

When Does Information Increase Electoral Accountability? Lessons from a Field Experiment in Mexico

Eric Arias, Horacio Larreguy, John Marshall,
and Pablo Querubín

In this chapter we report the results of a randomized intervention conducted during the 2015 Mexican elections. We teamed up with a local, nonpartisan, and transparency-focused nongovernmental organization (NGO) – Borde Político¹ – and conducted a large-scale informational campaign providing citizens with information from official audits revealing the way in which their municipal governments had spent funds intended for infrastructure projects benefiting impoverished localities. While official information on audits of government expenditures is publicly available in Mexico, citizens rarely have access to this information and when they do, the information is often not provided in a way that enables them to evaluate the performance of their politicians.

The material presented in this chapter closely follows Arias et al. (n.d.), where we originally reported our research design and the main results for a larger sample. However, in order to facilitate a comparison between our study and the other Metaketa studies presented in this volume, we focus on the variant of our intervention in which we presented information for the municipality’s incumbent party paired with information regarding other municipalities in the state, but governed by a different party.

¹ Borde Político is a leading NGO seeking to increase voter knowledge about the actions of their politicians in office, with significant experience in the development of web-based platforms providing politically relevant information to individuals (see <http://borde.mx>).

5.1 LOCAL ELECTIONS AND ELECTORAL ACCOUNTABILITY IN MEXICO

Mexico's federal system is divided into thirty-one states (and the Federal District of Mexico City) containing around 2,500 municipalities and 67,000 electoral precincts. Following major decentralization reforms in the 1990s,² municipal governments – the focus of this chapter – play an important role in delivering basic public services and managing local infrastructure, and account for 20 percent of total government spending. Municipalities are governed by mayors typically elected to three-year nonrenewable terms, although re-election will become possible for incumbents in some states starting in 2018. Traditionally, local political competition in Mexico 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 Party of the Democratic Revolution (PRD). Due to regional bases of political support and the often localized nature of municipal politics, municipal elections are typically dominated by one or two main parties. For example, the average Mexican municipality has only 2.5 effective parties.

A key component of a mayor's budget is the Municipal Fund for Social Infrastructure (FISM), which represents 24 percent of the average municipality's budget. According to the 1997 Fiscal Coordination Law (LCF), FISM funds are direct federal transfers reserved exclusively for infrastructure projects that benefit localities containing a significant population living in poverty. Eligible projects include investments in the water supply, drainage, electrification, health infrastructure, education infrastructure, housing, and roads. However, voters are poorly informed about both the resources available to mayors and the mayor's responsibility to provide basic public services.³

The use of FISM transfers is subject to independent audits. Responding to high levels of perceived mismanagement of public resources, the Federal Auditor's Office (ASF) was established in 1999 to audit the use of federal funds. Although the ASF reports to Congress, its management autonomy is constitutionally enshrined and it has the power to impose fines, recommend economic sanctions, and file or recommend criminal

² See Wellenstein, Núñez, and Andrés (2006).

³ Chong et al. (2015).

cases against public officials.⁴ The ASF selects around 150 municipalities for audit each year, based primarily on the relative contribution of FISM transfers to the municipal budget, historical performance and factors that raise the likelihood of mismanagement, and whether the municipality has recently been audited (including concurrent federal audits of other programs).⁵ Importantly, the municipalities to be audited in a given year are announced after spending has occurred.

Audits address the spending, accounting, and management of FISM funds from the previous fiscal year. And although the ASF's reports categorize the use of FISM funds in various ways, we focus on two key dimensions of mayoral malfeasance documented in the audit reports (that are not necessarily mutually exclusive): the share of funds spent on projects not directly benefiting the poor, and the share of funds spent on unauthorized projects. Spending not benefiting the poor entails the allocation of FISM funds to social infrastructure projects that do not benefit impoverished localities. Unauthorized spending primarily includes the diversion of resources for nonsocial infrastructure projects (e.g., personal expenses and election campaigns) and funds that are not accounted for. Such spending is akin to the corruption identified by Brazilian audit reports.⁶ The results for each audited municipality are reported to Congress in February the year after the audit was conducted. Thus, just as in the Metaketa studies in Benin and Uganda 2, the information reported in our leaflets comes from official government sources. However, the information we distribute is more publicly available since the ASF's reports are made available on its website, www.asf.gob.mx, and traditional media outlets often make reference to them.⁷

According to these metrics, mayoral malfeasance is relatively high. Between 2007 and 2015, 8 percent of audited funds were spent on projects not benefiting the poor, while a further 6 percent were spent on unauthorized projects. Given that the ASF captures only one dimension of malfeasance, it is thus unsurprising that 45 percent of voters do not believe that municipal governments use public resources honestly.⁸

Although economic and criminal punishments for misallocating funds are relatively rare, there are good reasons to believe that voters will hold

⁴ See Larreguy, Marshall, and Snyder (2016).

⁵ Auditoría Superior de la Federación (2014).

⁶ Ferraz and Finan (2008).

⁷ Larreguy, Marshall, and Snyder (2016).

⁸ Chong et al. (2015).

the incumbent party responsible in Mexico's party-centric electoral context, even without the possibility of reelecting individual mayors. First, voters are considerably better informed about political parties than about individual politicians.⁹ Crucially for political accountability, 80 percent 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 correlated attributes.¹⁰ For example, 74 percent of voters in the survey we conducted as part of this study believe that if the current mayor is malfeasant, another candidate from within the same party is at least somewhat likely to also be malfeasant. Moreover, the same survey shows that 74 percent and 72 percent of respondents in control precincts respectively regard fighting poverty and honesty as important or very important in deciding which candidate to vote for.

However, existing evidence of electoral sanctions against Mexico's incumbent parties in response to information on malfeasant behavior is mixed. Larreguy, Marshall, and Snyder (2016) observe large electoral penalties among urban voters with access to broadcast media outlets incentivized to report local news. Exploiting plausibly exogenous variation in the release of audit reports prior to elections and access to radio and television stations in urban electoral precincts across the country, they find that each additional local media station decreases the vote share of an incumbent party revealed not to have spent FISM funds on the poor or in an unauthorized manner before the election by around 1 percentage point.¹¹ This evidence supports the standard electoral accountability model.¹²

Conversely, in a field experiment conducted in twelve municipalities across three states, Chong et al. (2015) find evidence that providing malfeasant information breeds disengagement. That is, while incumbent support declines when the incumbent is revealed as highly malfeasant, challenger support declines at least as much. Further, they find reduced partisan attachment to the incumbent, suggesting an equilibrium where voters disengage because they believe that all politicians are malfeasant.

⁹ Chong et al. (2015).

¹⁰ Langston (2003).

¹¹ These results confirm the findings of Ferraz and Finan (2008) in Brazil, who found that the presence of local media is critical for the release of audit reports to affect electoral outcomes.

¹² Barro (1973), Fearon (1999), Ferejohn (1986), Rogoff (1990).

The disjuncture between these accountability and disengagement findings points to the importance of developing theories and interventions to better understand the role of information.

In sum, there are several reasons why Mexico constitutes a valuable case study for the Metaketa. After over seventy years of single-party rule, Mexico has become an increasingly vibrant multiparty democracy since the mid 1990s. Thus, just like in Benin, Brazil, Burkina Faso, and Uganda,¹³ we study the effect of information in the context of a new democracy where voters may be particularly uninformed and have relatively less experience with keeping politicians accountable via elections. Nonetheless, in spite of considerable progress (e.g., the creation of an independent federal electoral agency, the reduction of vote fraud, and the creation of a more level playing field for different parties), politics in Mexico – just like in the other Metaketa countries – remains highly clientelistic. The clientelistic nature of political transactions heavily undermines the capacity of voters to keep politicians accountable. Just like in the Uganda 2, Burkina Faso, and Brazil studies in this volume, we provide information on the performance of local politicians, as opposed to national legislators as in the Benin and Uganda 1 studies in this volume. Accountability may operate very differently for local politicians who have executive responsibilities and whose actions are potentially more observable to voters, than for national politicians. Finally, and unlike the other Metaketa studies where the politicians under consideration are allowed to be reelected, we study the accountability of incumbent parties rather than of incumbent individuals, given that Mexico did not allow the reelection of incumbent mayors.

