DOES THE CONTENT AND MODE OF DELIVERY OF INFORMATION MATTER FOR ELECTORAL ACCOUNTABILITY? EVIDENCE FROM A FIELD EXPERIMENT IN MEXICO*

ERIC ARIAS† HORACIO A. LARREGUY‡ JOHN MARSHALL§ PABLO QUERUBÍN¶

OCTOBER 2018

Evidence that information campaigns help voters select better politicians is mixed. We propose that comparative performance information and public dissemination may moderate information's effects on electoral accountability, by respectively helping voters to identify malfeasance incumbent parties and facilitating coordination around the information. We test these mechanisms using a large-scale field experiment that provided voters with the results of audit reports documenting mayoral malfeasance before the 2015 Mexican municipal elections. We find that neither benchmarking incumbent performance against mayors from other parties within the state, nor accompanying leaflet delivery with loudspeakers announcing the leaflets' delivery, significantly moderated the effects of information on voter beliefs or incumbent party vote share. Comparative performance information's ineffectiveness likely reflects voters' limited updating from the particular comparison provided, while the loudspeaker created common knowledge without meaningfully facilitating voter coordination. The results highlight challenges in designing informational campaigns to capture the theoretical conditions conducive to electoral accountability.

*We thank the steering committee and other team members of the EGAP Metaketa initiative for illuminating discussions and useful comments. We are extremely grateful to Anais Anderson, Adriana Paz, and Alejandra Rogel, and the Data OPM and Qué Funciona para el Desarrollo teams for their implementation of this project, as well as to Juan Carlos Cano Martínez, Executive Secretary of the Guanajuato Electoral Institute, for his assistance in responding to municipal governments that tried to prevent our treatment's dissemination. We are grateful to Tommaso Nannicini and Francesco Trebbi, and Frederico Finan and Laura Schechter, for sharing their survey instruments. This research was financed by the EGAP Metaketa initiative, and was approved by the Harvard Committee on the Use of Human Subjects (15-1068) and the New York University Committee on Activities Involving Human Subjects (15-10587). Our pre-analysis plan was pre-registered with EGAP, and is publicly available here.

†Department of Political Science, William and Mary College. Email: eric.arias@princeton.edu.

‡Department of Government, Harvard University. Email: hlarreguy@fas.harvard.edu.

§Department of Political Science, Columbia University. Email: jm4401@columbia.edu.

¶Department of Politics, New York University. Email: pablo.querubin@nyu.edu.

1 Introduction

Theoretical models of political accountability suggest that incumbent performance information is essential for enabling voters to identify and elect desirable politicians (Fearon 1999; Manin, Przeworski and Stokes 1999; Rogoff 1990). Such electoral accountability is especially important in developing contexts where weak political institutions may otherwise fail to constrain corruption, incompetence, and clientelism (Pande 2011). However, the experimental and quasi-experimental evidence that informing voters of poor performance results in electoral sanctioning that has now accumulated is markedly mixed. On one hand, some studies—often disseminating information via broadcast or print media—report that the revelation of relatively good performance is indeed rewarded, while sufficiently bad performance results in electoral sanctions (Banerjee et al. 2011; Ferraz and Finan 2008; Humphreys and Weinstein 2012; Larreguy, Marshall and Snyder 2018). On the other hand, a number of other studies observe little effect of providing information (de Figueiredo, Hidalgo and Kasahara 2014; Dunning et al. forthcoming), or even that poor incumbent performance disproportionately harms challengers when voters disengage (Chong et al. 2015).

In this article, we test two potentially-critical theoretical mechanisms that could help to account for the varying efficacy of informational interventions: the absence of benchmarked performance indicators; and the reliance on private, rather than public, modes of information dissemination. While both components are frequently touted as central ingredients in supporting electoral accountability, and are often included as part of multifaceted information campaigns, we are aware of no prior study that experimentally separates the causal contributions of each component.

First, if voters only receive information pertaining to their incumbent's performance (what we call *local* performance information), they may struggle to distinguish their incumbent's aptitude from common shocks, such as budgetary reforms, that affect the performance of all incumbents holding similar offices at the same time. This could lead naive voters to excessively punish (reward) bad (good) performance, or Bayesian voters to largely disregard performance signals driven by factors beyond their particular incumbent's control. However, by providing a relevant cross-sectional benchmark, our simple model shows how *comparative* performance information may influence voters through two channels: (a) enabling voters to more accurately evaluate their incumbent's performance by filtering out the common component of performance that is expected from both incumbent and challenger parties; and (b) directly informing voters about challenger parties' performance (Aytaç 2018; Besley 2006; Holmstrom 1982; Kayser and Peress 2012; Meyer and Vickers 1997). For example, consider performance indicators where higher values indicate greater aptitude. We anticipate that providing voters with comparative performance information will induce less (more) favorable appraisals of the incumbent's aptitude than simply providing local performance information when also providing challenger performance information causes voters to conclude that a larger (smaller) common shock was present than previously believed. Such updating about the common shock occurs when comparison units performed better (worse) than expected. Moreover, to the extent that comparison units are perceived to resemble local challenger parties, better (worse) than expected performance by comparison units will differentially decrease (increase) incumbent vote share, relative to non-benchmarked incumbent performance information.

