

PRIORS RULE: WHEN DO MALFEASANCE REVELATIONS HELP OR HURT INCUMBENT PARTIES?

Eric Arias

Independent researcher

John Marshall

Columbia University

Horacio Larreguy

ITAM

Pablo Querubín

New York University

Abstract

Effective policy-making requires that voters avoid electing malfeasant politicians. However, informing voters of incumbent malfeasance in corrupt contexts may not reduce incumbent support. As our simple learning model shows, electoral sanctioning is limited where voters already believed incumbents to be malfeasant, while information's effect on turnout is non-monotonic in the magnitude of reported malfeasance. We conducted a field experiment in Mexico, which informed voters about malfeasant mayoral spending before municipal elections, to test whether these Bayesian predictions apply in a developing context where many voters are poorly informed. Consistent with voter learning, the intervention increased incumbent vote share where voters possessed unfavorable prior beliefs and when audit reports caused voters to favorably update their posterior beliefs about the incumbent's malfeasance. Furthermore, we find that low and, especially, high malfeasance revelations increased turnout, while less surprising information reduced turnout. These results suggest that improved governance requires greater transparency and citizen expectations. (JEL: D72, D83)

The editor in charge of this paper was Imran Rasul.

Acknowledgments: We thank the steering committee and other team members of the EGAP Metaketa initiative for illuminating discussions and useful comments. We also thank Abhijit Banerjee, Esther Dufo, Georgy Egorov, Leopoldo Fergusson, Pablo Fernandez-Vazquez, Claudio Ferraz, Nikhar Gaikwad, Reema Hanna, Torben Iversen, Ethan Kaplan, Philip Keefer, Stuti Khemani, Julien Labonne, Marco Larizza, Maria Petrova, Vincent Pons, Massimo Pulejo, Tara Slough, Johannes Urpelainen, and participants at several workshops for their feedback and comments. We are extremely grateful to Anais Anderson, Adriana Paz, and Alejandra Rogel, and the Data OPM and Que Funciona para el Desarrollo teams for their implementation of this project. We are grateful to Tommaso Nannicini and Francesco Trebbi, and Frederico Finan and Laura Schechter, for sharing their survey instruments. We thank Taylor Boas and Danny Hidalgo for sharing their experimental data from Brazil. This research was financed by the EGAP Metaketa initiative. Larreguy gratefully acknowledges funding from France's Agence Nationale de la Recherche under the Programme des Investissements d'Avenir grant ANR-17-EURE-0010. The research design was approved by the Harvard CUHS (15-1068) and the NYU UCAIHS (15-10587). Our pre-analysis plan was pre-registered with EGAP, and is publicly available at egap.org/registration/760.

E-mail: ericarias302@gmail.com (Arias); horacio.larreguy@itam.mx (Larreguy); jm4401@columbia.edu (Marshall); pablo.querubin@nyu.edu (Querubín)

1. Introduction

Elected politicians around the world are expected to implement policies to support economic development and alleviate poverty. The median voter in developing countries is generally poor, and thus often stands to benefit substantially from anti-poverty programs. However, such programs can be beset by bribery (e.g. Hsieh and Moretti 2006), procurement and invoicing fraud (e.g. Ferraz and Finan 2008), and misallocated spending (e.g. Larreguy et al. 2020). While policy-makers and NGOs have increasingly sought to design institutions to mitigate these lasting concerns (Khemani et al. 2016), effective political accountability ultimately requires citizens to elect honest politicians. A key question is thus: when will voters hold their governments to account by punishing incumbent parties for malfeasant behavior in office?

A growing political economy literature has emphasized the importance of providing voters with information about incumbent performance in office. Exposure to negative information, such as reports revealing corruption, is expected to induce the electorate to screen out (e.g. Fearon 1999; Rogoff 1990) or sanction (e.g. Barro 1973; Ferejohn 1986) those responsible when it is believed that politicians or parties vary in competence or their efforts to represent voters' interests.

However, while several prominent studies have found that incumbent performance information promotes electoral accountability, the evidence supporting the voter learning logic is mixed. On one hand, Chang et al. (2010), Ferraz and Finan (2008), and Larreguy et al. (2020) find that media revelations of mayoral malfeasance reduce incumbent support in Italy, Brazil, and Mexico, respectively. Experimental studies by Banerjee et al. (2011) and Buntaine et al. (2018) further find that disseminating scorecards reporting incumbent activity can reduce the vote share of poorly performing elected officials and increase the vote share of highly performing elected officials in India and Uganda. On the other hand, other recent field experiments by Adida et al. (2020), Boas et al. (2019), Chong et al. (2015), and de Figueiredo et al. (2014) find that disseminating information about national and local incumbent performance in Benin, Brazil, and Mexico did little to affect incumbent electoral prospects. The effects on turnout of revealing incumbent malfeasance are similarly mixed: while Chong et al. (2015) suggest that unfavorable information may induce systemic disengagement in Mexico, Banerjee et al. (2011) observe increased turnout in India.

Even among the findings that information induces sanctions (rewards) for low (high)-performing incumbents, it is not obvious that information's effects actually reflect the learning mechanism underpinning theories of electoral accountability.¹ Since the studies reporting the largest effects of information campaigns typically involve mass media, it remains possible that information provision instead generates a public signal coordinating voters in favor of better candidates and against worse candidates without significantly updating their beliefs (e.g. Morris and Shin 2002).

1. In fact, the studies that administer post-election surveys suggest that voting behavior changed *without* substantially altering voter beliefs about incumbent performance (Banerjee et al. 2011; Buntaine et al. 2018).

Another recent strand of literature shows that informational interventions may impact electoral outcomes by triggering responses from incumbent and challenger parties or inducing a strategic reallocation of campaign resources (Banerjee et al. 2011; Bidwell et al. 2020; Bowles and Larreguy 2020; Cruz et al. 2021). Beyond its theoretical importance, whether belief updating or these alternative mechanisms drive the effects of providing information has important implications for the design and scale of information dissemination campaigns.

We argue that voters' *prior* beliefs can play a key role in rationalizing these mixed findings, and ultimately help to explain when and how providing information about incumbent performance in office impacts turnout and vote choice. We illustrate the importance of the direction and magnitude of belief updating in response to signals of incumbent malfeasance in a two-party model where expressive voters learn about the incumbent party's underlying malfeasance. Our simple model emphasizes that, if voters already believe that their incumbent party is malfeasant, even revelations of relatively severe malfeasance can fail to decrease incumbent support because voters do not update their posterior beliefs unfavorably. Accordingly, well-intentioned interventions can produce seemingly perverse consequences in terms of supporting malfeasant politicians.

Furthermore, the implications for turnout imply a testable non-monotonicity. Under relatively general and empirically plausible distributions of partisan attachments, we show that signals which induce moderate levels of updating can reduce turnout by shifting a large mass of weak supporters of one party to abstain when turning out is costly. However, sufficiently surprising revelations—whether favorable or unfavorable—increase turnout by shifting voters who previously abstained, and even supporters of the other party, to turn out for the party shown to be less malfeasant.

We test these theoretical predictions, which we pre-registered, using a field experiment conducted in Mexico around the 2015 municipal elections. Beyond its large population and recent shift towards a more pluralistic democracy, Mexico's relatively high—but substantially varying—levels of corruption and distrust in elected politicians across municipalities make it a well-suited location to test our argument. Although municipal mayors could not seek re-election at the time, voters hold parties responsible for incumbent performance in office in Mexico's party-centric system. Extending two recent studies examining electoral responses to municipal audit reports, but with markedly different findings (Chong et al. 2015; Larreguy et al. 2020), we examine how voters respond to leaflets revealing the extent to which municipal governments correctly spent federal transfers earmarked for social infrastructure projects benefiting the poor.

We partnered with a local civil society organization to disseminate leaflets documenting the results of independent municipal audit reports across 678 rural and urban electoral precincts in 26 municipalities from four central Mexican states in the weeks just before the election. Voters in treated precincts, where leaflets were delivered to up to 200 households, received one of two measures of incumbent malfeasance: the share of funds earmarked for social infrastructure projects that was spent on projects that did not benefit the poor, or the share of such funds spent on unauthorized projects.

These measures ranged from 0% to 58% in our sample, with significant variation around the mean of 21%. Since a baseline survey was not financially feasible, we use the municipal control group's post-election beliefs to proxy for the pre-treatment prior beliefs of treated and control voters within each municipality. We further proxy for belief updating among treated voters within each municipality by measuring belief updating by control group respondents when exposed to the leaflet information during the post-election survey. A variety of tests validate these proxies.

Consistent with the theory, we find that the impact of revealing municipal audit reports on voters' support for the incumbent party depends on how the information relates to their prior beliefs. On average, information did not affect voters' posterior beliefs regarding incumbent party malfeasance. Most likely by increasing the certainty of risk averse voters, treatment ultimately increased the incumbent party's vote share by two percentage points. However, our key finding is that voter learning is a central force driving voting behavior. At both the individual and precinct levels, we show that the average effects mask substantial heterogeneity in the response of a Mexican electorate generally skeptical that local politicians allocate funds as legally mandated. Specifically, the increase in incumbent support induced by our treatment is concentrated in municipalities in which audit reports revealed low malfeasance, and where voters already believed that the incumbent party was malfeasant, voters possessed less precise prior beliefs, and voters most favorably updated their posterior beliefs regarding incumbent party malfeasance upon receiving the information.

The non-monotonic effect of malfeasance revelations on electoral turnout is also supported, though changes in turnout were relatively small. Information provision produced heterogeneous effects on turnout, with relatively unsurprising information—20-30% of funds spent on projects that did not benefit the poor or on unauthorized projects—depressing turnout by around 1 percentage point and extreme cases of malfeasance—both 0% and above 50%—mobilizing turnout by around 0.5 percentage points. In contrast with the view that malfeasance revelations breed generalized disengagement (Chong et al. 2015), we find little evidence to suggest that revealing more severe cases of malfeasance to voters reduces confidence in the capacity of elections to select competent politicians.

Several further analyses suggest that these changes in beliefs and voting behavior were largely driven by voter learning. First, a number of robustness checks show that heterogeneity in response to treatment is not driven by potential confounds of voters' prior beliefs or the level of malfeasance reported. Second, although incumbent and especially challenger parties discredited or incorporated malfeasance reports into their election campaigns, these reactions are unlikely to be the primary determinant of voters' response to treatment. This is because the information treatment increased incumbent support on average and politician reactions did not respond differentially in municipalities where voters had more favorable prior beliefs or updated more unfavorably about the incumbent after receiving the information. Third, voter coordination was also second-order for understanding the intervention's positive effect on incumbent support overall as well as heterogeneity in its effect by voters' prior beliefs, voters' belief updating, and the level of malfeasance reported. As we show in

a separate paper, electoral precincts containing more highly-connected networks did respond to treatment by coordinating votes for challengers, but this occurred where voters already believed that challenger parties were less malfeasant than incumbent parties rather than in response to new malfeasance information (Arias et al. 2019).

By documenting electoral accountability and sophisticated learning by voters in response to receiving incumbent performance information, this article makes three main contributions. First, we provide the first clear evidence from a developing country of the Bayesian interaction between the provision of non-partisan information and prior beliefs for understanding voting behavior. While previous studies have highlighted the potential importance of voters' prior beliefs about incumbent performance (Banerjee et al. 2011; Buntaine et al. 2018; Chong et al. 2015; Ferraz and Finan 2008; Humphreys and Weinstein 2012; Larreguy et al. 2020), such studies either did not measure prior beliefs and updating or did not detect effects of information provision on posterior beliefs and actual vote choices.² By illustrating the voter learning channel, our findings help rationalize why Brazilian voters only punish incumbents responsible for more than one corruption violation (Ferraz and Finan 2008) and performance scorecards affect support for the best- and worst-performing Indian politicians (Banerjee et al. 2011). As well as differences in dissemination technologies and the relevance of the information provided, voters' low expectations of elected politicians may then explain the mixed impact of disseminating credible indicators of what may objectively be regarded as poor incumbent performance on electoral accountability.

Second, our focus on voters' prior beliefs in the context of non-partisan information provision complements previous studies highlighting the significance of prior beliefs for understanding voter responses to partisan campaign messages. In particular, Kendall et al. (2015) indicate that Italian voters in a single municipality updated from both valence and ideological messages delivered by an incumbent's canvassing team during an election campaign, although only the valence message—the high regional ranking of the mayor's development plan—influenced vote choice.³ Our results imply that such sophisticated learning in a developed context extends to performance indicators delivered by non-partisan sources in a developing context. This may be especially important for policy-makers and civil society organizations seeking to maintain credibility by avoiding partisan messaging.