5.2 EXPERIMENTAL DESIGN AND DATA

5.2.1 Experimental Design

We designed a field experiment to test the role of information on electoral accountability, focusing on Mexico's municipal elections held on Sunday, June 7, 2015.¹⁴ Specifically, we seek to identify the effect of providing voters in 678 electoral precincts with the results of audit reports documenting the municipal use of federal transfers designated for

¹³ See the Metaketa studies in this volume.

¹⁴ 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.




FIGURE 5.1 Mexico: Timeline of the experiment's implementation

infrastructure projects benefiting the poor. We first explain our sample selection, and then describe our information interventions, randomization, and estimation strategy. Figure 5.1 presents the overall timeline of the experiment's implementation.

Sample Selection

Our study focuses on twenty-six municipalities in the central states of Guanajuato (seven municipalities), México (fourteen municipalities), San Luis Potosí (four municipalities), and Querétaro (one municipality), as shown in Figure 5.2.¹⁵ Beyond holding elections in 2015, these four states were chosen for security and logistical reasons and because they contain internal variation in the municipal incumbent party and its performance on the ASF audits.

The twenty-six municipalities were selected from those where an audit was released in 2015 according to three criteria. First, we prioritized the safety of voters and our distribution and survey teams. Nevertheless, just like in the studies in Benin, India, and Uganda 1 (this volume), our field teams faced intimidation and interference in some areas. Immediately after receiving threats upon entering Aquismón and Villa Victoria, the precincts in these municipalities were replaced by Atlacomulco, Temoaya, and additional precincts from Tlalnepantla de Baz in the state of México. Second, we only selected municipalities where at least one of our two measures of reported malfeasance (percentage of misallocated FISM resources) was at least two percentage points lower or, more often, higher than the state average among audited opposition parties. Finally, municipalities were chosen to match the distribution of incumbent parties across audited municipal governments in these four states. Specifically, of our twenty-six municipalities, seventeen were governed by

¹⁵ The average municipality contains 259,000 registered voters.

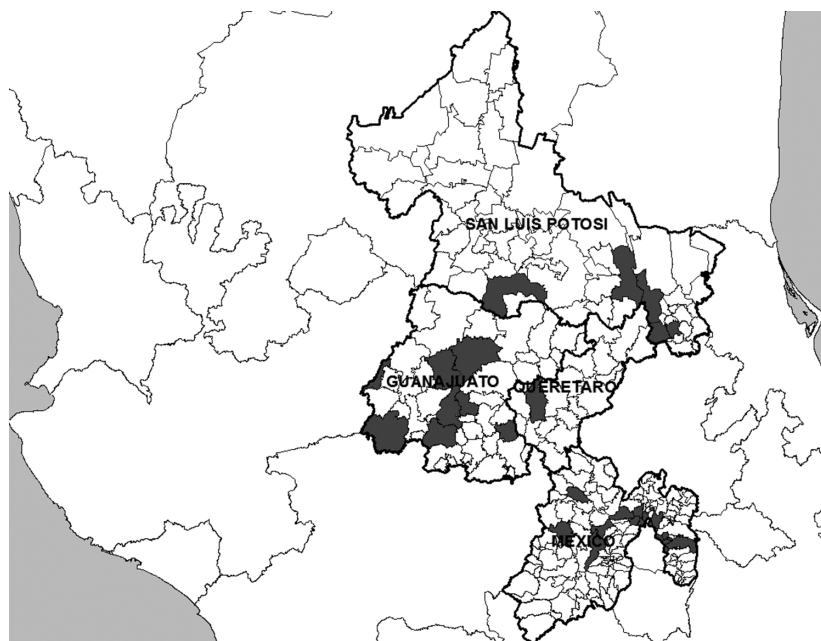


FIGURE 5.2 Mexico: The twenty-six municipalities in our sample

the PRI (including sixteen in coalition with the Teacher's (PANAL) and Green (PVEM) parties), five by the PAN (including two in coalition with PANAL), two by the PRD, and one by the Citizen's Movement (MC).

In order to ensure that our intervention would not affect aggregate electoral outcomes, we selected only up to one-third of the electoral precincts within each municipality.¹⁶ To ensure variation in the information reported in our leaflets, we oversampled precincts from municipalities with particularly high or low levels of incumbent malfeasance, and thus with strong contrasts with respect to opposition parties' malfeasance within the state. Within municipalities, we first prioritized accessible rural precincts, where possible, in order to minimize cross-precinct spillovers and maximize the probability that voters would not

¹⁶ Given our focus on rural precincts, effectively this means that fewer than one-third of voters in the municipality were subject to our treatment. In addition, as discussed above, in order to guarantee that our intervention would not unfairly affect any specific party, the distribution of incumbent parties in our sample closely matches the distribution of incumbent parties in the four states.

receive the information through other means.¹⁷ 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 areas, where we had more precincts to choose from, we restricted our sample to precincts with at most 1,750 registered voters, and designed an algorithm to minimize the number of neighboring urban precincts in our sample.¹⁸ This implies that our estimates are particularly informative for voters in relatively rural and sparsely populated areas.

Information Treatment

In partnership with Borde Político, we distributed to voters leaflets documenting the use of FISM funds in their municipality. For each municipality, the leaflet focused on either the proportion of spending that does not benefit the poor, or unauthorized spending. All treatments were delivered at the electoral precinct level, Mexico's lowest level of electoral aggregation.

Our leaflet was designed to be nonpartisan, accessible, and sufficiently intriguing so that voters would not immediately discard it. The particular design was produced by a local graphic designer based on feedback from multiple focus groups.¹⁹ We also sought legal advice to ensure that our leaflets did not constitute political advertisements, and thus were not subject to distribution restrictions stipulated in Mexican electoral law. To avoid suspicions of political motivation, neither the incumbent mayor nor its party are referred to directly, although as noted above the vast majority of voters can correctly identify the party of their incumbent mayor.

One of the variants of our intervention involved whether the leaflet included information only on the local incumbent party (local treatment) or whether it also reported the average outcome among all audited municipalities within the same state but governed by a different political party (benchmarked treatment). Figure 5.3 provides an example of a

¹⁷ Municipalities in Mexico are divided into multiple *localidades* (the equivalent of US Census blocks). Most municipalities consist of both urban and rural *localidades*.

¹⁸ The algorithm started with the set of neighboring precincts surrounding each precinct and identified all neighboring precincts that were eligible for our sample; we then iteratively removed the precinct with most "in-sample" neighbors until we reached the required number of precincts for that municipality. In most municipalities, the algorithm ensured that our sample contained no neighboring precincts.

¹⁹ To avoid any association with any specific political party, the leaflet design was in black and white. We used images to communicate in a simple and friendly way how FISM funds should be spent according to the law.



FIGURE 5.3 Mexico: Example of benchmarked leaflet in Ecatepec de Morelos, México

benchmarked 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 nonpartisan organization and that the information contained in the leaflet is based on the ASF's official audit reports 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 spent in an unauthorized way by their government (45 percent). The leaflet also shows the average percentage of unauthorized spending in other municipalities in the state governed by a different party (9 percent).

Block Randomization and Implementation

Our sample of 678 precincts was allocated to treatment conditions according to a factorial design with a pure control, as shown in Table 5.1. The 400 treated precincts were equally divided between the four variants of the information treatment.

Since all of the other Metaketa studies provided benchmarked information, the latter constitutes our common intervention arm and thus, throughout the rest of this chapter, we focus exclusively on the effects of

TABLE 5.1 *Mexico: Factorial design with a pure control*

	Control	Private	Public
Control	278 precincts		
Local		100 precincts	100 precincts
Benchmarked		100 precincts	100 precincts

the benchmarked treatment. For power considerations, and also because we find limited differences between the private and public treatment, in our baseline analysis we pool the private and public benchmarked treatments and use this as our common arm (in Section 5.4 we report separate estimates for the private/public variants). Nonetheless, in the online appendix we also report estimates using only the benchmarked-private treatment as our common arm. As such, the 200 precincts assigned to the local leaflet are dropped from our analysis and thus our effective sample covers 478 precincts. The control group comprising 278 electoral precincts reflects our sampling and block randomization design.