Second, models of citizen coordination suggest that informational interventions may fail if voters do not know that other voters have also received such information and do not expect other voters to change their voting behavior accordingly. However, public dissemination of performance information can induce tacit coordination among voters by establishing common knowledge and expectations or induce explicit coordination by stimulating discussion within communities (Chwe 2000; Morris and Shin 2002; Nickerson 2008). This could amplify the effects of information dissemination on electoral accountability to the extent that voters expect others to vote on the basis of their posterior beliefs about incumbent aptitude and wish to coordinate their voting behavior with others. Without this mechanism, there may be little incentive for individual voters to forgo clientelistic benefits to hold politicians to account. Alternatively, by increasing attention or discussion, public dissemination could increase the likelihood that voters receive and engage with performance indicators. Both mechanisms are often attributed to mass media (Adena et al. 2015; Arias forthcoming; Enikolopov, Makarin and Petrova 2016; Manacorda and Tesei 2016; Yanagizawa-Drott 2014), which may explain why extant mass dissemination studies generally report larger accountability-enhancing effects.

To test the relevance of these potential moderators of the effects of information campaigns, we conducted a large-scale field experiment around Mexico's 2015 municipal elections. Specifically, we randomly varied the *content* and *form* of leaflets containing incumbent performance information—relating to mayoral malfeasance—that were disseminated in treated electoral precincts. The experiment, which was conducted across 678 electoral precincts from 26 municipalities in four states of central Mexico, informed voters of the results of independent audit reports documenting unauthorized spending and illegal misallocation of funds to municipal projects not benefiting impoverished localities. We build on Arias et al.'s (2018) earlier analysis of the same experiment, which focused on demonstrating how voters' priors beliefs moderate their responses to receiving new information. Importantly for interpreting our findings, Arias et al. (2018) show that revealing often severe levels of mayoral malfeasance in a voter's municipality increased the incumbent party's vote share on average, due to voters' pessimistic prior beliefs and increased certainty about incumbent party malfeasance, as well as effective reactions to the treatment by incumbent parties. However, rather than pool across information treatments, this article instead investigates the distinct effects of experimentally varying core theoretical mechanisms underpinning information provision's effects on belief updating and electoral accountability.

Our primary contribution is thus to explore how benchmarked and publicly-disseminated variants of incumbent performance information affect voters' posterior beliefs and voting behavior. Specifically, in addition to delivering leaflets reporting the share of audited funds spent by the incumbent government that did not comply with federal government rules, we examine two additional treatment conditions that: (1) benchmarked municipal mayors against the average performance of mayors from other parties within their state; and (2) accompanied leaflet delivery with a loud speaker announcing that the leaflets were being delivered throughout the local community. Both conditions successfully achieved their desired "first stages": voters receiving these treatments were, respectively, significantly more likely to recall encountering benchmarked information or a loudspeaker, while the loudspeaker increased common knowledge that the leaflets were delivered throughout the community. To our knowledge, this article is the first to experimentally unpack the roles of cross-sectional benchmarks and common knowledge signals in examining the effects of providing incumbent performance information on accountability and electoral outcomes.

However, we find little evidence that either variant of the basic information treatment differentially influenced voting behavior. Failing to support our pre-registered theoretical expectations, comparative performance information—that generally detailed lower than expected levels of malfeasance among other parties within the state—was not more likely to decrease the vote share for the incumbent party than simply providing incumbent-only information, even when challengers' performance notably exceeded expectations. Our survey data indicate that voters did not differentially update their beliefs from such comparative information, suggesting that voters either did not understand the benchmark or did not regard it as relevant. Similarly, we find no evidence that public dissemination amplified, or otherwise moderated, the effect of distributing leaflets. Despite increasing common knowledge, we do not find evidence that the loudspeaker significantly enhanced voter coordination. In each case, our design is powered to exclude even relatively small effect sizes with significant confidence.

Our primary contribution is thus to show that these theoretically-motivated interventions are insufficient to break voters out of low-accountability political equilibria, at least in the context of Mexican municipal politics. As Khemani et al. (2016) optimistically argue, informed participation on the part of citizens has the potential to "make politics work for development" by changing the incentives for politicians to serve their constituents and helping voters to select those most likely and able to do so. Despite their theoretical promise in mitigating these agency problems, we find little evidence to suggest that either of our interventions—which, especially in the case of providing comparative performance information, could be scalable—meaningfully altered voter behavior. These rather pessimistic findings—which challenge widely-held assumptions about how information influences electoral accountability—suggest that information interventions may need to be more specifically targeted to ensure their relevance and comprehensibility, while more powerful and larger-scale interventions may be required to induce coordinated efforts. Nevertheless, our null findings advance the study of accountability and inform the design of future information dissemination campaigns by highlighting types of interventions unlikely to moderate the effects of simply providing incumbent performance indicators.