Third, we provide an alternative interpretation for extant results suggesting that revelations of malfeasance motivate voters to disengage from the political system and reduce turnout (Chong et al. 2015; de Figueiredo et al. 2014). Most notably, Chong et al. (2015) surprisingly find that revealing severe malfeasance reduced challenger turnout more than incumbent turnout. However, since this does not account for how the

2. Other studies in the EGAP Metaketa initiative, which this study was part of, also examined the updating of posterior beliefs (see Dunning et al. 2019). However, the other studies generally yielded relatively inconclusive evidence, possibly because they focused primarily on the *direction* of updating (rather than its extent) and took a different theoretical approach to turnout.

3. Other natural or field experiments also document the effectiveness of partisan campaign messaging (e.g. Larreguy et al. 2018; Pons 2018; Spenkuch and Toniatti 2018), but without examining voter beliefs.

information provided related to voters' prior beliefs, it can be rationalized within our theoretical framework by voters expecting particularly high levels of malfeasance by the incumbent party. Although we do not preclude disengagement, at least in theory, our approach nevertheless substantiates the claim that the mixed extant findings with respect to turnout may to a significant degree reflect Bayesian updating. The importance of belief updating in making turnout decisions also accords with Leon's (2017) finding that experimentally reducing voters' perception of fines for abstention reduced turnout in Peru, especially among the voters most indifferent between parties.

The article is structured as follows. Section 2 describes the Mexican municipal context motivating our argument. Section 3 presents a simple model highlighting the conditions under which information increases or decreases a voter's propensity to turn out and cast a ballot for the incumbent party. Section 4 explains and validates our experimental design. Sections 5 and 6, respectively, present the individual- and precinct-level results. Section 7 concludes.

2. Malfeasance, Audits, and Elections in Mexican Municipalities

Mexico's federal system is divided into 31 states (and the Federal District of Mexico City), which contain around 2,500 municipalities and 67,000 electoral precincts. Following major decentralization reforms in the 1990s (see Wellenstein et al. 2006), municipal governments—the focus of this article—have played an important role in delivering basic public services and managing local infrastructure. Municipalities, which account for almost 10% of total government spending, are governed by mayors who were typically elected to three-year non-renewable terms.⁴

2.1. Independent Audits of Municipal Spending

A key component of a mayor's budget is the Municipal Fund for Social Infrastructure (FISM), which represents 24% of the average municipality's budget. According to the 1997 Fiscal Coordination Law, FISM funds are direct federal transfers mandated exclusively for infrastructure projects that benefit populations living in poverty, as defined by localities deemed to be marginalized by the National Population Council (CONAPO). Eligible projects include investments in the water supply, drainage, electrification, health infrastructure, education infrastructure, housing, and roads. However, citizens are poorly informed about both the resources available to mayors and their responsibility to provide basic public services (Chong et al. 2015).

The use of federal funds, including FISM transfers, is subject to independent audits by the Federal Auditor's Office (ASF). Although the ASF reports to Congress, its autonomy is enshrined in the constitution, and it has the power to impose fines, recommend economic sanctions, and file or recommend criminal lawsuits against

4. Re-election became possible for incumbents in most states as of July 2018.

public officials. The ASF selects around 150 municipalities for audit each year, based primarily on the relative contribution of FISM transfers to the municipal budget, historical audit outcomes, factors that raise the likelihood of mismanagement, and whether the municipality has recently been audited (including concurrent federal audits of other programs—see Auditoría Superior de la Federación 2014). The municipalities to be audited in a given year are announced after the funds disbursed for a given fiscal year have been spent.

Audits address the spending, accounting, and management of FISM funds from the previous fiscal year. We focus on two key dimensions of mayoral malfeasance documented in these audit reports, which are not necessarily mutually exclusive: (i) the share of funds spent on social infrastructure projects that did not directly benefit the poor; and (ii) the share of funds spent on unauthorized projects, which includes the diversion of resources to non-social infrastructure projects (e.g. personal expenses and election campaigns) and funds that are not accounted for. Between 2007 and 2015, 8% of audited funds were spent on projects that did not benefit the poor, while 6% were spent on unauthorized projects. Since ASF reports capture only one dimension of malfeasance, it is not surprising that 42% of voters do not believe that municipal governments use public resources honestly (Chong et al. 2015).

At the time of this study, the results for each audited municipality were reported to Congress in February the year after the audit was conducted. All reports are available on the ASF's website, <http://www.asf.gob.mx/Default/Index.asf.gob.mx>. Despite their public release, most voters are poorly informed about the ASF and media coverage of individual municipalities is mixed.

2.2. Municipal Elections

Traditionally, local political competition has been between either the populist Institutional Revolutionary Party (PRI) and the right-wing National Action Party (PAN), or between the PRI and its left-wing offshoot, the Party of the Democratic Revolution (PRD). Due to regional bases of political support and highly localized influence within municipalities, local politics is typically dominated by one or two main parties.⁵ In the municipal elections that we study, the average effective number of political parties by vote share at the precinct and municipal levels remains consistently around 2.5.⁶ Moreover, as Appendix Figure A.1 shows, this two-party dominance is reflected in the generally bimodal distribution of voter partisanship within municipalities.

5. In order to get elected, the three large parties often subsume smaller parties into municipal-level coalitions. The National Regeneration Movement (MORENA) also stood for the first time in 2015, and made headway against this hegemony at the national level, obtaining 9% of the federal legislative vote. However, it was not until 2018 that MORENA obtained large vote shares across the country.

6. The effective number of parties is given by $1 / \sum_{j \in J} V_j^2$, where V_j is party j 's vote share (Laakso and Taagepera 1979).

Although economic and criminal punishments for misallocating funds are relatively rare, there are good reasons to believe that voters might hold the incumbent party responsible, even before mayors could seek re-election. First, voters are considerably better informed about political parties than about individual politicians (e.g. Chong et al. 2015; Larreguy et al. 2018). Crucially for political accountability, 80% of voters in our survey can correctly identify the party of their municipal incumbent. Second, Mexico's main parties have differentiated candidate selection mechanisms that deliver candidates with similar attributes (Langston 2003). For example, 74% of voters in our survey believe that if the current mayor is malfeasant, then another candidate from the same party is likely to also be malfeasant. Third, citizens care about how their governments allocate resources. The surveys we conducted for this study show that 74% and 72% of respondents in control precincts, respectively, regard fighting poverty and honesty as important or very important when deciding which candidate to vote for.

However, extant evidence of electoral sanctioning in response to revelations of malfeasant behavior in Mexico is mixed. Among voters with access to more broadcast media outlets incentivized to report local news, Larreguy et al. (2020) observe larger electoral penalties (rewards) in municipalities where the ASF reported high (low) levels of malfeasance just before municipal elections. Conversely, in a field experiment conducted in 12 municipalities across three states, Chong et al. (2015) suggest that flyers documenting severe incumbent malfeasance breed generalized disengagement: while incumbent support declined when the incumbent was revealed to be highly malfeasant, challenger support also declined at least as much. The disjuncture between these accountability and disengagement findings, which cover the same information over the same period, exemplifies the need for a more refined theory capable of explaining when and why different types of information impact voters differently.

3. Information, Prior Beliefs, and Voting Behavior

Our theoretical framework explores the implications for electoral accountability of providing information about incumbent malfeasance. A simple learning model first shows that the impact of information on voters' posterior beliefs—and ultimately their vote choice—depends on how the information revealed relates to voters' prior beliefs. Our second insight concerns turnout: with a positive cost of voting and an empirically plausible distribution of voter partisan attachments, information relatively close to voters' prior beliefs may reduce turnout, while major departures can increase turnout by causing wholesale shifts in support between incumbent and challenger parties.

3.1. Model

We consider a simple decision-theoretic model in which voters in a given municipality—or part of a municipality—update their posterior beliefs about a party's malfeasance based on informative signals, and choose between voting for

incumbent party I , voting for challenger party C , and abstaining.⁷ Since two-party competition predominates in most parts of Mexico, this assumption approximates political competition in most Mexican municipalities.

We assume that voters receive expressive utility from voting for the *relatively* less malfeasant party, and only turn out if parties are sufficiently different in terms of the utility that voters expect to obtain from either of them.⁸ For analytical simplicity, we model the expected utility that voter i receives from voting for party $j \in \{I, C\}$ as the sum of a fixed partisan benefit deriving from voting for I and j 's expected level of malfeasance:⁹

$$U_{ij} = \begin{cases} \delta_i + E[-\theta_I] & \text{if } j = I \\ E[-\theta_C] & \text{if } j = C \end{cases} \quad (1)$$

where $\delta_i \in \Gamma \subseteq R$ is a partisan bias towards the incumbent that can be positive or negative, and party j 's type $\theta_j \in R$ captures its underlying level of malfeasance.¹⁰ The partisan bias δ_i is independently and identically distributed across voters according to a twice-differentiable cumulative distribution function F , and could reflect durable partisan attachments or shocks occurring before the election that are uncorrelated with prior beliefs and signals of malfeasance. For simplicity, voters are assumed to be risk-neutral regarding the costs of expected malfeasance θ_j , although similar results hold when voters are risk averse. Finally, $c > 0$ is a constant cost of turning out to vote.

A voter only turns out to vote if the difference in expected utility between the two parties is large enough. Conditional on voting, individuals vote for their most preferred party: i votes for incumbent party I if $\Delta_i := U_{iI} - U_{iC} \geq c$, votes for challenger party C if $-\Delta_i \geq c$, and abstains if $|\Delta_i| < c$.¹¹

Voters are uncertain about the underlying malfeasance θ_j of both the incumbent and challenger parties, and learn from a signal about incumbent party malfeasance in a Bayesian fashion. We assume that all voters in a municipality share a prior belief about the malfeasance of each party j that is normally distributed according to $N(\mu_j, \sigma_j^2)$, where $\lambda_j := 1/\sigma_j^2$ denotes the precision of the prior beliefs. Heterogeneity in prior beliefs across municipalities could emanate from differences in the (realization or number of) private signals pertaining to incumbent and challenger

7. In the model, we abstract from party attempts to counteract the effect of scandal exposure. Empirically, we find some evidence of such responses. However, as explained below, this operates alongside, rather than in place of, voter updating of posterior beliefs.

8. In the relatively large municipalities of our sample, voters are unlikely to perceive themselves as pivotal. In such contexts, expressive voting is a standard motive for voting (e.g. Brennan and Hamlin 1998).

9. The theory could be extended to incorporate a ban on re-election by allowing for imperfect within-party candidate correlations. Provided that candidates within parties are sufficiently similar, the forces underpinning our results remain.

10. The latent malfeasance dimension, and the signal described below, are modeled with unbounded support to simplify the analysis using a normal learning framework.

11. An alternative specification of expressive utility, in which voters vote for j if $U_{ij} > \max\{U_{i,-j}, c\}$, would complicate our analysis but yield qualitatively similar comparative statics.

party malfeasance that voters were previously exposed to. When voters receive an audit report documenting malfeasance that pertains to the incumbent, they observe a common signal s_I drawn from a normal distribution $N(\theta_I, \tau_I^2)$ centered on the incumbent's true (but unknown) malfeasance level θ_I . The known precision of this signal, $\rho_I := 1/\tau_I^2$, could derive from the audit report only capturing one dimension of an incumbent's malfeasance. For simplicity, we consider the case where the malfeasance of each party is known to be independently distributed, but show similar results for correlated prior beliefs in Appendix Section A.6.

After receiving a signal s_I of underlying incumbent malfeasance, voters' posterior beliefs about θ_I become:

$$N\left(\mu_I + \kappa_I(s_I - \mu_I), \frac{1}{\lambda_I + \rho_I}\right) \quad (2)$$

where $\kappa_I := \rho_I/(\lambda_I + \rho_I)$ captures the relative precision of the signal. Higher values of κ_I increase voter updating because the signal is precise relative to voters' prior beliefs, while positive values of $s_I - \mu_I$ —which we henceforth refer to as the extent of unfavorable updating—indicate that the signal exceeds the prior expectation of incumbent malfeasance among voters. The difference in the expected utility of voting for I relative to C for voter i then becomes: $\Delta_i = \delta_i - E[\theta_I | s_I, \mu_I] + E[\theta_C | \mu_C] = \delta_i - [\mu_I + \kappa_I(s_I - \mu_I)] + \mu_C$.

Integrating over the distribution of voter partisan biases, we obtain the following results pertaining to the share of voters V_I that turn out for the incumbent party.

PROPOSITION 1 (Incumbent vote share). *Receiving a signal s_I of incumbent malfeasance increases incumbent party vote share V_I , relative to receiving no signal, if and only if $s_I < \mu_I$. This difference in V_I is decreasing in s_I and increasing in μ_I (provided that κ_I is sufficiently large), and the magnitude of the difference is decreasing in λ_I .*

Proof. All proofs are in Appendix Section A.1.

The effect of different signals on the incumbent party's vote share is illustrated in Figure 1, which plots the distribution of voters by their relative preference Δ_i for the incumbent for bimodal and unimodal cases of F . Voters for whom $\Delta_i > c$ turn out for I , while voters for whom $\Delta_i < -c$ vote for C ; the voters for whom $\Delta_i \in [-c, c]$ abstain. We analyze how the key parameters in our model affect voting behavior by shifting the distribution of voters along the Δ_i axis.

