial Segregation in the Cities of Latin America.

Skogan, Wesley G. 2005. "Citizen Satisfaction with Police Encounters." *Police Quarterly* 8 (3): 298–321.

Slough, Tara, and Scott A. Tyson. 2023. "External Validity and Meta-Analysis." *American Journal of Political Science* 67 (2): 440–455.

Taber, Charles S., and Milton Lodge. 2006. "Motivated Skepticism in the Evaluation of Political Beliefs." *American Journal of Political Science* 50 (3): 755–769.

Vilalta, Carlos. 2020. "Violence in Latin America: An Overview of Research and Issues." *Annual Review of Sociology* 46 (1): 693–706.

Villarreal, Andrés. 2010. "Stratification by Skin Color in Contemporary Mexico." *American Sociological Review* 75 (5): 652–678.

Visconti, Giancarlo. 2020. "Policy Preferences after Crime Victimization: Panel and Survey Evidence from Latin America." *British Journal of Political Science* 50 (4): 1481–1495.