For the randomization, precincts were first stratified into rural or urban blocks of six or seven similar precincts within a given municipality.²⁰ Precinct similarity was defined by the Mahalanobis distance between twenty-three social, economic, demographic, and political variables provided by Mexico's National Statistical Agency (INEGI) and the National Electoral Institute (INE).²¹ 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. Our block randomization ensures that different municipalities do not receive different treatment proportions and maximizes the power of the experiment by minimizing differences between treated and control precincts within blocks.

The leaflets were distributed by our implementing partners Data OPM and Qué Funciona para el Desarrollo using precinct maps provided by state electoral institutes. Our distribution teams delivered one leaflet to a

²⁰ Subject to there existing sufficient precincts, and the total treated precincts not exceeding one-third of all precincts, we used blocks of seven precincts.

²¹ We used the R package `blockTools` to assign precincts to blocks. The algorithm is "greedy" in that it creates the most similar group first. Where a surplus of potential precincts were available, we used the most similar blocks to maximize statistical efficiency.

maximum of 200 households in the largest locality (in rural blocks) and in randomly selected city blocks (in urban blocks) within each treated precinct. Within our sample, the median precinct contained 353 households (according to the 2010 Census), 420 private dwellings, and 1,056 voters registered for the 2015 election. Where possible, leaflets were delivered in person with a short message explaining the leaflet's provenance. When no adult was available, leaflets were left in mail boxes or taped to the recipient's front door in a waterproof bag. Leaflet delivery took several hours per precinct, and was implemented over three weeks concluding at the legally designated end of the campaign four days before the election. Our team recorded where leaflets were distributed in order to return to these same places for our follow-up survey.

While compliance with the delivery of our treatments was very good in general, we nevertheless encountered some issues in the field. In a couple of cases, some leaflets were delivered to voters outside the precinct or adverse weather and poor road conditions prevented us from reaching a precinct. To preserve the randomization, we focus on estimating intention to treat (ITT) effects and thus code our treatment dummy according to the original assignment.

Our experimental design, sample, and replacements are illustrated in CONSORT-style Figure 5.4.

5.2.2 Data

Throughout this section we make specific reference to the notation used for different measures in the meta-analysis pre-analysis plan (MPAP).

We collected two sources of data to measure our main outcomes. First, we use official precinct-level electoral returns to measure incumbent party vote share (as a fraction of those who voted) and turnout.²² These correspond to measures M2 and M4, respectively, in the MPAP.

Second, we conducted a postelection survey interviewing 10 voters from each of the treated precincts and 100 randomly selected control precincts.²³ Despite its success in the administration of the Mexico Panel

²² We had to drop the three precincts in our sample that were merged with another precinct, due to having fewer than 100 registered voters.

²³ For details on the field protocol used for selecting survey respondents see Arias et al. (n.d.).

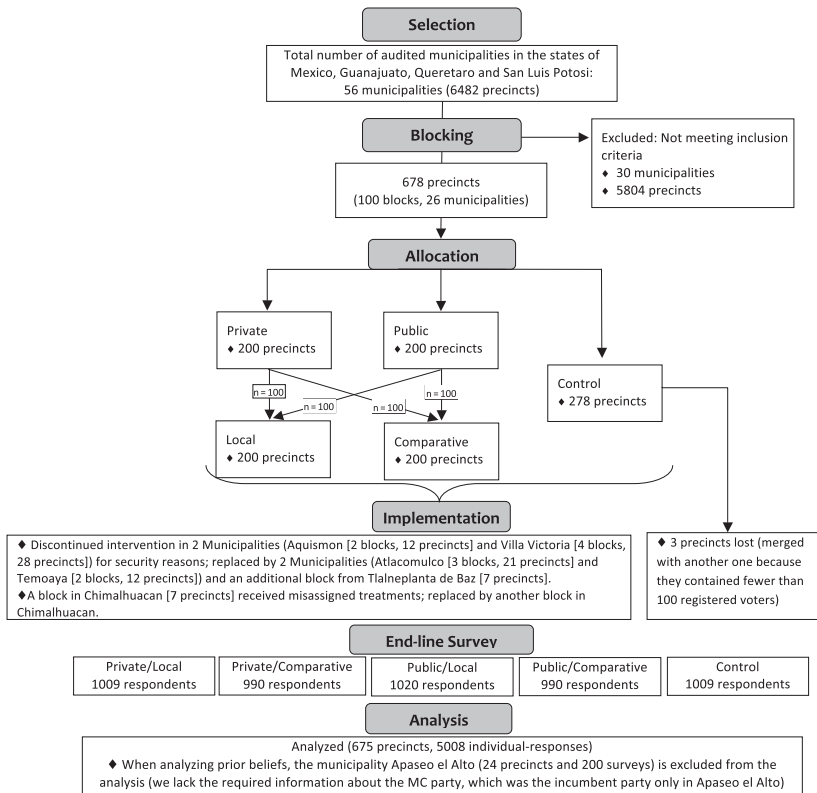


FIGURE 5.4 Mexico: CONSORT diagram of research design

Surveys and Comparative Study of Electoral Systems modules and in the Metaketa study in Burkina Faso, our attempts to gauge individual vote choice by simulating the electoral process with an urn during the survey had little success. Many voters felt uneasy in the aftermath of surprising electoral results, and refused to participate believing that our survey team were working on behalf of a party to identify individual vote choices or to trick voters into casting a different ballot. Furthermore, enumerators reported little confidence that respondents had truly reported the party they voted for. Thus, throughout the baseline analysis we rely only on official precinct-level measures of vote share and turnout, but report results with the self-reported individual-level data in the online appendix.

To measure respondents' evaluations of the incumbent and challenger parties' malfeasance (MPAP measure M6), at the beginning of the survey we asked respondents to rate, on a five-point scale from very low (−2) to

very high (2), the level of corruption or the level of interest in supporting the poor of each major party (depending on the measure of malfeasance provided in that municipality). This provides us with a measure of voters' posterior beliefs/evaluations (Q_{ipbm}) where higher values correspond to higher perceptions of malfeasance.²⁴

In order to test the hypotheses preregistered in the MPAP, we need to assess how the information provided in the leaflets compares with voters' prior beliefs. This is necessary in order to establish whether the information provided was interpreted as "good news" or "bad news" by voters. However, this is particularly challenging in our context since unlike the other Metaketa studies, due to financial constraints, we were unable to conduct a baseline survey. To address this issue, we use the (postelection) surveys conducted in each block's control precincts to proxy for the average pretreatment beliefs of the treated and control voters within the same block.²⁵ This provides us with a block-level measure of voters' priors about the malfeasance of the incumbent and challenger parties, P_{pbm} (MPAP measure M9).

Using postelection surveys among the control group to proxy for pretreatment beliefs requires several assumptions. First, control group respondents are similar to treatment group respondents. Second, control group respondent beliefs are consistent across the month between the intervention and the postelection survey. In Arias et al. (n.d.), we show that these assumptions are plausible in the context of this study. First, randomization ensures that treated and control precincts are identical in expectation and relatively similar in realization due to our blocking strategy. Second, we show that municipal-level electoral outcomes do not systematically affect the level or strength of beliefs about incumbent malfeasance among respondents in the control group (indicating that the election outcome itself does not significantly influence voter beliefs). Third, the 2012 Mexican Panel Survey shows that voter assessments of

²⁴ In the MPAP, Q makes reference to the information provided in the leaflet, in our case, the percentage of malfeasant spending. However, we did not ask voters about their expectations of this continuous measure (fraction of resources misspent) as this would have been very hard for our respondents to assess. Thus we define Q as voters' posterior evaluations of each party (after having seen the information in the leaflet) such that we can compare it to our measure of priors on the same five-point scale.