The overall effect of information dissemination depends on how the signal relates to voters' prior beliefs. As illustrated by the three distributions of voter preferences, a signal that the incumbent is less malfeasant than voters initially believed (i.e. $s_I < \mu_I$) raises Δ_i to increase I 's vote share by producing a commensurate shift in the distribution of relative voter preferences to the right. A comparison of the medium and thickest distributions shows that a larger favorable update—due to an especially low s_I —further increases I 's vote share. The magnitude of the distribution

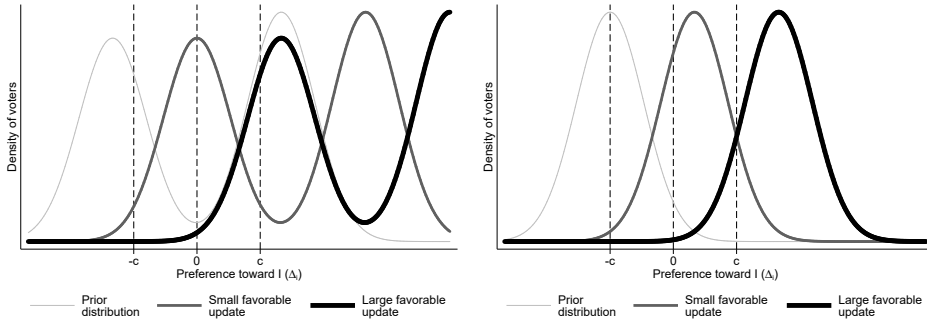


FIGURE 1. Vote choice for different signals and distributions of voters. Left panel: bimodal distribution, where each mode in the prior distribution votes for a different party. Right panel: symmetric unimodal distribution.

shift, in either direction, is decreasing in λ_I because relatively precise prior beliefs reduce the weight attached to the signal in a voter's posterior belief. Finally, the degree to which the distribution shifts after voters receive the public signal is also influenced by μ_I : where voters possess more unfavorable prior beliefs, the effect of the signal is more favorable toward the incumbent because a larger mass of voters can be persuaded to support the incumbent party.

While the incumbent vote share results hold for any distribution F of partisan attachments, the effect of providing information about the incumbent on overall turnout T depends on the shape and position of F and the extent to which information induces updating.

PROPOSITION 2 (Turnout). *Receiving a signal s_I of incumbent malfeasance ambiguously affects turnout: T increases (decreases) when $F(\bar{\delta}_C) - F(\hat{\delta}_C) - [F(\bar{\delta}_I) - F(\hat{\delta}_I)] > (<)0$, where $\bar{\delta}_j$ and $\hat{\delta}_j$ denote the points of indifference between voting for party j and not voting, respectively, with and without the signal. This effect is increasing (decreasing) in s_I when $F'(\bar{\delta}_C) - F'(\bar{\delta}_I) > (<)0$.*

To illustrate the intuition, consider the case where voters receive $s_I < \mu_I$. This signal of lower-than-expected incumbent malfeasance produces two effects: (i) it induces some voters who would not otherwise have voted to turn out for I ; and (ii) it induces some voters who would otherwise have voted for C not to turn out. The relative masses of voters associated with these conflicting effects on turnout determine whether turnout increases or decreases. Intuitively, a sufficiently extreme realization of s_I in either direction will eventually increase turnout because all voters will support or oppose the incumbent party. However, the effect of more moderate signals depends on both F and $s_I - \mu_I$.

To produce sharper empirical predictions, we gain insight by focusing on two empirically-plausible distributions. Motivated by the distribution of partisan

attachments in the two-party races that are common in Mexican municipalities (see Appendix Figure A.1), Figure 1 first considers the case where F is bimodally distributed and, absent a signal, voters at each mode turn out for different parties. Signals that constitute small and large favorable updates demonstrate how the effect on turnout can be non-monotonic: while the medium thickness distribution associated with a small favorable update reduces turnout by shifting the mode where voters initially supported C to a point of abstention, the thickest distribution associated with a large favorable update increases turnout by shifting the same mode to vote for I . In less polarized contexts, the distribution of partisanship could be unimodal. Figure 1 then considers a symmetric unimodal distribution where the distribution is not centered on zero. This case shows that a sufficiently moderate signal again decreases turnout.

More generally, providing incumbent performance information can produce a non-monotonic effect on turnout whenever the initial distribution of voters allows for a sufficiently large mass of voters to instead abstain—inducing an aggregate decrease in turnout—in response to at least some signals. Building from our two example distributions, the following proposition establishes sufficient conditions for a non-monotonic effect of providing information on incumbent malfeasance on turnout that are relatively general.

PROPOSITION 3 (Non-monotonic effects on turnout). *The following conditions guarantee that the effect of receiving a signal s_I of incumbent malfeasance on turnout is positive for $s_I \leq s^*$ and $s_I \geq s^{**} > s^*$, and is negative for some $s_I \in (s^*, s^{**})$:*

- F is unimodal and the distribution of voter prior beliefs does not minimize or maximize T .
- F is bimodal with modes m_C and m_I , where $m_C \leq \hat{\delta}_C < \hat{\delta}_I \leq m_I$ and $F'(\hat{\delta}_C) \neq F'(\hat{\delta}_I)$.

Where F has more than two modes, which is rare in the empirical context of this study, the effect of a common signal may vary across intervals of s_I .

3.2. Empirical Implications

We focus our comparative static predictions on the effect of providing voters with a common signal of incumbent malfeasance, s_I , via a treatment containing information pertaining to mayoral malfeasance. We now enumerate the core hypotheses that motivate our field experiment and that we registered in our pre-analysis plan.

Comparing the expected prior belief (μ_I) and posterior belief ($\mu_I + \kappa_I(s_I - \mu_I)$) shows that the average effect of providing information on voters' posterior beliefs—and vote choice, as Proposition 1 shows—depends on how the signal relates to voters' prior expectations of malfeasance (i.e. $s_I - \mu_I$) in the average municipality. While the average treatment effect is context-dependent and hard to anticipate until prior beliefs are measured in the municipalities under study, there are clear second-order predictions for the effect of information dissemination on voters' posterior beliefs. First, this effect is smaller where voters already believe that the incumbent party is

malfeasant (i.e. high μ_I), since the signal drags a voter's posterior belief away from their prior belief. Second, the posterior beliefs of voters where voters already possess precise prior beliefs about incumbent malfeasance (i.e. low κ_I , or high λ_I) are less responsive to new information in either direction. Third, voters update their posterior beliefs more favorably (unfavorably) about the incumbent party's malfeasance upon learning that the incumbent is relatively clean (malfeasant) (i.e. low (high) s_I).

These effects on voter's posterior beliefs are summarized in the following hypothesis.

H1 (Posterior beliefs). *The effect of providing information about an incumbent's malfeasance on voters' posterior beliefs about whether the incumbent party is malfeasant is:*

- (a) *Decreasing in voters' prior beliefs that the incumbent party is malfeasant.*
- (b) *Decreasing in magnitude with the precision of voters' prior beliefs.*
- (c) *Increasing in the severity of the reported malfeasance.*

In addition to establishing the extent to which voters update their beliefs, these empirical predictions regarding voters' posterior beliefs imply the following effects on the incumbent party's vote share.

H2 (Incumbent party vote share). *The effect of providing information about an incumbent's malfeasance on the incumbent party's vote share is:*

- (a) *Increasing in voters' prior beliefs that the incumbent party is malfeasant.*
- (b) *Decreasing in magnitude with the precision of voters' prior beliefs.*
- (c) *Decreasing in the severity of the reported malfeasance.*
- (d) *Decreasing in the extent to which the information unfavorably updates voters' prior beliefs.*

As shown above, new information is predicted to produce non-monotonic effects on turnout under empirically plausible distributions of voter partisanship. In particular, highly favorable or unfavorable revelations motivate voters who previously abstained to turn out to vote and induces voters to switch parties, while relatively unsurprising—but nevertheless informative—favorable (unfavorable) information induces challenger (incumbent) partisans to become relatively indifferent between the parties and abstain from voting. While this logic does not yield clear predictions for the average effect of new information or its linear interaction with the level of malfeasance reported, it clearly predicts the following.

H3 (Turnout). *Providing information reporting sufficiently high and low levels of incumbent malfeasance increases electoral turnout, while some intermediate levels of reported malfeasance decrease turnout.*

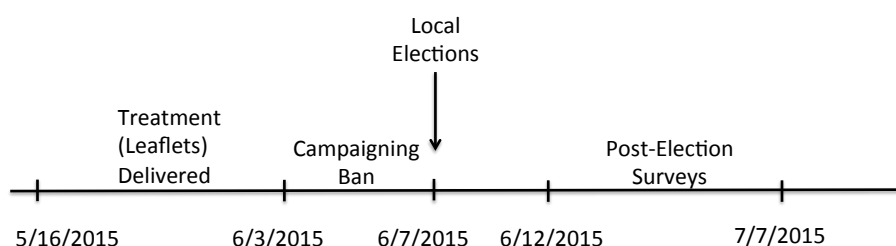


FIGURE 2. Timeline of the experiment's implementation.

4. Experimental Design

We designed a field experiment to test the model's predictions around Mexico's June 7, 2015 municipal elections, which were held concurrently with state and federal legislative elections. We examine the effect of providing voters in 678 electoral precincts across 26 municipalities with the results of audit reports documenting the municipal use of federal transfers designated for infrastructure projects benefiting the poor. We first explain our sample selection, and then outline the intervention, randomization, and estimation strategy. Figure 2 summarizes the project's timeline.

4.1. Sample Selection

Our study focused on 26 municipalities in the central states of Guanajuato, México, San Luis Potosí, and Querétaro. In addition to the fact that they held elections in 2015,¹² these four states were chosen for security and logistical reasons, and because they exhibit variation in the municipal incumbent party. The municipalities were selected to ensure: (i) the safety of voters and our implementing team; (ii) that the level of malfeasance reported by the ASF differed from the average of other incumbent parties elsewhere in the state; and (iii) that the proportion of municipal governments run by each party matched the proportion of municipalities audited by the ASF. The average municipality contained 259,000 registered voters. Appendix Section A.3.3 explains the selection of municipalities in detail.

After immediately receiving threats upon entering Aquismón and Villa Victoria, these municipalities were replaced by Atlacomulco, Temoaya, and an additional block from Tlalnepantla de Baz in the state of México. Since our blocking strategy—explained below—ensured that treatments were randomized within municipalities, excluding these problematic municipalities does not affect the study's internal validity.

12. Municipal elections reflect state electoral cycles, which are staggered across years. On June 7, 2015, 15 states and the federal district held simultaneous local elections.

Within each municipality, we selected up to one third of the electoral precincts. To generate variation in the level of malfeasance reported, we oversampled precincts from municipalities with particularly high or low levels of incumbent malfeasance and starker contrasts with opposition party malfeasance within the state. Within municipalities, we first prioritized accessible rural precincts to mitigate the risk of cross-precinct spillovers and the possibility that voters had already encountered the audit information. Moreover, to maximize the share of households that we could reach with a fixed number of leaflets, attention was restricted to precincts with fewer registered voters. In urban precincts, which constitute 49% of our sample, we restricted attention to precincts with at most 1,750 registered voters, and minimized the number of neighboring urban precincts in our sample. Appendix Table A.2 shows that our final sample of precincts is similar to the national distribution according to various socioeconomic indicators from the 2010 Census.

4.2. Information Treatment

In partnership with the Mexican NGO Borde Político,¹³ we sought to evaluate the impact of distributing leaflets to voters that documented the use of FISM funds in their municipality. For each municipality, the leaflet focused on one of two indicators of incumbent malfeasance—which represent signals of θ_I in our model—documented by the ASF that observational studies find voters care about (Larreguy et al. 2020): the proportion of unauthorized spending or the proportion of spending that did not benefit the poor. For each municipality, we chose the malfeasance measure that maximized the difference from other parties within the municipality's state. All treatments were delivered at the electoral precinct level, Mexico's lowest level of electoral aggregation.

The leaflet was designed to be non-partisan, accessible, and sufficiently intriguing that voters would not discard it.¹⁴ Figure 3 provides an example of a leaflet focusing on a severe case of unauthorized spending in the municipality of Ecatepec de Morelos in the state of México. The front page explains that Borde Político is a non-partisan organization and that the information contained in the leaflet is based on the ASF's official audit reports, which are available online.

The main page first states that FISM funds should only be spent on social infrastructure projects, and provides graphical examples of such projects on the right. The leaflet then informs recipients of the total amount of money their municipality received (146.3 million pesos, in this case), and the percentage of this money that was spent in an unauthorized way (45%). To avoid suspicions of political motivation, neither the incumbent mayor nor their party is referred to directly. Figure 4 shows that

13. Borde Político is a leading NGO seeking to increase voter knowledge about the actions of their politicians in office, with significant experience in developing web-based platforms to provide politically relevant information to voters (see <http://borde.mxborde.mx>).