²⁵ We prespecified that prior beliefs would be defined by control voters at the municipal level. However, we focus on the block-level controls to produce more precise measures. Our results are similar if we define priors using control respondents within municipality (rather than within block).

politicians are relatively persistent.²⁶ Fourth, we show that there is no compelling evidence of spillovers to neighboring precincts outside our sample. Finally, we also rely on data from the Brazil study in order to test the extent to which our approach of measuring voter priors based on endline responses of control respondents is valid. In the Brazil study, both baseline and endline surveys were conducted, making it possible to directly test the block-level (average) correlation between endline evaluations of control respondents and baseline evaluations of treated respondents. The correlation is close to 0.7, which suggests our approach is indeed valid.

Next, we construct a measure of the gap between the voters' priors (P_{pbm}), and their evaluations (posteriors) after observing the information reported in the leaflet (Q_{pbm}). This allows us to establish whether the treatment caused voters to update favorably or unfavorably about the malfeasance of the different parties. Specifically, at the end of the survey we showed voters the leaflet corresponding to their municipality and again asked them how they perceived the incumbent/challenger party on the same five-point scale. To measure how the information provided compared to their priors, we simply consider the average change in perceptions before and after showing the corresponding leaflets to voters within the block's control precincts. This allows us to compute a block-level measure of the gap between the voters' posterior and prior evaluations of incumbent/challenger malfeasance, namely:

$$G_{pbm} = Q_{pbm} - P_{pbm}$$

Positive (negative) values of this measure imply that the party was more (less) malfeasant than what voters anticipated. We use this continuous measure to generate block-level indicators for whether voters interpreted the information in our treatment as good news or bad news:

*Good News*_{pbm} = 1 if $G_{pbm} < 0$ and zero otherwise.

*Bad News*_{pbm} = 1 if $G_{pbm} \geq 0$ and zero otherwise.

The MPAP hypothesized how good or bad news about the incumbent's party affect the latter's evaluations and vote share. Thus, in our baseline analysis we report the effect of good and bad news (and the gap

²⁶ Voter opinions of the presidential candidates before and after the election – three months apart, in contrast to the three to four weeks apart we examine – exhibit a 0.4 correlation.

between posteriors and priors) for the incumbent party. However, given that our treatment is benchmarked and provides information on both the incumbent and the challenger parties, a natural exercise is to explore how relative good/bad news affect voters' evaluations of the incumbent relative to the challenger (relative posteriors).²⁷ While we did not preregister hypotheses about relative updating, we present some of these results in the online appendix.

Finally, we construct several measures to analyze the extent to which our information treatment has differential effects depending on context-specific factors. As before, we rely on the postelection survey responses from control precincts in each surveyed block to construct an average pretreatment proxy for our variables of interest. In all cases, for ease of interpretation we standardize all these measures to have a mean of 0 and a standard deviation of 1.

First, we examine the moderating role of partisan identity (MPAP measure M19). Voters with stronger partisan preferences or attachment may be less likely to update their beliefs or voting behavior in response to new information. Here, we define "Incumbent Partisanship" as the block-level fraction of control respondents who identified themselves as sympathizers of the incumbent party.²⁸ In a similar spirit, we explore the competitiveness of the electoral context (MPAP measure M25). In competitive political environments, newly revealed information on the malfeasance of incumbent and challenger parties is more likely to be pivotal in voters' electoral decision. For this, we create an "Electoral Competitiveness" measure at the precinct level. Using results from the 2012 election, we define competitiveness as 1 minus the margin of victory of the incumbent, where the margin of victory is defined as the vote share of the election winner minus the vote share of the runner up. We also examine the extent to which our treatment effects vary by the

²⁷ More concretely, let G_{pbm}^I measure the gap between posteriors and priors about the incumbent and G_{pbm}^C the corresponding measure for the challenger. This allows us to define the measure $Relative\ Good\ News_{pbm} = 1$ if $(G_{pbm}^I - G_{pbm}^C) < 0$ and $Relative\ Bad\ News_{pbm} = 1$ if $(G_{pbm}^I - G_{pbm}^C) > 0$. We can also use as outcome variable a measure of relative evaluation (posterior) of the incumbent's malfeasance defined as $Q_{pbm}^I - Q_{pbm}^C$.

²⁸ More concretely, the survey question asked: "In general, with which political party do you sympathize?" and coded responses as 1 if the party coincided with the incumbent, and 0 otherwise.

prevalence of clientelistic practices within the community (MPAP measure M22). Clientelism fundamentally undermines the ability of voters to keep politicians accountable²⁹ and thus informational interventions may be less effective in environments in which the voter-politician relation is mediated by contingent exchange. To measure this, we created a proxy for vote buying as follows: we asked respondents to assess how frequently the municipal incumbent party offered gifts, services, and other favors in exchange for votes, with responses ranging from ‘not frequently at all’ (1) to ‘very frequently’ (4), and again estimated the block-level average from control respondents.

Second, we also construct several measures of the informational environment. Our intervention may be particularly effective in places where voters have limited access to other sources of information and thus where our treatment is more likely to provide new information the voters were not already aware of. We create two indexes. First, we construct a political knowledge index. To do so, our survey included three questions evaluating factual questions about local politics, namely (i) how long is the term of the municipal mayor, (ii) which is the party of the outgoing mayor, and (iii) which party won the last municipal elections. We code correct answers with a value of 1 (0 otherwise) and then aggregate the three questions for an individual-level measure ranging between 0 and 3. We then take the average of this measure across control respondents to obtain a block-level proxy of political knowledge. Second, we also create a media consumption index. For this, we asked respondents how often they followed electoral news over TV, radio, newspapers, and internet and social media, respectively, with possible responses ranging from “never” to “daily” (5). We take the mean of these four responses to create an individual-level measure, and once again, take the block-level average across respondents in control precincts.

Finally, related to both the electoral and societal context, we also examine the extent to which our information treatment varies with the level of trust in elections (MPAP measures M26 and M27). If voters are disillusioned and do not believe that elections are an effective mechanism to keep politicians accountable, they may be less likely to change their behavior in response to information on malfeasance. For this, our survey retrieved people’s opinions on how clean they perceived the election to

²⁹ Stokes (2005).

be. More specifically, we asked people to assess how likely it was that the vote count during the June 7 election was free and fair, with possible responses ranging from “not at all likely” (1) to “extremely likely” (5). That is, higher values represent higher beliefs that the electoral process is clean. Again, we take the mean of control respondents to create our block-level proxy.

5.3 ANALYSIS OF TREATMENT EFFECTS ON THE COMMON TREATMENT ARM

We present the main treatment effects of our intervention, focusing on the common treatment arm. We focus on providing estimates for the different hypotheses registered in the MPAP. First, we provide estimates for the effect of our intervention on voters’ posterior beliefs on the integrity of candidates from the incumbent party (H₃). Next, we report estimates on vote share and turnout (H₁ and H₂). In all tables, we first report estimates of the average effect of our treatment (i.e., the effect of providing a voter with a leaflet) when we pool all randomization blocks (panel A). We then report separate estimates for blocks in which our information was interpreted by voters as bad news (panel B) and good news (panel C). We explore the effect of different variants of our treatment such as mode of delivery (private vs. public) and the type of audit information reported in the leaflet (unauthorized spending vs. spending on the poor) and finally test for heterogeneous effects on several context-specific factors.

We estimate OLS regressions of the form:

$$Y_{pbm} = \beta T_{pbm} + \eta_{bm} + \epsilon_{pbm}, \quad (5.1)$$

where Y_{pbm} is an outcome variable in precinct p , in randomization block b , in municipality m (individual level outcome variables have an additional subindex i) and T_{pbm} is an indicator for whether precinct p was assigned to treatment.