14. The leaflet was produced by a local graphic designer based on feedback from multiple focus groups. We also sought legal advice to ensure that the leaflets did not constitute political advertisements, and thus were not subject to distribution restrictions stipulated in Mexican electoral law.



FIGURE 3. Example of local information leaflet in Ecatepec de Morelos, México.

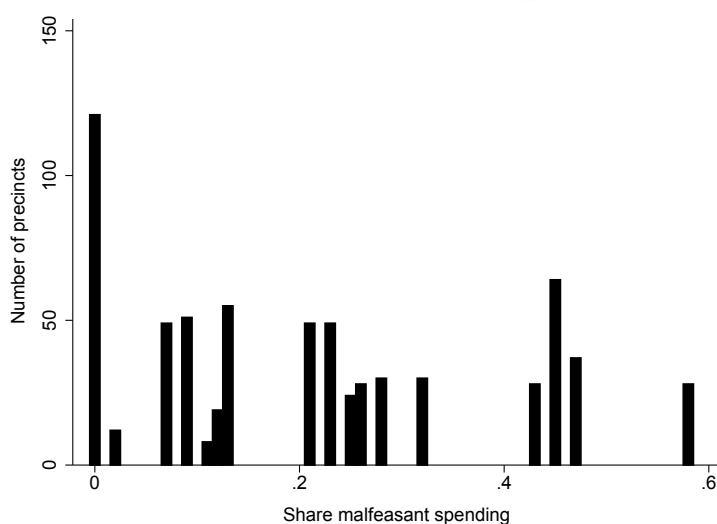


FIGURE 4. Precincts by share of malfeasant spending in our sample.

the average precinct in our sample was informed of 21% malfeasant spending within their municipality.

TABLE 1. Factorial design with a pure control.

	Control	Private	Public
Control	278 precincts		
Local		100 precincts	100 precincts
Comparative		100 precincts	100 precincts

The experiment also incorporated two variants of this information treatment. First, to examine the effect of providing voters with a benchmark against which to compare their municipality's malfeasance, we supplemented the leaflet by providing the mean outcome among all audited municipalities within the same state governed by a different political party; Appendix Figure A.3 provides an example of such a leaflet. Second, to vary the extent to which the distribution of the leaflets was common knowledge among voters within the precinct, we also varied whether leaflet delivery was accompanied by a loudspeaker informing voters that their neighbors would also receive the information and encouraging them to share and discuss it. These treatment variants did not generate different effects, as shown in Tables A.32 and A.33. We present more details about the treatment variants and additional results in a short companion paper (Arias et al. 2018).

4.3. Block Randomization and Implementation

Our sample of 678 precincts was randomly assigned to receive treatments according to the factorial design described in Table 1. The 400 treated precincts were divided equally between the four versions of the information treatment. Given that neither the comparative nor public information components significantly moderated our treatment effects, and all leaflets contained the same baseline information pertaining to incumbent malfeasance, we proceed by pooling all treatment conditions. The control group, comprising 278 electoral precincts, received no leaflets.

For the randomization, precincts were first stratified into blocks each containing six or seven similar precincts *within* a given municipality.¹⁵ Within each block, we then randomly assigned precincts to each of the treatment conditions and, depending on the availability of an additional precinct, either two or three pure control precincts. Because blocks lie strictly within municipalities, malfeasance information always pertains to the same municipal incumbent party and dimension of malfeasance for all precincts within a block.

Our distribution teams delivered one leaflet to a maximum of 200 randomly selected households in the largest locality in rural blocks and randomly selected city blocks in

15. If there were sufficient precincts, and the total number of treated precincts did not exceed one-third of all precincts, we used blocks of seven precincts. Precinct similarity was defined by the Mahalanobis distance between 23 social, economic, demographic, and political variables provided by Mexico's National Statistical Agency and the National Electoral Institute (INE).

urban blocks within each treated precinct.¹⁶ Within our sample, the median precinct contained 353 households (according to the 2010 Census) and 1,056 voters registered for the 2015 election. Where possible, leaflets were delivered in person with a short verbal explanation of the leaflet's provenance. When no adult was available, leaflets were left in mailboxes or taped to the recipient's front door in a waterproof bag. Leaflet delivery took several hours per precinct, and was conducted over a period of three weeks, concluding at the legally designated end of the election campaign four days before the election.

While compliance with the delivery of our treatments was very good in general, we also encountered several issues in the field. In a couple of cases, some leaflets were delivered to voters outside the precinct or adverse weather conditions and poor road conditions prevented us from reaching a precinct.¹⁷ To preserve the randomization, we focus on estimating intent to treat (ITT) effects, which are arguably also the most policy relevant estimand.

4.4. Precinct- and Individual-Level Data

We collected two sources of data to measure our main outcomes. First, using publicly available results and freedom of information requests, we collected official precinct-level electoral returns from each state's electoral institute to compute three pre-registered precinct-level outcomes: incumbent party vote share (as a share of turnout), incumbent party vote share (as a share of registered voters), and turnout. Measuring incumbent party vote share using the share of registered voters allows us to abstract from changes in turnout. We drop the three precincts in our sample that the INE merged with another precinct because they contained fewer than 100 registered voters, which produces a final sample of 675 electoral precincts.¹⁸ We complement the 2015 precinct-level electoral returns with covariates from the 2010 Census and 2012 electoral returns.

Second, we conducted a post-election survey that interviewed ten voters from each of the treated precincts and ten voters from a randomly selected control precinct within each block.¹⁹ At the beginning of the survey, we measured voters' posterior beliefs about each major party's level of corruption or level of interest in supporting the poor (depending on the measure of malfeasance that the leaflets reported on in that

16. Since randomization blocks consist of either only rural or only urban precincts, block fixed effects fully account for any sampling differences across rural and urban precincts.

17. The results are robust to dropping the misassigned precincts from our sample.

18. In two of these cases, the precinct was merged with another precinct that remains in our sample; where the treatment condition conflicts, we retain the larger precinct's treatment status. We were not aware of these merges when the experiment was designed.

19. For treated precincts, enumerators were instructed to survey the localities and city blocks where our informational treatment was delivered. In control precincts, respondents were chosen according to the same protocol used to determine the delivery of leaflets in treated rural and urban blocks.

municipality) on a five-point scale from very low (-2) to very high (2).²⁰ Higher values of this variable indicate that voters believed a party was more malfeasant.²¹ To gauge the precision of these beliefs we then asked respondents to report how certain they were about this belief on a four-point scale ranging from very uncertain (1) to very certain (4). Summary statistics for the main variables are provided in Appendix Table A.3.

4.5. Estimation and Balance

Following our pre-analysis plan, we estimate the average ITT effect of providing any type of information using OLS regressions of the form:

$$Y_{pbm} = \alpha_{bm} + \beta \text{Treatment}_{pbm} + \varepsilon_{pbm}, \quad (3)$$

where Y_{pbm} is an outcome for electoral precinct p within randomization block b in municipality m . For individual-level survey outcomes, Y_{ipbm} also includes an i subscript. Block fixed effects, α_{bm} , are included to adjust for the differential probabilities of treatment assignment across blocks, arising from different block sizes, and to increase efficiency by absorbing block-specific characteristics, such as race-specific differences across municipalities. Including block fixed effects also ensures that we only compare precincts that chose between the same candidates. Throughout, standard errors are clustered at the municipality-treatment level.

We use equation (3) to validate the randomization. Appendix Table A4 demonstrates that the treatment is well-balanced across 46 precinct and survey respondent-level covariates. As usual, there are some significant differences, most notably with respect to incumbent vote share in the previous elections in 2012. However, Appendix Table A.26 shows that our estimates are robust—and, if anything, more precisely estimated—when we adjust for the 40 precinct-level pre-treatment variables.

4.6. Heterogeneous Effects

To test our core hypotheses examining how the effects of providing malfeasance information vary with voters' prior beliefs, the level of malfeasance reported, and the extent of voter updating, we further estimate interactive specifications of the form:

$$Y_{pbm} = \alpha_{bm} + \beta \text{Treatment}_{pbm} + \gamma (\text{Treatment}_{pbm} \times X_m) + \varepsilon_{pbm}, \quad (4)$$

where X_m is a municipality-level variable capturing the heterogeneous effects enumerated in hypotheses H1-H3. Since X_m is not randomly assigned, we also

20. We did not ask about Movimiento Ciudadano, which was the incumbent party only in Apaseo el Alto. Consequently, the 24 precincts from this municipality are dropped from analyses examining prior beliefs.

21. We did not elicit perceptions of the exact share of funds that respondents believe each party spends in a malfeasant way, as we believed this would be hard for respondents to understand.

show the robustness of these specifications to interacting our treatment with potential confounders of X_m .

Measuring the prior beliefs and voter updating required to test parts of H1 and H2 is challenging in our context. Since we could not conduct a baseline survey due to financial constraints, we use the post-election responses from each municipality's surveyed control precincts to proxy for the average pre-treatment beliefs of the treated and control voters within the same municipality. Specifically, to measure the level of voters' prior beliefs—a proxy for the parameter μ_I in the model—we use the mean belief about the incumbent party's malfeasance reported in a municipality's control precincts. For the precision of such prior beliefs (λ_I), we similarly use the mean precision of the incumbent malfeasance perceptions reported in a municipality's control group.

To proxy for the overall extent to which voters in a given municipality updated their posterior beliefs about incumbent party malfeasance in response to treatment, we measure the average change in the control group's beliefs after being exposed to the treatment information during the post-election survey. Specifically, we showed all voters the leaflet corresponding to their municipality at the end of the survey and asked them again how they perceived the incumbent party on the same five-point scale. The average change within each municipality approximates $E[\theta_I | s_I, \mu_I] - E[\theta_I | \mu_I]$ in our model, where positive (negative) values imply that voters updated unfavorably (favorably) relative to their prior beliefs. Given that control group respondents had less time to internalize the information than those in treated precincts, we focus on the *slope* with respect to updating, rather than relying on the levels of updating to categorize favorable and unfavorable updating.

Using post-election surveys from the control group to proxy for pre-treatment beliefs and belief updating by treated voters in a municipality relies on two assumptions: (i) that control group respondents are similar to treatment group respondents; and (ii) that control group respondent beliefs are persistent and not subject to spillovers between the intervention and the post-election survey. Appendix Section A.4.3 provides extensive support for these assumptions. In short, our randomization and the lack of selection into the endline sample support assumption (i), while our blocking strategy ensures that treated and control respondents within municipalities are similar in practice. In support of assumption (ii), we show that municipal-level electoral outcomes do not influence control group beliefs, that there is no evidence of cross-precinct spillovers, that control group respondents update more than treated respondents upon being shown the leaflet, and that a validation exercise we conducted using a panel survey in Brazil suggests limited changes in politician assessments just before and after elections among control voters. Provided these assumptions hold, we further prove in Appendix Section A.4.3 that our estimates represent a lower bound on the magnitude of the precinct- and individual-level heterogeneous effects, to the extent that aggregating to the municipal level adds classical measurement error to precinct- and individual-specific prior beliefs.

TABLE 2. Effect of information treatment on self-reported engagement with leaflet.

	Remember leaflet (1)	Remember reading leaflet (2)	Correctly remember content (3)	Leaflet influenced vote (4)
Information treatment	0.247*** (0.022)	0.171*** (0.018)	0.138*** (0.019)	0.051*** (0.010)
Outcome range	{0,1}	{0,1}	{0,1}	{0,1}
Control outcome mean	0.09	0.05	0.06	0.02
Control outcome std. dev.	0.28	0.22	0.25	0.14
R^2	0.11	0.09	0.10	0.06
Observations	4,958	4,958	4,958	4,958

Notes: All specifications include block fixed effects, and are estimated using OLS. Standard errors clustered by municipality-treatment are in parentheses. * denotes $p < 0.1$, ** denotes $p < 0.05$, *** denotes $p < 0.01$.

5. How do Voters Interpret the Information Treatment?

Before examining the precinct-level electoral results, we start by examining how the information treatment affected voters' actions and posterior beliefs using our post-election survey.

5.1. Manipulation Checks

The four self-reported outcomes in Table 2 show that treated voters received and engaged with the information distributed. Column (1) finds that treated voters were 25 percentage points more likely to remember receiving our leaflet, relative to a control mean of 9% of voters.²² Column (2) next shows that voters in treated precincts were 17 percentage points more likely to report having read the leaflet, while column (3) demonstrates that treated voters were 14 percentage points more likely to correctly answer a multiple choice question asking what issue was covered in the leaflet. Finally, column (4) indicates that 7% of treated voters reported that the leaflet influenced their vote choice, which is 5 percentage points higher than for voters located in control precincts.

Voters generally did not believe that the leaflet was politically motivated. Among treated precincts, 44% of voters correctly believed that the leaflet came from a non-partisan NGO. This response was more than twice as likely as any particular political party, while 33% did not know. The difference was even greater among those who remembered the leaflet. Moreover, neither the comparative nor public treatment variants—which could have been perceived as more political—differentially affected

22. The non-zero control mean likely reflects respondents mistaking our leaflet for another leaflet. Appendix Tables A6 and A7 find no evidence to suggest that this is explained by cross-precinct spillovers.

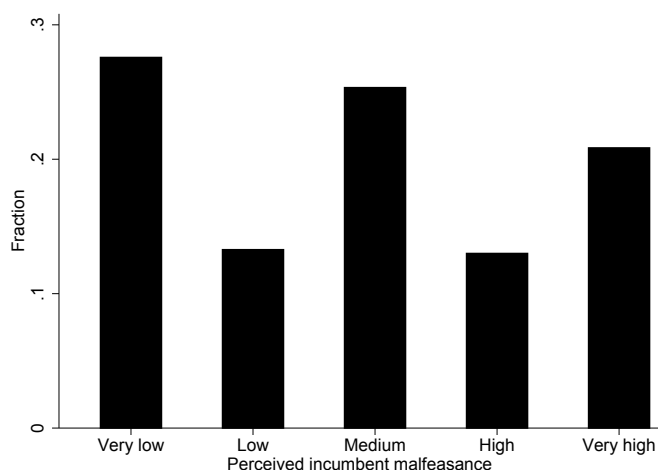


FIGURE 5. Perceived incumbent party malfeasance in control precincts.

the perception that the treatment emanated from a government or political source. Finally, as Appendix Tables A.10 and A.11 show, such perceptions about the leaflet are not correlated with municipal-level prior beliefs, the precision of those beliefs, or belief updating.

5.2. The Effect of Information on Voters' Posterior Beliefs

The distribution of prior beliefs about the municipal incumbent party's malfeasance in the control group indicates that voters had relatively low expectations of incumbent parties. Figure 5 shows that around 60% of respondents reported that they believe the incumbent party engaged in medium to very high levels of corruption or misallocated spending. These expectations of non-trivial incumbent malfeasance in office are consistent with the prior beliefs of Mexican voters reported in Chong et al. (2015). Nevertheless, a significant fraction of voters also perceived that their incumbent party engaged in limited malfeasant behavior. Interestingly, however, voters' prior beliefs are not significantly correlated with the malfeasance levels documented in the ASF reports.

Given voters' low expectations, it is important to assess whether voters favorably or unfavorably update their posterior beliefs about the incumbent party's malfeasance to understand how the information treatment will affect incumbent party support on average. Taking such posterior beliefs as our outcome, Table 3 presents estimates from equations (3) and (4). The negligible and far from statistically significant, coefficient in column (1) shows that treated voters did not increase their posterior beliefs about their incumbent party's malfeasance upon learning of relatively high levels of malfeasance, on average. As in Banerjee et al. (2011), this finding suggests that the information provided broadly aligned with what the *average* voter already believed.

TABLE 3. Effect of information treatment on voters' posterior beliefs about incumbent party malfeasance.

	Perceived incumbent party malfeasance (very low to very high)					
	(1)	(2)	(3)	(4)	(5)	(6)
Information treatment	-0.001 (0.040)	-0.015 (0.037)	0.427 (0.476)	0.016 (0.067)	0.848* (0.452)	-0.096** (0.047)
× Incumbent malfeasance prior		-0.126*** (0.035)			-0.151*** (0.033)	
× Incumbent prior precision			-0.132 (0.149)		-0.258* (0.139)	
× Incumbent malfeasance spending				-0.083 (0.214)	-0.137 (0.165)	
× Unfavorable incumbent updating						0.102*** (0.030)
Outcome range	{-2,-1,0,1,2}	{-2,-1,0,1,2}	{-2,-1,0,1,2}	{-2,-1,0,1,2}	{-2,-1,0,1,2}	{-2,-1,0,1,2}
Control outcome mean	-0.14	-0.14	-0.14	-0.14	-0.14	-0.14
Control outcome std. dev.	1.48	1.48	1.48	1.48	1.48	1.48
Interaction range		[-1.4,1.1]	[2.4,3.8]	[0,0.58]		[-0.6,2.7]
Interaction mean		-0.09	3.23	0.21		0.91
Interaction std. dev.		0.82	0.26	0.17		1.00
R ²	0.29	0.29	0.29	0.29	0.30	0.29
Observations	4,624	4,624	4,624	4,624	4,624	4,624

Notes: All specifications include block fixed effects, and are estimated using OLS. Lower-order interaction terms are absorbed by the block fixed effects. Standard errors clustered by municipality-treatment are in parentheses. * denotes $p < 0.1$, ** denotes $p < 0.05$, *** denotes $p < 0.01$.

However, the lack of updating among treated voters *on average* masks substantial heterogeneity in responses across voters that possessed different prior beliefs. Consistent with hypothesis H1, column (2) indicates that treated voters in municipalities that had unfavorable prior beliefs (i.e. pre-existing expectations of high levels of malfeasance) about the incumbent favorably updated those beliefs about the incumbent, while treated voters in municipalities that had favorable prior beliefs (i.e. expectations of low malfeasance) were more likely to report perceiving their incumbent as corrupt or neglectful of the poor.²³ For the average leaflet, the difference in responses to the treatment between the municipalities with the most favorable and most unfavorable prior beliefs is almost one third of a standard deviation in the posterior belief. Moreover, column (3) shows that treated voters within municipalities with relatively weak prior beliefs are not significantly more likely to unfavorably update their posterior beliefs about their incumbent party. Given the lack of an effect on the average posterior belief, this null finding is also consistent with the model's prediction that the magnitude of the average effect only significantly varies with the precision of voters' prior beliefs when the magnitude of the average effect is not zero. Column (5) reports similar estimates when each of the main interactions are included simultaneously.

23. Appendix Table A.30 shows that this finding is robust to splitting the sample between municipalities with above- and below-median prior beliefs. This check addresses the concern that the results in column (2) could arise mechanically because the municipal-level incumbent malfeasance prior regressor is an aggregation of the individual level posterior belief outcomes in the control group.

The insignificant interaction in column (4) initially provides surprisingly little evidence that the share of misspent funds differentially influences the posterior beliefs of treated voters, as predicted by part (c) of hypothesis H1. However, the precinct-level electoral results described below strongly support this hypothesis. Moreover, posterior beliefs do change once we account for how the information provided relates to prior beliefs. The statistically significant positive coefficient on the interaction between the treatment indicator and our measure of voter updating in column (6) demonstrates that treated voters in municipalities where voters unfavorably (favorably) update their posterior beliefs about the incumbent display substantially more unfavorable (favorable) opinions of the incumbent party. Substantively, a one-standard-deviation difference in updating translates to around a 0.1-standard-deviation change in posterior beliefs among treated voters.²⁴

Our information treatment could, in theory, have also affected posterior beliefs about challengers (e.g. Kendall et al. 2015). Appendix Tables A.16 and A.17 show that treated voters in municipalities with unfavorable prior beliefs about the challenger were also more likely to favorably update their posterior beliefs about the main challenger's malfeasance. Given that such effects are similar across the local and comparative variants of the treatment (Arias et al. 2018), this suggests that voters in our sample primarily updated their posterior beliefs about challengers from the information they received about the incumbent, and that voters believed incumbent and challenger types to be positively correlated.²⁵ To the extent that voters updated similarly about challengers, our estimates are likely to understate the effect of information only inducing voters to update about the incumbent party. However, justifying our focus on incumbent parties, Tables A.18 and A.19 show that voting behavior is driven primarily by how the treatment relates to voters' prior beliefs about the incumbent party rather than challenger parties. Appendix Section A.6 discusses these results in greater detail.

Together, these results confirm that voters meaningfully updated their posterior beliefs about the incumbent party in response to our information treatment. Although reported malfeasance conformed with prior beliefs on average, voters nevertheless updated in a direction that depended on how the information received related to their prior beliefs. We next examine whether such belief updating translates into precinct-level vote choices.

24. This result is not mechanical because municipal-level measures of unfavorable updating are based only on responses from voters in control precincts upon receiving the leaflet.

25. Using a 5-point scale of whether voters believed other candidates of the same party would behave similarly to the incumbent, ranging from not all probable (1) to extremely probable that they will behave similarly in office (5), we find that voters on average believed candidates of all parties to be very similar to the incumbent, though this perceived similarity was slightly higher for incumbent party candidates (3.12) than for those of challenger parties (3.06).

6. Precinct-Level Election Results

We now present our three main precinct-level findings. First, reflecting a combination of voters' low expectations, as well as uncertainty-reduction and potentially differentially effective incumbent responses, we first show that the information treatment increased the incumbent's vote share on average. Second, and consistent with our theoretical model, this effect is greatest where voters updated their posterior beliefs about the incumbent party most favorably based on the information received. Third, we find a non-monotonic effect of information on electoral turnout where revelations of intermediate levels of malfeasance reduced turnout, but extreme levels—low and particularly high malfeasance—increased turnout.

6.1. Average Effects of Information on Incumbent Vote Share

We first document that information about incumbent malfeasance *increased* the incumbent party's vote share, on average within our sample. Column (1) of panel A in Table 4 reports that our intervention significantly increased the incumbent party's vote share, as a proportion of those that turned out, by an average of 2 percentage points. Column (1) of panel B similarly shows that this translates into a 0.8 percentage point increase in the incumbent party's vote share, as a proportion of all registered voters in the precinct. The latter estimate indicates that the information caused the incumbent party to gain more voters, rather than simply demobilized challenger supporters. Relative to the mean vote share in the control group, the information treatment increased the incumbent party's vote share by 5%, or around a sixth of a standard deviation.

Although voters' expectations were sufficiently low that malfeasance revelations did not shift their posterior beliefs on average, there are several reasons why the incumbent party might still benefit from information provision. Specifically, voter risk aversion or political campaign responses to information dissemination could account for the increased incumbent party vote share. We conduct several exploratory analyses, which were not prespecified, to assess these potential mechanisms.

Incorporating voter risk aversion into our stylized model generates another channel through which signals of incumbent malfeasance could influence voters.²⁶ In particular, such information could have increased incumbent party support by reducing posterior uncertainty about the party's type (see also Kendall et al. 2015). We find evidence consistent with this risk-reduction interpretation in Table 5. While the average effect of providing information on posteriors' precision reported in column (1) is zero—likely due to a ceiling effect on high reported levels of precision (mean precision was 3.25 on the four-point scale in control precincts)—column (2) intuitively shows that the greatest increases in posterior precision indeed occurred among respondents in municipalities

26. Risk aversion could be incorporated into the model by allowing the politician type component of a voter's utility function to be represented by $E[-\exp(\theta_j)]$. Our pre-analysis plan noted this theoretical extension may prove relevant.

TABLE 4. Effect of information treatment on incumbent party vote share.

	(1)	(2)	Incumbent party vote share		(5)	(6)
			(3)	(4)		
Panel A: Incumbent party vote share (share of turnout)						
Information treatment	0.020*** (0.004)	0.019*** (0.004)	0.146*** (0.044)	0.031*** (0.006)	0.137*** (0.036)	0.026*** (0.004)
× Incumbent malfeasance prior		0.009* (0.005)			0.005 (0.003)	
× Incumbent prior precision			-0.040*** (0.014)		-0.033*** (0.010)	
× Incumbent malfeasant spending				-0.052** (0.023)	-0.051*** (0.016)	
× Unfavorable incumbent updating						-0.009** (0.004)
Outcome range	[0.07,0.85]	[0.07,0.85]	[0.07,0.85]	[0.07,0.85]	[0.07,0.85]	[0.07,0.85]
Control outcome mean	0.38	0.39	0.39	0.38	0.39	0.39
Control outcome std. dev.	0.12	0.12	0.12	0.12	0.12	0.12
R ²	0.61	0.59	0.59	0.61	0.59	0.59
Panel B: Incumbent party vote share (share of registered voters)						
Information treatment	0.008*** (0.002)	0.008*** (0.002)	0.054** (0.025)	0.014*** (0.003)	0.047** (0.019)	0.012*** (0.003)
× Incumbent malfeasance prior		0.005** (0.002)			0.004* (0.002)	
× Incumbent prior precision			-0.014* (0.008)		-0.010* (0.006)	
× Incumbent malfeasant spending				-0.029** (0.013)	-0.028*** (0.010)	
× Unfavorable incumbent updating						-0.005*** (0.002)
Outcome range	[0.03,0.47]	[0.03,0.47]	[0.03,0.47]	[0.03,0.47]	[0.03,0.47]	[0.03,0.47]
Control outcome mean	0.19	0.20	0.20	0.19	0.20	0.20
Control outcome std. dev.	0.07	0.07	0.07	0.07	0.07	0.07
R ²	0.62	0.61	0.61	0.62	0.61	0.61
Interaction range		[-1.4,1.1]	[2.4,3.8]	[0.58]		[-0.6,2.7]
Interaction mean		-0.10	3.23	0.21		0.91
Interaction std. dev.		0.83	0.26	0.17		1.00
Observations	675	651	651	675	651	651

Notes: All specifications include block fixed effects, and are estimated using OLS. Lower-order interaction terms are absorbed by the block fixed effects. The smaller sample in Columns (2), (3), and (5) reflect the lack of data on prior beliefs about the incumbent party in Apaseo el Alto. Standard errors clustered by municipality-treatment are in parentheses. * denotes $p < 0.1$, ** denotes $p < 0.05$, *** denotes $p < 0.01$.

where voters possessed less precise prior beliefs. Furthermore, columns (3) and (4) shows that there was a significant increase in incumbent vote share among respondents in municipalities with below-median prior precision (i.e. below 3.25) and no detectable effect in municipalities where the precision of prior beliefs was greatest.²⁷ Together, this evidence suggests that, given their low expectations of politicians, voters did not substantially improve their perception of incumbent party malfeasance, on average,

27. As with column (2) of Table 3, this split sample approach addresses the concern that a combination of block fixed effects perfectly explain control group responses.