To estimate differential effects for variants of our intervention or context-specific factors we estimate regressions of the form:

$$Y_{pbm} = \beta T_{pbm} * \mathbf{Z}_{bm} + \eta_{bm} + \epsilon_{pbm}, \quad (5.2)$$

where \mathbf{Z}_{bm} contains different measures on the mode of delivery, type of information of the leaflet and block characteristics.

All our regressions include randomization block fixed effects (η_{bm}) and standard errors are clustered at the municipality-treatment level.³⁰ Precinct-level observations are weighted by the share of voters to whom we delivered a leaflet (in control precincts, we use the number of leaflets delivered to the average treated precinct). This weighting scheme was not preregistered but permits more precise estimates by de-weighting large precincts where only a small fraction of voters could receive the leaflet. In the online appendix we report similar estimates from unweighted regressions.

5.3.1 Validity Checks

Before reporting the estimates for our main outcome variables, we provide evidence to validate our randomization. In Online Appendix Table B1, we use the basic specification in Equation 5.1 to show that treatment is well-balanced across thirty-seven precinct and survey respondent characteristics, with only a handful of statistically significant differences. However, all of our estimates are robust to controlling for precinct level characteristics (results not reported).

We also conducted some basic manipulation checks reported in Arias et al. (n.d.). Respondents in treated precincts are substantially more likely to report having received a leaflet, remember the content of the leaflet, and say that the leaflet influenced their voting decision. Moreover, respondents assigned to the benchmarked treatment (the focus of this chapter) are more likely to report that information on the opposition parties was also included in the leaflet.

5.3.2 Effects on Beliefs about Party Integrity

First, we examine how our treatment impacted voters' evaluations of the incumbent party's level of malfeasance (Q_{ipbm}). Recall that this variable is measured on a five-point scale and higher values imply higher perceptions of incumbent malfeasance.

The coefficient in column 1 of Table 5.2 suggests that, on average, our intervention improved voters' perceptions about candidate integrity, though it is not statistically significant at conventional levels. However,

³⁰ We cluster at the municipality \times treatment rather than the block level since the information reported in the leaflets (i.e., incumbent level of malfeasance) is the same across blocks within a municipality.

TABLE 5.2 *Mexico: Effect of information treatment on voter posterior beliefs about incumbent party malfeasance*

	Perceived incumbent party malfeasance (very low – very high)			
	(1)	(2)	(3)	(4)
Panel A: Pooled				
Information treatment	-0.030 (0.040)	-0.193*** (0.060)		
Information treatment × Gap between posteriors and priors		0.183*** (0.037)		
Private information treatment			-0.048 (0.048)	
Public information treatment			-0.012 (0.050)	
Unauthorized spending information treatment				-0.039 (0.043)
Not-spending on the poor information treatment				-0.024 (0.060)
Observations	2816	2816	2816	2816
R ²	0.30	0.30	0.30	0.30
Panel B: Bad News				
Information treatment	0.035 (0.048)	-0.285*** (0.092)		
Information treatment × Gap between posteriors and priors		0.237*** (0.050)		
Private information treatment			0.001 (0.055)	
Public information treatment			0.070 (0.057)	
Unauthorized spending information treatment				0.048 (0.070)
Not-spending on the poor information treatment				0.031 (0.059)
Observations	2048	2048	2048	2048
R ²	0.24	0.24	0.24	0.24

TABLE 5.2 (continued)

	Perceived incumbent party malfeasance (very low – very high)			
	(1)	(2)	(3)	(4)
Panel C: Good News				
Information treatment	−0.202** (0.081)	−0.203 (0.127)		
Information treatment × Gap between posteriors and priors		−0.006 (0.211)		
Private information treatment			−0.179* (0.094)	
Public information treatment			−0.225** (0.100)	
Unauthorized spending information treatment				−0.099 (0.071)
Not-spending on the poor information treatment				−1.025*** (0.157)
Observations	768	768	768	768
R ²	0.12	0.12	0.12	0.13

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 $p < 0.01$.

this average masks substantial variation since the information provided may have represented good news for some voters and bad news for others, relative to their prior beliefs. Consistent with H₃, the estimates in column 1 of Panels B and C suggest that good news leads voters to favorably update their beliefs about the incumbent's party integrity, while bad news leads voters to update unfavorably (though the latter coefficient is not statistically significant).

A complementary approach is to explore how our intervention varies with the continuous measure of the gap between voters' prior beliefs and the information reported in the leaflets, G_{pbm} . The coefficient for the interaction term in column 2 of Panel A is positive, which suggests that voters' perceptions of the incumbent party get worse when the information in the leaflet is more unfavorable relative to what voters expected. This is consistent with the different average treatment effects

found in the good and bad news samples and provides evidence in support of H12. The estimates in column 2 of Panels B and C, show that once we break down the samples by bad and good news, the gap between posteriors and priors has a statistically significant effect only in the former case. However, the MPAP did not hypothesize about different marginal effects of this or other variables across the bad and good news samples.

Importantly, notice that the different treatment effects across the bad and good news sample as well as the positive interaction with the gap between posteriors and priors, while intuitive, are not mechanical since our block-level measures of priors and posteriors are based on responses from voters in control precincts within the block.

5.3.3 Effect on Incumbent Party's Vote Share and Turnout

Next we look at treatment effects on precinct-level incumbent vote share and turnout. The estimate in column 1 of Table 5.3 is positive, which suggests that on average, our treatment increased the incumbent party's vote share by 2.9 percentage points. This is consistent with the estimate in column 1 of Table 5.2 that, while not statistically significant, suggests that our intervention led voters to update favorably about the incumbent party on average. In our pre-analysis plan we conjectured that our intervention would have, on average, a negative effect on the incumbent's party. However, this was under the assumption that the relatively high levels of malfeasance on average reported in the leaflets would cause the average voter to update unfavorably about the incumbent party.

In Panels B and C we once again break down the effect by whether our treatment was interpreted as bad or good news by voters in the randomization block. Consistent with the estimates in Table 5.2, the electoral reward for the incumbent party is almost 50 percent higher in blocks with good news relative to blocks with bad news. In column 2 we interact our treatment indicator with the continuous measure of the gap between posteriors and priors. The coefficient on the interaction term is negative (consistent with the idea that incumbents are less likely to be rewarded the worse posteriors are relative to priors) though not close to being statistically significant.

The positive and statistically significant coefficient for blocks that interpreted our treatment as bad news for the incumbent may appear somewhat surprising. A first important issue is that the average effect

TABLE 5.3 *Mexico: Effect of information treatment on incumbent party vote share and turnout (weighted estimates)*

	Incumbent party vote share				Turnout
	(1)	(2)	(3)	(4)	(5)
Panel A: Pooled					
Information treatment	0.029*** (0.007)	0.027*** (0.007)			-0.002 (0.005)
Information treatment × Gap between posteriors and priors		-0.003 (0.005)			
Private information treatment			0.029** (0.013)		
Public information treatment			0.028*** (0.009)		
Unauthorized spending information treatment				0.025*** (0.008)	
Not-spending on the poor information treatment				0.031*** (0.011)	
Observations	475	459	475	475	475
R ²	0.67	0.67	0.67	0.67	0.72
Panel B: Bad News					
Information treatment	0.021** (0.009)	0.016 (0.018)			-0.004 (0.007)
Information treatment × Gap between posteriors and priors		0.004 (0.009)			
Private information treatment			0.019 (0.016)		
Public information treatment			0.022* (0.013)		
Unauthorized spending information treatment				0.006 (0.013)	

TABLE 5.3 (continued)

	Incumbent party vote share				Turnout
	(1)	(2)	(3)	(4)	(5)
Not-spending on the poor information treatment				0.026** (0.011)	
Observations	332	332	332	332	332
R ²	0.59	0.59	0.59	0.59	0.75
Panel C: Good News					
Information treatment	0.034*** (0.008)	0.019 (0.015)			-0.005 (0.008)
Information treatment × Gap between posteriors and priors		-0.044 (0.034)			
Private information treatment			0.027** (0.012)		
Public information treatment			0.040*** (0.008)		
Unauthorized spending information treatment				0.038*** (0.008)	
Not-spending on the poor information treatment				-0.006 (0.012)	
Observations	127	127	127	127	127
R ²	0.79	0.80	0.79	0.80	0.68

Notes: All specifications include block fixed effects, weighted by the share of the precinct that was treated, 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 lack of prior data in Apaseo el Alto. Standard errors clustered by municipality × treatment are in parentheses. * denotes $p < 0.1$, ** denotes $p < 0.05$, *** denotes $p < 0.01$.

for bad and good news is harder to interpret since it may depend on the specific cutoff used to define these categories. It is likely that whenever posteriors and priors are very close to each other (i.e., G_{pbm} very close to zero), voters interpret this as “neutral news” that doesn’t change their

voting decision. Thus, what is particularly important is the difference (or slope) in the effect between the estimate for good and bad news, which suggests a more positive effect for the former, consistent with a basic model in which voters respond to how they update based on the information received.