TABLE 5. Effect of information treatment on the precision of voters' posterior beliefs about incumbent party malfeasance.

	Precision of perceived incumbent party malfeasance (very low - very high)			
	(1)	(2)	(3) Above-mean incumbent prior precision	(4) Below-mean incumbent prior precision
Information treatment	0.016 (0.024)	0.675** (0.265)	-0.020 (0.041)	0.050* (0.026)
× Incumbent prior precision		-0.204** (0.084)		
Outcome range	{1,2,3,4}	{1,2,3,4}	{1,2,3,4}	{1,2,3,4}
Control outcome mean	3.25	3.25	3.51	2.94
Control outcome std. dev.	0.85	0.85	0.72	0.88
Interaction range		[2.4,3.8]		
Interaction mean		3.23		
Interaction std. dev.		0.26		
R^2	0.11	0.11	0.06	0.06
Observations	4,673	4,673	2,429	2,244

Notes: All specifications include block fixed effects, and are estimated using OLS. Lower-order interaction terms are absorbed by the block fixed effects. Standard errors clustered by municipality-treatment are in parentheses. * denotes $p < 0.1$, ** denotes $p < 0.05$, *** denotes $p < 0.01$.

but nevertheless became more likely to vote for an incumbent party that represented a less risky option.

Another possible explanation is that voting behavior reflected general equilibrium considerations, including incumbent and challenger parties reactions to the information's provision. Previous studies have found that information provision reduced vote buying in India (Banerjee et al. 2011), although the opposite occurred in the Philippines (Cruz et al. 2021). Bowles and Larreguy (2020) and Bidwell et al. (2020) also suggest that candidates adjusted their on-the-ground campaigning after debates in Liberia and Sierra Leone. While incorporating such reactions in our overall point estimates may capture the primary parameter of policy interest, it remains important to understand whether the mechanism reflects belief updating in responses to the information's content or campaign responses to the information disseminated.

We examine this systematically by asking voters whether incumbents and challengers referred to the information reported in our leaflets in any of the following (non-exclusive) ways: (i) campaign activities; (ii) partisan leaflets; (iii) visits from local political actors; (iv) advertisements; or (v) through the media. Around 17% of voters reported experiencing at least one type of incumbent response, and 16% reported at least one type of challenger response. According to our respondents, incumbents most frequently claimed that all parties were equally bad, while opposition parties were

TABLE 6. Effect of information treatment on political party responses.

	(1)	(2)	Total party activities		(5)	(6)
			(3)	(4)		
Panel A: Incumbent reactions						
Information treatment	0.032 (0.043)	0.034 (0.043)	0.681* (0.348)	-0.131* (0.077)	0.439 (0.296)	-0.001 (0.069)
× Incumbent malfeasance prior		0.024 (0.038)			0.018 (0.032)	
× Incumbent prior precision			-0.204* (0.111)		-0.177* (0.096)	
× Incumbent malfeasant spending				0.766*** (0.258)	0.755*** (0.230)	
× Unfavorable incumbent updating						0.036 (0.040)
Control outcome mean	0.43	0.46	0.43	0.43	0.43	0.43
Control outcome std. dev.	1.18	1.17	1.18	1.18	1.18	1.18
R^2	0.12	0.12	0.12	0.12	0.12	0.12
Panel B: Challenger reactions						
Information treatment	0.102** (0.039)	0.105*** (0.039)	0.609 (0.398)	-0.024 (0.060)	0.400 (0.384)	0.089 (0.060)
× Incumbent malfeasance prior		0.033 (0.043)			0.029 (0.038)	
× Incumbent prior precision			-0.159 (0.122)		-0.132 (0.116)	
× Incumbent malfeasant spending				0.591*** (0.204)	0.588*** (0.187)	
× Unfavorable incumbent updating						0.014 (0.036)
Control outcome mean	0.40	0.48	0.40	0.40	0.40	0.40
Control outcome std. dev.	1.17	1.24	1.17	1.17	1.17	1.17
R^2	0.12	0.12	0.12	0.12	0.12	0.12
Outcome range	{0,1,2,3,4,5}	{0,1,2,3,4,5}	{0,1,2,3,4,5}	{0,1,2,3,4,5}	{0,1,2,3,4,5}	{0,1,2,3,4,5}
Interaction range		[-1.4,1.1]	[2.0,3.8]	[0.0,0.58]	[-0.6,2.7]	[0.0,0.58]
Interaction mean		-0.09	3.18	0.21		0.90
Interaction std. dev.		0.80	0.35	0.17		0.97
Observations	4,958	4,958	4,958	4,808	4,958	4,958

Notes: All specifications include block fixed effects, and are estimated using OLS. Lower-order interaction terms are absorbed by the block fixed effects. Standard errors clustered by municipality-treatment are in parentheses. * denotes $p < 0.1$, ** denotes $p < 0.05$, *** denotes $p < 0.01$.

more likely to emphasize the content of the leaflets. Our outcome of interest is the total number of politician responses reported by the respondent, ranging from 1 to 5.

Column (1) of Table 6 shows that incumbents, and especially challengers, engaged in more campaign activities in treated precincts.²⁸ The increase is only statistically significant among challengers, and the effect magnitude is relatively small in each case. For politician responses to explain the positive average effect, the incumbent's resource advantage would need to make their responses substantially more effective (e.g. Cruz et al. 2021) or challenger reactions would need to have backfired. Thus, while campaigns did respond somewhat to information dissemination, it is unlikely that these responses drove average voter behavior. Moreover, our evidence of voter learning—to

28. The non-zero number of activities in the control group likely reflects recall failures.

which we soon turn—suggests that political responses cannot account for heterogeneity in treatment effects by prior beliefs, updating and incumbent-malfeasance levels.

It is also unlikely that the positive average effect reflects other potential explanations. One possibility is that voters (wrongly) credited the incumbent party for attracting FISM resources to their municipality. However, we find little support for this interpretation in Appendix Table A.20, which shows no heterogeneous effects by the quantity of FISM funds received by the municipality, in either absolute or per voter terms. Another possibility is that the intervention may have been perceived as a smear campaign against the incumbent party. However, as shown above, voters nevertheless updated their beliefs and generally thought the information came from a non-partisan source. Finally, the treatment could have altered the weight that voters attached to different issues when deciding how to cast their ballot. However, Appendix Table A.21 finds no evidence to suggest that a candidate's honesty or likelihood of addressing poverty became more important to treated voters.

6.2. *Heterogeneous Effects of Information on Incumbent Vote Share*

Although treated precincts somewhat surprisingly rewarded incumbent parties *on average*, we next demonstrate that—as theorized (and prespecified)—voting behavior on the margin varies with the information content received and in line with the changes in posterior beliefs documented in our survey data.

First, supporting hypothesis H2, the information treatment's largest positive effects were detected where voters initially believed that their incumbent was more malfeasant. Across both panels in Table 4, column (2) shows that the treatment's positive effect on incumbent party vote share was significantly greater (smaller) in precincts within municipalities where the control group had more unfavorable (favorable) prior beliefs regarding the incumbent party's level of malfeasance. Our estimates indicate that moving from the municipality with the most favorable prior beliefs about the incumbent party (-1.4) to the municipality with the most unfavorable prior beliefs (1.1) increased the effect of providing information on the incumbent party's vote share from 0.6 to 2.9 percentage points and the effect on the incumbent's share of registered votes from 0.1 to 1.4 percentage points. Also consistent with H2 and the risk-reduction explanation for the positive average effect on the incumbent's vote share, column (3) reports a significantly smaller positive effect of the information in precincts where the municipality's control respondents had more precise prior beliefs.

Second, and further supporting H2, treated voters were more likely to vote for incumbents overseeing lower levels of malfeasance. The significant negative interaction in column (4) between the treatment and the share of malfeasant spending reported in the leaflet implies that a one-standard-deviation increase in the share of malfeasance reduced the positive effect of treatment on the incumbent party's vote share (as a share of turnout) by 0.9 percentage points. As illustrated in Figure 6, revealing any level of malfeasant spending below 35% significantly increased the incumbent's vote share. The effect of providing information is never meaningfully negative in our sample, where the highest level of reported malfeasance is 58%. However, between

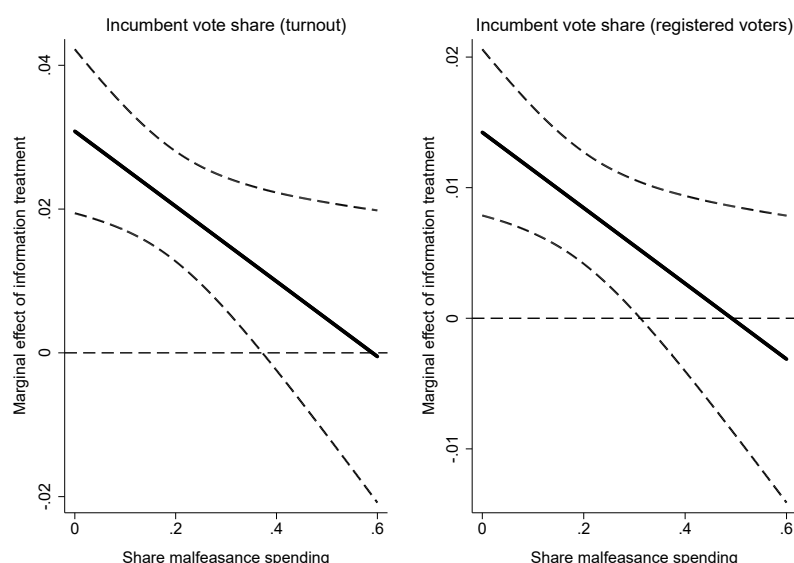


FIGURE 6. Marginal effect of information treatment on incumbent party vote share, by share of malfeasance spending (95% confidence intervals).

2007 and 2015, 46 audited municipalities (3.4% of all audited municipalities) across the country registered malfeasance rates exceeding 60%. Electoral sanctions could occur in such settings where more extreme levels of malfeasance are reported. Column (5) demonstrates that these results are robust to simultaneously adjusting for interactions with voters' prior beliefs.

Third, and combining the preceding heterogeneous effects, the effect of revealing incumbent malfeasance information decreases with the overall extent to which voters unfavorably updated their beliefs about the incumbent party's malfeasance. Column (6) of both panels reports a significant negative interaction between the treatment and our measure of unfavorable updating of posterior beliefs in each municipality's control precincts. A one-standard-deviation increase in unfavorable updating induced by the information reduced the incumbent party's vote share (as a share of turnout) by 0.9 percentage points in treated precincts.

We observe broadly similar responses to revelations concerning spending that did not benefit the poor and revelations pertaining to unauthorized spending. Appendix Table A.22 splits the sample between municipalities that received information about not spending FISM funds on projects that benefited the poor and spending on unauthorized projects. There is a clear positive average effect of our treatment across both types of malfeasance spending, while the interactions with the share of malfeasance spending and

the extent of unfavorable updating point in the same direction across subsamples.²⁹ Moreover, we show in Appendix Table A.25 that the results in the full sample are driven by the specific dimension of malfeasance that was reported to voters, rather than treatment raising the salience of other dimensions of malfeasance more generally.

These findings fit closely with voting behavior reflecting updated beliefs driven by signals of incumbent malfeasance. However, it is also possible that vote choices could instead reflect voter reactions to differences in party campaign strategies across different types of treated and control precincts. Indeed, the modest increases in political activity documented in column (1) of Table 6 mask significant heterogeneity in response to the distribution of Borde Político's leaflets. The large and significant positive interactions in column (4) demonstrate that, for both incumbents and challengers, party activity increased substantially in municipalities in which high levels of malfeasance were revealed. In a treated precinct within a municipality with 50% malfeasant spending, activity almost doubled relative to a municipality with 0%. However, these party responses cannot fully account for the heterogeneous effects of treatment attributed to voters learning from the leaflets. A comparison of column (4) with columns (2), (3), (5), and (6) in Table 6 shows that political responses are driven by the *level* of malfeasance reported, rather than voters' prior beliefs or the extent to which these were updated based on this information. This suggests that incumbent parties and their operatives may not know the extent to which voters expect their representatives to engage in minimal malfeasant spending while in office. The preceding evidence of increased party activity thus suggests that party responses could play a role in shaping how information dissemination impacts incumbent party support, but an important component is nevertheless explained by voters updating directly from the leaflets.

6.3. Robustness Tests

Table 7 demonstrates that the incumbent party vote share results are robust to several alternative specifications. We focus on vote share as a share of turnout, but Appendix Table A.28 reports similar results for incumbent party vote share using registered voters in the denominator.

First, we address the concern that aggregating prior beliefs and belief updating at the municipal level fails to capture meaningful variation in beliefs across precincts within municipalities. To the extent that using municipal-level aggregates of the prior beliefs introduces classical error in the precinct level analysis, such aggregation may lead to the underestimation of heterogeneous effects across precincts. We combat this issue by using our survey data from control precincts, which are uncontaminated by treatment, to impute precinct-level prior beliefs across the sample. The 14 covariates underpinning

29. The lack of heterogeneity in electoral response by precinct socioeconomic development also indicates that misallocating funds to projects that did not benefit the poor is no less salient where voters were less likely to directly benefit from FISM projects themselves. This suggests that voters, at least in our experimental sample, primarily worry about malfeasance in terms of incumbent integrity or competence, rather than its distributive implications.

our prediction model are described in Appendix Section A.9.2, and explain at least 50% of variation in precinct-level prior beliefs, prior belief precision, and belief updating upon viewing the leaflet. Panel A of Table 7 shows that these precinct-specific predicted beliefs yield similar results to our main estimates, suggesting that measurement error in the precinct-level regressions due to municipal-level aggregation of prior beliefs is not a major challenge for estimation. Appendix Table A.27 shows that a similar individual-level prediction exercise also generates similar results.

Second, since only the provision of audit information was randomized, it is possible that our heterogeneous effects could be confounded by correlates of voters' prior beliefs and the level of municipal malfeasance. In particular, our estimates could be biased if voters' prior beliefs correlate with potential confounds relating to the extent of treatment dissemination, the ease with which our information treatment could be relayed through local networks, alternative sources of our information, and the level of political polarization. Similarly, the content of the ASF's report could be correlated with structural factors that affect voters' expectations of government service provision or the welfare consequences of malfeasant spending, and in turn shape their response to treatment.

To address these concerns, we adjust for the interaction between our information treatment and potential confounders. We start with the following (demeaned) precinct-level covariates: share of the precinct electorate that received a leaflet, distance to the municipality center, whether a precinct is rural, population density, number of radio and television stations covering the precinct that transmit from within the municipality, percentage of households with access to a television, percentage of households with access to the internet at home, and municipal winning margin in the previous election. An important caveat is that some of these variables could themselves determine voters' prior beliefs, and this could contaminate our estimates by partialing out part of the effect of prior beliefs. In addition, we examine robustness to including the following (demeaned) municipal-level covariates: the number of registered voters; population density; working age share of the population; average years of schooling; the share of households with televisions; the share of households with access to the internet; and municipal incumbent vote share and victory margin at the previous election.

The results in panels B and C demonstrate that our heterogeneous effects are generally robust, supporting our interpretation that the results reflect voter learning. Only in the case of the interaction with the precision of prior beliefs in panel B are the results somewhat sensitive to the inclusion of interactive controls. Further analyses suggest that this sensitivity primarily reflects interactively adjusting for our proxies for ruralness, which could be determinants of prior precision as well as information flows within voter networks.

Third, we show that the effects become somewhat larger when we account for heterogeneity in the share of voters that actually received leaflets. Specifically, we weight precinct-level observations by the share of voters to whom we delivered a leaflet. In control precincts, we compute this share based on the average number of leaflets delivered to treated precincts within the same block. This weighting scheme downweights large precincts in which only a small fraction of voters could receive

TABLE 7. Robustness of information treatment on incumbent party vote share (share of turnout).

	(1)	Incumbent party vote share (share of turnout)				
	(2)	(3)	(4)	(5)	(6)	
Panel A: Predicted precinct-level prior beliefs and updating						
Information treatment	0.020*** (0.004)	0.022*** (0.004)	0.135*** (0.045)	0.031*** (0.006)	0.128*** (0.037)	0.033*** (0.005)
× Incumbent malfeasance prior (predicted)		0.009* (0.005)			0.006 (0.004)	
× Incumbent prior precision (predicted)			-0.036** (0.014)		-0.030** (0.011)	
× Incumbent malfeasant spending				-0.052** (0.023)	-0.052*** (0.017)	
× Unfavorable incumbent updating (predicted)						-0.009** (0.004)
Panel B: Adjusting for (demeaned) precinct-level covariates interacted with information treatment						
Information treatment	0.018*** (0.003)	0.018*** (0.003)	0.105** (0.052)	0.030*** (0.006)	0.131** (0.052)	0.026*** (0.004)
× Incumbent malfeasance prior		0.008** (0.004)			0.007* (0.004)	
× Incumbent prior precision			-0.027* (0.016)		-0.031* (0.016)	
× Incumbent malfeasant spending				-0.058** (0.025)	-0.068*** (0.021)	
× Unfavorable incumbent updating						-0.010*** (0.003)
Panel C: Adjusting for (demeaned) municipal-level covariates interacted with information treatment						
Information treatment	0.020*** (0.004)	0.019*** (0.003)	0.146*** (0.037)	0.040*** (0.005)	0.151*** (0.032)	0.031*** (0.006)
× Incumbent malfeasance prior		0.011 (0.008)			0.003 (0.006)	
× Incumbent prior precision			-0.040*** (0.012)		-0.035*** (0.010)	
× Incumbent malfeasant spending				-0.093*** (0.020)	-0.093*** (0.016)	
× Unfavorable incumbent updating						-0.013** (0.006)
Panel D: Weighting observations by the (expected) share of the precinct that received a leaflet						
Information treatment	0.026*** (0.006)	0.024*** (0.005)	0.168*** (0.062)	0.042*** (0.007)	0.144*** (0.046)	0.034*** (0.006)
× Incumbent malfeasance prior		0.012** (0.006)			0.008* (0.004)	
× Incumbent prior precision			-0.045** (0.019)		-0.033** (0.014)	
× Incumbent malfeasant spending				-0.072*** (0.027)	-0.071*** (0.018)	
× Unfavorable incumbent updating						-0.012** (0.004)

Notes: All specifications include block fixed effects, and are estimated using OLS. See text for interactive covariates included in panels B and C. Observations in panel D are weighted by the share of the precinct that was treated. Lower-order interaction terms are omitted. Standard errors clustered by municipality-treatment are in parentheses. * denotes $p < 0.1$, ** denotes $p < 0.05$, *** denotes $p < 0.01$.

the leaflet. Consistent with our results being driven by exposure to treatment, panel D reports larger point estimates across all specifications. Similarly, we show in Appendix Table A.29 that the magnitude of the average and heterogeneous effects of information provision generally increase with the share of voters in a precinct that received a leaflet.

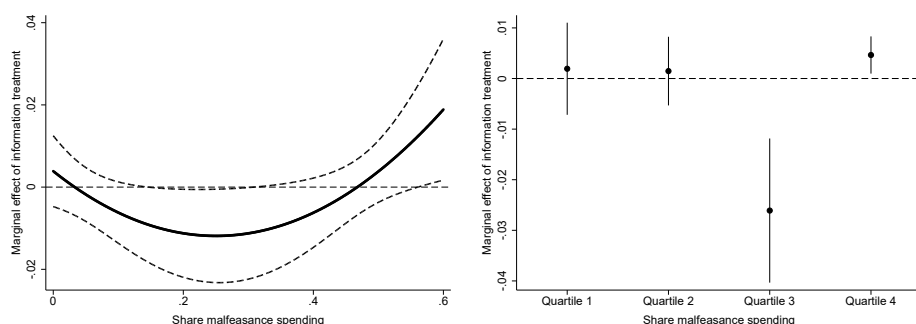


FIGURE 7. Marginal effect of information treatment on turnout, by share of malfeasant spending (95% confidence intervals).

6.4. Non-monotonic Effects of Information on Turnout

A distinctive feature of our theory is the non-monotonic relationship between the extent of malfeasance and turnout captured in H3. In particular, we predicted that revelations of either extremely low or high levels of malfeasance would induce significant masses of voters to strongly prefer a particular party and thereby reduce the number of largely indifferent voters that abstain. Signals of incumbent malfeasance that induce small shifts in the distribution of voter preferences could instead reduce turnout when a mode of voters become relatively indifferent between parties.

The heterogeneous effects reported in Table 8 find support for these predictions. Since we expected non-monotonic effects, it is not surprising to observe in columns (1) and (2) that the average ITT effect is close to zero and does not vary linearly with the share of malfeasant spending that was reported. Rather, we first focus on the interaction between treatment and a quadratic operationalization of reported malfeasance that allows us to detect the prespecified non-monotonic effect. At malfeasance levels close to 0%, the lower-order treatment term in column (3) shows that turnout increased by 0.4 percentage points. The negative linear and positive quadratic interactions with the share of malfeasant spending demonstrate that turnout decreased at interim levels of malfeasance—which conform more closely with voters’ prior beliefs—but increased by more than a percentage point at high levels of malfeasance. In this specification, the positive effect at the lowest and highest levels of malfeasance in our sample is not statistically significant, although we observe significant increases of at least a percentage point at each extremity when observations are weighted by the share of the precinct that received a leaflet, as Appendix Table A.31 shows. Column (4) reports similar results—with a statistically significant increase in turnout for high levels of malfeasance—when splitting the sample into quartiles by level of reported malfeasance, and thus demonstrates that the results are not an artifact of imposing a quadratic specification. Figure 7 depicts both the quadratic and non-parametric non-monotonic relationships graphically.

TABLE 8. Effect of information treatment on turnout and confidence in the electoral process.

Panel A: Turnout	Turnout					
	(1)	(2)	(3)	(4)		
Information treatment	-0.004 (0.003)	-0.005 (0.004)	0.004 (0.004)	0.002 (0.005)		
× Incumbent malfeasance spending		0.002 (0.012)	-0.126** (0.059)			
× Incumbent malfeasance spending squared			0.251** (0.111)			
× Incumbent malfeasance spending quartile 2				-0.000 (0.006)		
× Incumbent malfeasance spending quartile 3				-0.028*** (0.008)		
× Incumbent malfeasance spending quartile 4				0.003 (0.005)		
Outcome range	[0.21,0.79]	[0.21,0.79]	[0.21,0.79]	[0.21,0.79]		
Control outcome mean	0.50	0.50	0.50	0.50		
Control outcome std. dev.	0.10	0.10	0.10	0.10		
Interaction range		[0,0.58]	[0,0.58]			
Interaction mean		0.21	0.21			
Interaction std. dev.		0.17	0.17			
R ²	0.68	0.68	0.68	0.68		
Observations	675	675	675	675		
Panel B: Confidence in the system	Elections help to select competent candidates (did not help at all - helped a lot)					
	(1)	(2)	(3)	(4)	(5)	(6)
Information treatment	0.008 (0.042)	-0.000 (0.041)	0.389 (0.511)	0.052 (0.078)	0.712 (0.517)	-0.044 (0.054)
× Incumbent malfeasance prior		-0.078 (0.049)			-0.100** (0.048)	
× Incumbent prior precision			-0.118 (0.158)		-0.205 (0.163)	
× Incumbent malfeasant spending				-0.209 (0.255)	-0.247 (0.229)	
× Unfavorable incumbent updating						0.057 (0.038)
Outcome range	{1,2,3,4,5}	{1,2,3,4,5}	{1,2,3,4,5}	{1,2,3,4,5}	{1,2,3,4,5}	{1,2,3,4,5}
Control outcome mean	2.86	2.86	2.84	2.86	2.86	2.86
Control outcome std. dev.	1.40	1.40	1.41	1.40	1.40	1.40
Interaction range		[-1.4,1.18]	[2.4,3.8]	[0,0.58]		[-0.6,2.7]
Interaction mean		-0.10	3.23	0.21		0.91
Interaction std. dev.		0.82	0.26	0.17		1.00
R ²	0.06	0.06	0.06	0.06	0.06	0.06
Observations	4,615	4,615	4,615	4,615	4,615	4,615

Notes: All specifications include block fixed effects, and are estimated using OLS. Lower-order interaction terms are absorbed by the block fixed effects. Standard errors clustered by municipality-treatment are in parentheses. * denotes $p < 0.1$, ** denotes $p < 0.05$, *** denotes $p < 0.01$.