In addition, there are several reasons as to why the provision of information may have a positive effect on the incumbent party vote share irrespective of the direction in which voters update. As we show in Arias et al. (n.d.), a basic model of risk-averse voter learning suggests that the provision of information reduces voter uncertainty about the incumbent party's malfeasance, which can in turn increase support for the incumbent. In addition, as we document below, both incumbent and challengers respond strategically to our intervention, which could offset negative effects of unfavorable information on the incumbent's vote share if the incumbent's response were larger or more effective.

Finally, in column 5 of Table 5.3 we report the estimates for the average effect of our intervention on turnout. The coefficients in all panels are very small and statistically insignificant. In this respect, our results do not support MPAP's H2a and H2b that predicted that good news would increase turnout while bad news would decrease it. However, in Arias et al. (n.d.) we present a formal model that predicts that the effect of new information on turnout should be nonlinear. In particular, when voter partisan attachments are bimodally distributed and voters at each mode turn out for different parties,³¹ shockingly favorable or unfavorable revelations lead voters at one mode to switch parties and thereby increase turnout, while relatively unsurprising but nevertheless informative favorable (unfavorable) information induces challenger (incumbent) partisans to become relatively indifferent between the parties and instead abstain when faced with a cost of turning out. This logic does not yield clear predictions for the average effect of new information. This specific prediction was preregistered in the preanalysis plan of our study, and the results in Arias et al. (n.d.) support it.

5.3.4 Robustness Checks

The patterns reported in Tables 5.2 and 5.3 are broadly robust to various changes in our empirical specification. In Online Appendix

³¹ Furthermore, similar results obtain for unimodal distributions when the modal voter initially turns out.

Table B2, we report unweighted estimates for our vote share regressions (recall that the weighting scheme was not registered in our pre-analysis plan). The main patterns in the data remain relatively unchanged though the treatment effect in the good news sample becomes slightly smaller.

The similar treatment effects across the private and public variants of our treatment documented in Section 5.4.1 suggests that our decision to pool both variants of our benchmarked treatment as the common arm is reasonable. Nevertheless, as robustness checks, in Online Appendix Tables B3 and B4, we report our baseline estimates on posteriors and vote share where we use the benchmarked-private treatment as our common arm and compare it against control precincts (i.e., we drop precincts with the benchmarked-public treatment from our sample). While this further decreases the sample size in all regressions, the estimates are similar in magnitude to those in our baseline specification. In fact, the interaction between the treatment and the gap between priors and posteriors is now negative and statistically significant in Panels A and B of the vote share regressions in Online Appendix Table B4, suggesting that the electoral reward for incumbents is smaller in precincts where the information reported in the leaflets is worse than what voters expected.

In Online Appendix Table B5, we report estimates equivalent to those reported in Table 5.3 but where we use individual-level self-reported vote and turnout data (as opposed to official, precinct level aggregates). The results in this case are totally inconsistent with the findings based on official statistics. The average treatment effect in the good news sample (column 1, Panel C) is negative, large and statistically significant suggesting that good news about the incumbent party decreases the likelihood that voters vote for it. Moreover, the interaction of the treatment with the gap between posteriors and priors is positive, on average, and in the bad news sample, suggesting that whenever voters interpret the information as worse than what they expected, they are more likely to vote for the incumbent. However, as discussed in Section 5.2.2, there are several reasons why these self-reported individual-level measures are not reliable (and certainly less reliable than official precinct-level statistics) since many voters refused to report who they voted for, and enumerators expressed little confidence in the veracity of the information provided by those who did fill out the mock ballot.

Finally, in Online Appendix Tables B6 and B7 we report estimates when we focus on relative measures of posteriors and good/bad news.

Results using relative posteriors as the outcome variable in Online Appendix Table B6 show that respondents update more unfavorably about the incumbent, relative to the challenger party, whenever news about the incumbent are worse than about the challenger (Panel B). The coefficient in Panel C for the relative good news sample is negative (as one would expect) though not statistically significant. The interaction with the relative gap between posteriors and priors variable in column 2 is positive and statistically significant in all panels, which is consistent with the idea that voters update more negatively about the incumbent (relative to the challenger) whenever the gap between posteriors and priors is worse (bigger) for the incumbent than for the challenger. On the other hand, the estimates for the vote share regressions in Online Appendix Table B7 show that the effect of the treatment in the relative bad news sample is very close and, if anything, slightly larger, than the effect on the relative good news sample.

5.3.5 Heterogeneous Effects

Finally, we look at heterogeneous effects according to several contextual characteristics specified in MPAP H6–H11.³² As discussed in Section 5.2.2, given the absence of individual-level voting outcomes, we focus on average precinct or block-level characteristics based on control-group respondents. The results are presented in Table 5.4. All block-level characteristics used in the interaction are normalized to have mean zero and standard deviation equal to one. For reference, in column 1 we report the average treatment effects in each sample.

In column 2 we explore the role of partisanship. Consistent with MPAP H7, we find that, in blocks where voters are more likely to share the partisan affiliation of the incumbent, good news has a weaker effect on the incumbent's vote share. The interaction in the bad news sample is very small and statistically insignificant. These results are consistent with partisan voters having strong ideological attachments to the incumbent, and thus being less likely to update their posteriors or change their voting decisions in response to new information.

³² Because ethnic identity is not particularly salient in Mexico's political context, we did not collect any information about it and do not report any heterogeneous effects on this variable.

TABLE 5.4 Mexico: Effect of information treatment on incumbent party vote share, by context (weighted estimates)

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Panel A: Pooled							
Information treatment	0.029*** (0.007)	0.029*** (0.007)	0.028*** (0.007)	0.029*** (0.007)	0.029*** (0.007)	0.029*** (0.007)	0.029*** (0.007)
Information treatment \times Inc. Partisanship		-0.017*** (0.006)					
Information treatment \times Vote Buying			-0.014** (0.007)				
Information treatment \times Competitiveness				-0.014** (0.006)			
Information treatment \times Political Knowledge					0.002 (0.007)		
Information treatment \times Media Consumption						-0.000 (0.008)	
Information treatment \times Trust in Elections							0.000 (0.007)
Observations	475	475	475	475	475	475	475
R ²	0.67	0.68	0.68	0.68	0.67	0.67	0.67

Panel B: Bad News						
Information treatment	0.021** (0.009)	0.021** (0.009)	0.020** (0.009)	0.020** (0.010)	0.021** (0.010)	0.021** (0.010)
Information treatment \times Inc. Partisanship		-0.007 (0.008)				
Information treatment \times Vote Buying				-0.012 (0.008)		
Information treatment \times Competitiveness				-0.018*** (0.007)		
Information treatment \times Political Knowledge					0.004 (0.007)	
Information treatment \times Media Consumption					0.008 (0.010)	
Information treatment \times Trust in Elections						0.003 (0.009)
Observations	332	332	332	332	332	332
R ²	0.59	0.59	0.59	0.59	0.59	0.59
<i>(continued)</i>						

TABLE 5.4 (continued)

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Panel C: Good News							
Information treatment	0.034*** (0.008)	0.038*** (0.007)	0.033*** (0.009)	0.031*** (0.006)	0.033*** (0.008)	0.032*** (0.007)	0.036*** (0.008)
Information treatment \times Inc. Partisanship		-0.027*** (0.008)					
Information treatment \times Vote Buying			0.005 (0.007)				
Information treatment \times Competitiveness				0.016 (0.010)			
Information treatment \times Political Knowledge					-0.022* (0.010)		
Information treatment \times Media Consumption						-0.016* (0.008)	
Information treatment \times Trust in Elections							-0.018* (0.010)
Observations	127	127	127	127	127	127	127
R ²	0.79	0.80	0.79	0.80	0.80	0.80	0.80

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

To study the role of clientelism, in column 3 we look at the interaction of our treatment dummy with the block-level measure of vote buying. The interaction coefficient in Panel A is negative (consistent with MPAP H8) and statistically significant. However, the negative effect seems mostly concentrated in the bad news sample (even though the coefficient is not statistically significant). One may have expected incumbents in clientelistic settings of being more effective at counteracting the dissemination of bad news about them in which case we would have expected a positive (or null) interaction term. That said, hypotheses in the MPAP do not make differential predictions for the bad and good news sample. Moreover, since self-reported questions on vote buying are subject to social desirability bias and other problems, these results should be interpreted cautiously.

In column 4 we look at electoral competitiveness. Consistent with H10, we find that the incumbent party was less likely to be rewarded following the dissemination of bad news in more competitive precincts. A one-standard deviation increase in political competition decreases the effect of the treatment by 1.8 percentage points. In other words, a more competitive environment makes bad news more damaging (or less rewarding) for the incumbent. Conversely, the interaction in Panel C for the good news sample is positive (though not statistically significant), which suggests that good news increases the electoral reward of incumbent parties in more competitive environments. Overall our results do suggest that informational interventions may improve accountability (i.e., lead to the electoral sanction of bad behavior and the reward of good behavior) in electorally competitive environments, where it can play a pivotal role.

Finally, in columns 5–7 we explore the moderating effect of political knowledge, media consumption, and trust in elections. The interactions in Panel B for the bad news sample are very small and statistically insignificant. The interactions in Panel C on the other hand, suggest that when our intervention represented good news, it had a smaller effect in places where voters were more knowledgeable about politics or had higher levels of media consumption. This is consistent with MPAP H9, predicting stronger treatment effects in informationally weak environments. This is intuitive since voters may be less likely to update and change their voting behavior if they were already well informed about the incumbent party's performance (through the media or other means). The interaction for our measure of trust in elections is also negative and statistically significant, which goes against MPAP H11.

5.4 TREATMENT VARIANTS

5.4.1 Private vs. Public Treatment

To vary the extent to which the distribution of the leaflets is common knowledge among voters within the precinct, we also introduced a “public treatment” variant where leaflet delivery was accompanied by a loudspeaker informing voters that their neighbors were also receiving the information and encouraging them to share it and discuss it. This feature is also shared by the studies in Benin and Uganda 2, where they increase the public nature of their intervention by, respectively, broadcasting their information in a public location and varying the saturation of treatment. In contrast, we refer to leaflet delivery without the loudspeaker as the “private treatment.” The objective of this variant was to make our intervention more public and in this way induce common knowledge and coordination among voters – something that could strengthen the intervention’s effect as MPAP H15 suggests.

However, the coefficients in column 3 of Panel A of Tables 5.2 and 5.3 suggest that, on average, the effect of our common treatment arm was similar in places with the private and public treatment. The estimates for the sample of precincts with good news (Panel C) suggest that the improvement in voters’ evaluations and vote share of the incumbent were slightly larger for the public than for the private variant of our treatment though the differences are not too large. In sum, these estimates suggest that the use of the loudspeaker did not have a major influence on the effect of our intervention.

There are several reasons why the public variant of our intervention may have had a limited additional effect. First, the use of a loudspeaker may simply be ineffective to generate common knowledge or coordination. Many of the leaflet recipients were not at home during leaflet delivery and thus could not have heard the loudspeaker. The fraction of survey respondents who recall listening to a loudspeaker in control precincts and precincts subject to the private treatment was very small: 3 percent. The same average in those precincts subject to the public treatment was indeed higher and statistically different from zero but still somewhat small; only 10 percent of survey respondents in these households remember the loudspeaker. This suggests that our public treatment was too weak and many respondents in precincts that received this treatment simply did not hear the loudspeaker. It may take stronger dissemination mechanisms in traditional media outlets such as radio or

TV.³³ It is also worth noting that loudspeakers like the ones we used are commonly used by politicians during their campaigns. Thus, one possibility is that the loudspeaker increased the perception that our leaflets were being delivered by a political party and not by an NGO, which may have attenuated its effect on coordination. While plausible, this hypothesis does not receive much support in our data either; respondents in precincts with the public treatment were not more or less likely to believe that the leaflet came from a political party or an NGO.

A complementary explanation is that our private treatment was already sufficiently “public” in the sense that it generated the expectation that all other residents in their community received the leaflet. We asked survey respondents to report the fraction of community members that they believe received the leaflet. The majority of respondents in control and treated precincts believed “a very small number” of community members received the leaflet, though the fraction is lower in treated precincts (78 percent in control precincts vs. 62 percent in treated precincts). However, there is no difference in beliefs regarding the fraction of members who received the leaflet across the public and private treatments. Thus, while our treatment on average did not generate the expectation that everyone received the leaflet (consistent with common knowledge), the public treatment had no effect on this perception relative to the private one. It therefore appears to be more likely that the leaflet reached an insufficient numbers of voters to induce collective action.

5.4.2 Spending on the Poor vs. Unauthorized Spending

Across all treatments we use official information from audits conducted by the ASF. Thus, we cannot explore heterogeneous effects according to the credibility or reliability of the source (MPAP H14). However, recall that for each municipality, the leaflet focused on either the proportion of spending that does not benefit the poor, or unauthorized spending. Both types of spending go against the rules stipulated for FISM funds. But while unauthorized spending should be a valence issue opposed by all voters, it is possible that spending on the poor is interpreted by some voters as reflecting the preferences of the incumbent party for redistribution, rather than malfeasance. In this case, we would expect a stronger effect for leaflets with information on unauthorized spending, a hypothesis closely related to that in MPAP H13.

³³ See Larreguy, Marshall, and Snyder (2016).

In column 4, Panel A of Tables 5.2 and 5.3, we show that the coefficients for both types of leaflets are roughly similar, on average. This finding is consistent with the results in Larreguy, Marshall, and Snyder (2016) and suggests that voters interpreted both types of spending as malfeasance. This is not particularly surprising given the relatively high prevalence of poverty in our sample of precincts. In Panels B and C we split the samples between bad and good news, respectively. However, any differential effects of different types of information across these samples must be interpreted cautiously due to power considerations. For example, there are only fifteen precincts in the good news “spending on the poor” cell which may explain the rather large coefficient in column 4 of Table 5.2, Panel C. In the vote share regressions in Table 5.3, the coefficient is larger for the “not spending on the poor” variant in the bad news sample, but larger for the “unauthorized spending” variant in the good news sample. We did not preregister any hypotheses about differences in these treatment variants across the good and bad news samples and thus we prefer not to speculate about these differences, particularly given that some of these coefficients rely on very small samples.

5.5 POLITICAL BACKLASH

In this section we provide both quantitative and qualitative evidence on the reaction of politicians to our intervention. In our survey, we asked respondents whether the incumbent or opposition parties made reference to our leaflets via different (non-mutually exclusive) means such as (a) flyers, (b) campaign rallies, (c) door-to-door visits of campaign operatives, (d) loudspeakers, or (e) traditional media outlets such as radio or TV. We also asked them whether the party was trying to either take attention away from or direct attention towards the audit information reported in the leaflets, apologize for the information reported, or simply claim that all parties are the same.