While the preceding results support our model's emphasis on the importance of voters' prior expectations, it is possible that malfeasance revelations could induce voters to disengage with politics. This is most plausible where high levels of malfeasance are reported (Chong et al. 2015). However, we find no evidence to suggest that information about an incumbent's malfeasance induced a general form of disengagement with the political system. As previously noted, column (2) of panel A in Table 8 shows that turnout does not linearly decrease with the level of malfeasant spending. Furthermore, we turn to our survey data to examine voter responses on a five-point scale rating their belief that elections help to select honest and competent

politicians. Columns (4) and (6) of panel B show that neither high levels of malfeasance nor unfavorable updating significantly altered treated voters' faith that elections can select good candidates.³⁰

Another possibility is that the non-monotonic effects on turnout could reflect party mobilization strategies, rather than voter updating. However, the results in Table 6 indicate that party campaign responses were concentrated in high-malfeasance municipalities, and thus cannot explain how the treatment could induce high turnout in low-malfeasance municipalities. In sum, our findings suggest that malfeasance revelations which substantially deviate from the average voter's prior belief can increase turnout by inducing voters to shift not just towards indifference but to support other parties instead.

7. Conclusion

This article demonstrates the importance of voters' prior beliefs in understanding when incumbent malfeasance revelations affect electoral accountability. We find that Mexican voters—who, like voters in many developing contexts, have low expectations that their incumbents will correctly allocate resources—on average actually *reward* municipal incumbent parties revealed to have engaged in non-trivial levels of malfeasance in office. Although the increase in incumbent support on average likely reflects uncertainty-reduction among risk averse voters, we also document considerable support for our simple learning model. In particular, information provision was significantly more likely to increase incumbent support among voters who possessed unfavorable prior beliefs about incumbent party malfeasance, who learn of lower incumbent malfeasance, and who update their beliefs about the incumbent most favorably. Furthermore, and consistent with our theoretical model, the effect of information provision on turnout varies non-monotonically with the signal: surprising information increases turnout by shifting voters between parties, and relatively unsurprising information shifts voters towards indifference. By emphasizing voters' prior beliefs, and their relationship with the content of the information, these findings can help explain the mixed evidence that information induces electoral sanctioning or impacts political participation in developing democracies.

The implications of our findings for using information interventions to improve governance are mixed. A clear reason for optimism is that voters are able to understand signals of incumbent malfeasance and incorporate them into their voting behavior in an approximately Bayesian manner. Fixing voters' expectations of the parties, information thus helps voters to choose between candidates. However, the fact that some voters are so pessimistic that the misallocation of up to 40% of funds is considered good news is worrying for proponents of good governance. As the mixed evidence from previous

30. Unreported results including a quadratic interaction with incumbent malfeasance spending provide no evidence to suggest that confidence mirrors the non-monotonic relationship with turnout.

studies suggests, such beliefs may not be uncommon in developing contexts—and may be consistent with incumbent behavior (e.g. Caselli and Morelli 2004).

In this light, our findings suggest a need to improve voters' expectations of their elected representatives, which could also induce politicians to perform better in office in the long run (Barro 1973; Ferejohn 1986), the need for better politicians to stand for office, or the need for more effective audits and legal sanctions. Civic education or a critical media may be required to help voters understand what good performance entails (e.g. Adida et al. 2020; Botero et al. 2015; Gottlieb 2016). Higher-quality candidates should also be encouraged to stand for office; some evidence suggests that increased wages can help (Caselli and Morelli 2004; Gagliarducci and Nannicini 2013). More effective audits and legal sanctions may also help improve politicians' performance by causing parties to believe that they will be electorally sanctioned for malfeasance in office (Avis et al. 2018; Bobonis et al. 2016; Olken 2007; Zamboni and Litschig 2018).

Finally, our study underscores the importance of investigating equilibrium political responses for understanding the impact of informational interventions. As with several other recent studies (Banerjee et al. 2011; Bowles and Larreguy 2020; Bidwell et al. 2020; Cruz et al. 2021), we document evidence that politicians respond to such interventions. Although the patterns of political responses in this study do not confound our capacity to isolate effects attributable to belief updating, such responses could be consequential in terms of partially explaining the average effects. Moreover, it is interesting to find that, while politicians do respond to informational interventions in an attempt to counteract their electoral consequences, their responses do not address the sophisticated way in which voters process the information provided. To better understand when informational interventions are effective, political responses to non-partisan informational dissemination demand further attention.

References

- Adida, Claire, Jessica Gottlieb, Eric Kramon, and Gwyneth McClendon (2020). "When Does Information Influence Voters? The Joint Importance of Salience and Coordination." *Comparative Political Studies*, 53(6), 851–891.
- Arias, Eric, Pablo Balán, Horacio Larreguy, John Marshall, and Pablo Querubín (2019). "Information Provision, Voter Coordination, and Electoral Accountability: Evidence from Mexican Social Networks." *American Political Science Review*, 113(2), 475–498.
- Arias, Eric, Horacio Larreguy, John Marshall, and Pablo Querubín (2018). "Does the Content and Mode of Delivery of Information Matter for Electoral Accountability? Evidence from a Field Experiment in Mexico." Working Paper, Harvard University.
- Auditoría Superior de la Federación (2014). "Informe del Resultado de la Fiscalización Superior de la Cuenta Pública 2012.", URL <https://goo.gl/fR5nmt>. Audit Summary Report.
- Avis, Eric, Claudio Ferraz, and Frederico Finan (2018). "Do Government Audits Reduce Corruption? Estimating the Impacts of Exposing Corrupt Politicians." *Journal of Political Economy*, 126(5), 1912–1964.
- Banerjee, Abhijit V., Selvan Kumar, Rohini Pande, and Felix Su (2011). "Do Informed Voters Make Better Choices? Experimental Evidence from Urban India.", URL <https://www.povertyactionlab.org/>

- org/sites/default/files/research-paper/142\%20-\%20informed\%20voters\%20Nov2011.pdf. Working Paper, Pennsylvania State University.
- Barro, Robert J. (1973). "The Control of Politicians: An Economic Model." *Public Choice*, 14(1), 19–42.
- Bidwell, Kelly, Katherine Casey, and Rachel Glennerster (2020). "Debates: Voting and Expenditure Responses to Political Communication." *Journal of Political Economy*, 128(8), 2880–2924.
- Boas, Taylor C., F. Daniel Hidalgo, and Marcus André Melo (2019). "Norms versus Action: Why Voters Fail to Sanction Malfeasance in Brazil." *American Journal of Political Science*, 63(2), 385–400.
- Bobonis, Gustavo J, Luis Cámara R Fuertes, and Rainer Schwabe (2016). "Monitoring Corruptible Politicians." *American Economic Review*, 106(8), 2371–2405.
- Botero, Sandra, Rodrigo Castro Cornejo, Laura Gamboa, Nara Pavao, and David W. Nickerson (2015). "Says Who? An Experiment on Allegations of Corruption and Credibility of Sources." *Political Research Quarterly*, 68(3), 493–504.
- Bowles, Jeremy and Horacio Larreguy (2020). "Who Debates, Who Wins? At-Scale Experimental Evidence on the Supply of Policy Information in a Liberian Election." *Working Paper, Université Toulouse 1 Capitole*.
- Brennan, Geoffrey and Alan Hamlin (1998). "Expressive Voting and Electoral Equilibrium." *Public Choice*, 95(1-2), 149–175.
- Buntaine, Mark T., Ryan Jablonski, Daniel L. Nielson, and Paula M. Pickering (2018). "SMS texts on corruption help Ugandan voters hold elected councillors accountable at the polls." *Proceedings of the National Academy of Sciences*, 11(26), 6668–6673.
- Caselli, Francesco and Massimo Morelli (2004). "Bad politicians." *Journal of Public Economics*, 88(3), 759–782.
- Chang, Eric C.C., Miriam A. Golden, and Seth J. Hill (2010). "Legislative Malfeasance and Political Accountability." *World Politics*, 62(2), 177–220.
- Chong, Alberto, Ana De La O, Dean Karlan, and Leonard Wantchekon (2015). "Does Corruption Information Inspire the Fight or Quash the Hope? A Field Experiment in Mexico on Voter Turnout, Choice and Party Identification." *Journal of Politics*, 77(1), 55–71.
- Cruz, Cesi, Philip Keefer, and Julien Labonne (2021). "Buying Informed Voters: New Effects of Information on Voters and Candidates." *Economic Journal*, 131(635), 1105–1134.
- de Figueiredo, Miguel F.P., F. Daniel Hidalgo, and Yuri Kasahara (2014). "When Do Voters Punish Corrupt Politicians? Experimental Evidence from Brazil." URL <https://goo.gl/rWXBbm>. Working Paper, Harvard University.
- Dunning, Thad, Guy Grossman, Macartan Humphreys, Susan Hyde, Craig McIntosh, and Gareth Nellis (2019). *Information, Accountability, and Cumulative Learning: Lessons from Metaketa I*. Cambridge University Press.
- Fearon, James D. (1999). "Electoral Accountability and the Control of Politicians: Selecting Good Types versus Sanctioning Poor Performance." In *Democracy, Accountability, and Representation*, edited by Adam Przeworski, Susan Stokes, and Bernard Manin. Cambridge University Press.
- Ferejohn, John (1986). "Incumbent Performance and Electoral Control." *Public Choice*, 50(1), 5–25.
- Ferraz, Claudio and Frederico Finan (2008). "Exposing corrupt politicians: The effects of Brazil's publicly released audits on electoral outcomes." *Quarterly Journal of Economics*, 123(2), 703–745.
- Gagliarducci, Stefano and Tommaso Nannicini (2013). "Do Better Paid Politicians Perform Better? Disentangling Incentives from Selection." *Journal of the European Economic Association*, 11(2), 369–398.
- Gottlieb, Jessica (2016). "Greater Expectations: A Field Experiment to Improve Accountability in Mali." *American Journal of Political Science*, 60(1), 143–157.
- Hsieh, Chang-Tai and Enrico Moretti (2006). "Did Iraq Cheat the United Nations? Underpricing, Bribes, and the Oil for Food Program." *Quarterly Journal of Economics*, 121(4), 1211–1248.

- Humphreys, Macartan and Jeremy Weinstein (2012). "Policing Politicians: Citizen Empowerment and Political Accountability in Uganda Preliminary Analysis.", URL <https://goo.gl/EiVitc>. Working Paper, Columbia University.
- Kendall, Chad, Tommaso Nannicini, and Francesco Trebbi (2015). "How Do Voters Respond to Information? Evidence from a Randomized Campaign." *American Economic Review*, 105(1), 322–53.
- Khemani, Stuti, Ernesto Dal Bó, Claudio Ferraz, Frederico S. Finan, Johnson Stephenson, Louise Corinne, Adesinaola M. Odugbemi, Dikshya Thapa, and Scott D. Abrahams (2016). "Making politics work for development: Harnessing transparency and citizen engagement.", URL <https://goo.gl/5n58cG>. World Bank Policy Research Report 106337.
- Laakso, Markku and Rein Taagepera (1979). "'Effective' Number of Parties: A Measure with Application to West Europe." *Comparative Political Studies*, 12(1), 3–27.
- Langston, Joy (2003). "Rising from the Ashes? Reorganizing and Unifying the PRI's State Party Organizations after Electoral Defeat." *Comparative Political Studies*, 36(3), 293–318.
- Larreguy, Horacio, John Marshall, and Jr. Snyder, James M. (2018). "Leveling the Playing Field: How Equalizing Access to Campaign Advertising Helps Locally Non-Dominant Parties in Consolidating Democracies." *Journal of the European Economic Association*, 16(6), 1812–1849.
- Larreguy, Horacio, John Marshall, and Jr. Snyder, James M. (2020). "Publicizing Malfeasance: When the Local Media Structure Facilitates Electoral Accountability in Mexico." *Economic Journal*, 130(631), 2291–2327.
- Leon, Gianmarco (2017). "Turnout, Political Preferences and Information: Experimental Evidence from Peru." *Journal of Development Economics*, 127, 56–71.
- Morris, Stephen and Hyun Song Shin (2002). "Social Value of Public Information." *American Economic Review*, 92(5), 1521–1534.
- Olken, Benjamin A. (2007). "Monitoring Corruption: Evidence from a Field Experiment in Indonesia." *Journal of Political Economy*, 115(2), 200–249.
- Pons, Vincent (2018). "Will a Five-Minute Discussion Change your Mind? A Countrywide Experiment on Voter Choice in France." *American Economic Review*, 108(6), 1322–1363.
- Rogoff, Kenneth (1990). "Equilibrium Political Budget Cycles." *American Economic Review*, 80(1), 21–36.
- Spenkuch, Jörg L and David Toniatti (2018). "Political Advertising and Election Results." *Quarterly Journal of Economics*, 133(4), 1981–2036.
- Wellenstein, Anna, Angélica Núñez, and Luis Andrés (2006). "Social Infrastructure: Fondo de Aportaciones para la Infraestructura Social (FAIS)." In *Decentralized service delivery for the poor, Volume II: Background papers*, edited by The World Bank, pp. 167–222. Mexico City: The World Bank.
- Zamboni, Yves and Stephan Litschig (2018). "Audit risk and rent extraction: Evidence from a randomized evaluation in Brazil." *Journal of Development Economics*, 134, 133–149.