To explore more systematically whether our intervention triggered a reaction from politicians, we code an individual-level measure for the number of means through which politicians mentioned the leaflet ranging from 0 (none) to 5 (all means). In Table 5.5, we report estimates from OLS regressions to test whether respondents in treated precincts reported more reactions from the incumbent (columns 1 and 2) or opposition (challenger) parties (columns 3 and 4). The estimates in columns 1 and 3 suggest that, on average, our intervention triggered a political response

TABLE 5.5 *Mexico: Effect of information treatment on political party responses*

	Total party activities			
	Incumbent reactions (1)	(2)	Challenger reactions (3)	(4)
Information treatment	-0.010 (0.043)	-0.166** (0.078)	0.074* (0.039)	-0.038 (0.052)
Information treatment × Incumbent malfeasance spending		0.727** (0.272)		0.522** (0.205)
Observations	3019	3019	3019	3019
R ²	0.15	0.16	0.13	0.13
Control outcome mean	0.43	0.43	0.40	0.40
Control outcome std. dev.	1.18	1.18	1.17	1.17
Treatment mean	0.62	0.62	0.62	0.62
Treatment std. dev.	0.48	0.48	0.48	0.48
Interaction mean		0.21		0.21
Interaction std. dev.		0.17		0.17

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 $p < 0.01$.

by opposition but not by incumbent parties (coefficient in column 1 is small and statistically insignificant). However, these average effects once again mask substantial variation in the malfeasance of the incumbent party reported in the leaflet. In columns 2 and 4 we interact our treatment dummy with the level of malfeasance reported in the leaflet. The interaction terms are positive and statistically significant, suggesting that politicians from both incumbent and challenger parties were more likely to react to the leaflets and reach out to voters whenever the leaflets reported high levels of malfeasance. Interestingly, while the average (level) effect is lower for the incumbent than for the challenger, this pattern is reversed for the interaction (slope), which suggests that incumbent parties' decisions to react depends to a large extent on their malfeasance levels.³⁴ These findings suggest that politicians strategically

³⁴ We also conducted a list experiment in order to assess whether vote buying was higher in treatment precincts. We do not find any evidence of this. However, the list experiment we used for this exercise may not have been effective since it detected overall vote buying rates much lower than those found by previous studies of clientelism in Mexico.

reallocated campaign efforts across precincts, which could – as noted above – explain the positive average treatment effect on the incumbent’s vote share, if incumbent responses were more effective than challenger responses.

The quantitative evidence presented in Table 5.5 illustrates the political reactions experienced by voters in treated precincts. However, our staff and implementation teams also faced significant political push-back from municipal incumbents and local political party operatives. From the moment leaflet distribution began, staff from *Borde Político* started receiving calls with complaints about the content of the leaflet, and inquiring about the possibility of producing additional leaflets disseminating negative information about the opposing party.

While initially our field team was welcomed in most locations, we faced some difficulties early on. On five occasions – Dolores Hidalgo and Juventino Rosas, Guanajuato; Cuatiltlán Izcali, Ixtapaluca, and Villa Victoria, Estado de México – several of our team members were taken by the local police under the excuse that people – often local political party operatives – complained that they were disrupting public order, or under false accusations that they were distributing leaflets with political advertisement and that they lacked a commercial permit to do so. In Villa Victoria, residents threatened our team by detaining them until the police arrived and burnt their car. A police unit temporarily detained our team under the accusation of distributing negative information about the municipality. In all cases, our team members were released several hours later since no charges could be pressed. Nonetheless, in all of these places our teams were threatened by the municipal authorities if leaflet distribution continued. We contacted the corresponding State Electoral Institutes. In the state of Guanajuato, officials contacted PRI representatives and the mayor of Dolores Hidalgo with warnings of pressing electoral charges if they continued to prevent the leaflet distribution. Our team was eventually allowed to conclude the delivery of the leaflets. In Juventino Rosas the police intervened after leaflet distribution was completed so this did not interfere with the implementation. However, we could not resume leaflet delivery in Cuatiltlán Izcali, Ixtapaluca, and Villa Victoria since the officials from the State Electoral Institute never replied to requests from *Borde Político*.³⁵

³⁵ In addition, we consulted with our teams in the field and they considered that resuming leaflet delivery would not be advisable for security reasons.

More concerning, in Aquismón, San Luis Potosí, our field staff were intercepted by alleged local police officials in black SUVs with polarized windows, and were told that their safety could not be guaranteed if they continued distributing the leaflets. Following our security protocol, our team members left the municipality immediately but were followed by the SUVs for more than 60 kilometers to the closest city. That same evening, our field staff suffered further intimidation. Leaflet delivery was temporarily suspended in all municipalities in San Luis Potosí and was resumed only after a proper safety assessment by local experts on a municipality-by-municipality basis. Our team, however, never returned to Aquismón.

On other occasions, while local police and party operatives did not directly confront our team, they did take actions towards sabotaging or neutralizing our intervention. In Juventino Rosas, Guanajuato, after our team completed the intervention, PAN representatives removed the leaflets left on the doors of households who were absent during the delivery. Lastly, in Cuatiltlán Izcali, Estado de Mexico, following the detainment of our team, local police also confiscated the leaflets they were distributing. Days later, forged versions of those leaflets, attacking another party (the PAN) were apparently distributed on behalf of Borde Político. Figure 5.5 shows both the original and forged version of the leaflet.

However, the political backlash we experienced in the field provided a unique opportunity to raise attention to this issue. In light of the events in Cuatiltlán Izcali, and in order to prevent a defamation lawsuit from the PAN, Borde Político took immediate actions. First, it reached out to PAN officials to clarify that they were not responsible for the forged leaflets that had been distributed in Cuatiltlán Izcali. Second, it filed protests with the Prosecutor for Electoral Crimes (FEPADE) of the PGR (Attorney General's Office) not only about the situation in Cuatiltlán Izcali, but also about all the other incidents discussed above. Third, Borde Político wrote a press release explaining those incidents which led legislators and senators from various parties to push for a resolution of the Permanent Commission of Mexican Congress.³⁶ The resolution, which was ultimately approved (http://www.senado.gob.mx/64/gaceta_del_senado/documento/55379), exhorted the FEPADE of the PGR to investigate Borde Político's charges

³⁶ The Permanent Commission is formed by legislators and senators and takes care of urgent matters during recesses.



FIGURE 5.5 Mexico: Original and forged leaflet in Cuautitlán Izcalli, Estado de México

against the municipal authorities that abused their power in order to guarantee Borde Político's right to freedom of expression and our team's safety.

5.6 CONCLUSION

The results of our intervention highlight the importance of considering how information provided to voters compares with their prior beliefs. Contrary to our initial expectations, even moderate levels of malfeasance were often interpreted by voters as good news about the incumbent party. This resulted in higher electoral support, on average, for the incumbent party in treated precincts. Future research should try to understand the determinants of voters' prior beliefs. Interventions aimed at increasing voters' expectations and demands on politicians may also be a prerequisite for other informational interventions to produce the desired effect.

While we focus on the differential effect (slope) of relatively better news, our estimates also suggest that, surprisingly, even information interpreted by voters as bad news had a positive effect on the incumbent's vote share. We conjecture that this result may be partly explained by the decreased uncertainty regarding the parties' malfeasance levels that our intervention generates, and also, potentially, by the responses of politicians to the dissemination of such information.

Our heterogeneous effects also suggest that informational interventions may be particularly effective in electorally competitive environments (where voters are more likely to change their voting decision in response to new information) and in places where voters don't already have access to other media outlets or sources of political information.

Finally, our findings on politician's reactions show that researchers should consider the general equilibrium effects of their interventions.³⁷ Politicians, who are often the target of many of the interventions conducted, often respond strategically in an attempt to prevent any detrimental effect of the intervention on their political careers. Since this is not always the case,³⁸ our understanding of equilibrium political behavior may merit further research examining the conditions under which incumbents and challengers engage with informational interventions in different ways. This would be particularly interesting to explore in the context of earlier interventions occurring before candidate and platforms are selected.

³⁷ Cruz, Keefer, and Labonne (2016).

³⁸ See, for example, Banerjee et al. (2011).