2 Theoretical framework

There are good reasons to believe that informing voters about incumbent performance in office could affect voting behavior. An influential theoretical literature argues that signals of incumbent performance can update voter beliefs, and thereby help voters to prospectively identify and elect competent politicians more likely to represent their interests (Fearon 1999; Rogoff 1990), while also potentially mitigating moral hazard once in office (Barro 1973; Ferejohn 1986).[1] Several studies provide empirical evidence consistent with such a belief updating channel (e.g. Arias et al. 2018; Banerjee et al. 2011; Ferraz and Finan 2008; Humphreys and Weinstein 2012; Kendall, Nannicini and Trebbi 2015).

In this article, we empirically extend the learning framework to incorporate the provision of comparative performance information and the public dissemination of performance information. The following subsections discuss the theoretical rationale for these extensions to the basic Bayesian framework, and state the implications for voter beliefs and voting behavior that we pre-registered ahead of implementing our experiment.[2]

2.1 Comparative performance information

We consider the role of comparative performance information—in our setting, the provision of information that compares the incumbent party's observed malfeasance to that of potential challenger parties that are incumbents in comparable municipalities—in the context of a learning framework

[1] However, prospective voters may not be able to commit to punishing low effort in office (Fearon 1999).
[2] The few deviations from our pre-analysis plan are justified in Appendix section A.3 (p. 9).

where observed malfeasance reflects: i) the unobserved underlying malfeasance of an incumbent party, and ii) unobserved factors equally influencing the malfeasance indicators of all incumbents. This framework characterizes voters as Bayesians seeking to select the least malfeasant politician on the basis of the information available, and emphasizes the importance of how voters' prior beliefs relate to indicators of malfeasance for determining the effect of information on support for the incumbent party. We use this framework to predict how the effect of comparative performance information differs from that of providing information only about the incumbent without a benchmark (local performance information). Appendix section A.1 (pp. 3-7), which builds on our pre-analysis plan, formalizes the theoretical expectations we describe below.

The signal extraction problem for voters is to separate an incumbent party's propensity to be malfeasant from the effects of "common shocks" that afflict all municipalities (Holmstrom 1982; Meyer and Vickers 1997). For example, nationwide budgetary shifts, decentralization reforms, or economic pressures could represent common shocks that influence malfeasance indicators in all municipalities without reflecting an individual incumbent party's malfeasance. In the absence of performance benchmarks, voters' evaluations of the incumbent become less favorable to the extent that reported malfeasance indicators, adjusted for prior expectations of common shocks affecting all incumbents, exceed prior expectations.

Information about the malfeasance of incumbents from different parties in other municipalities provides voters with a second signal of performance that relates to challengers, at least where party candidates' characteristics are believed to be correlated across municipalities. This signal most obviously informs voters about challenger malfeasance, and how this compares with indicators of incumbent malfeasance. However, it also enables voters to more accurately update their posterior beliefs about their incumbent party's underlying malfeasance by filtering out their updated belief about the common component of malfeasance driving the observed performance of both incumbent and challenger parties. In particular, as our Appendix proves, if challengers perform better (are less malfeasant) than voters expected, then Bayesian voters will infer that the common shock partly driving the observed malfeasance level of the incumbent is lower than they had expected. Because of this, comparative performance information would thus lead voters to update more unfavorably about the incumbent relative to only receiving a signal of incumbent malfeasance in this example.

In contrast with posterior beliefs about levels of incumbent malfeasance, voting behavior reflects a voter's beliefs about the *relative* levels of incumbent and challenger party malfeasance. We assume that voters will re-elect the incumbent party to the extent that they believe—conditional on their partisan attachments—that the challenger would be less malfeasant than the incumbent. In addition to enabling a more precise assessment of the incumbent's malfeasance (by allowing voters to filter out the common shock), comparative performance information also helps voters to update their beliefs about the challenger's malfeasance. Thus, comparative performance information will induce greater electoral sanctioning of the incumbent, relative to only receiving a signal of incumbent malfeasance, when such information induces voters to update favorably about the challenger.[3]

Based on the specific informational content that voters would receive, this theoretical framework implies a number of testable implications. We focus on comparing the effects of comparative vs. local information, rather than on the overall effect of these treatments relative to the control group.[4] As detailed in section 4, the typical voter in our sample received information revealing severe incumbent malfeasance in comparison with zero or low malfeasance among challengers (municipalities governed by different parties within the same state). Applied to our framework, we assumed—at the point of pre-registering our hypotheses—that indicators of incumbent malfeasance would exceed prior expectations, while indicators of challenger malfeasance would fall below prior expectations. We thus conjectured that, for the average voter, the provision of comparative information would lead voters to update relatively more unfavorably about the incumbent party than the provision of local information. More concretely:

H1. *Relative to local malfeasance information, comparative malfeasance information will, on*

[3] To the extent that a benchmarked signal increases the precision of voters' posterior beliefs, there is also a second-order effect such that sanctioning is also increasing in unfavorable updating about the incumbent.

[4] Appendix section A.1 (pp. 3-7) also derives testable implications for the effect of comparative information relative to a no-information control group.

average, increase voters' posterior beliefs that the incumbent party is malfeasant and decrease voters' posterior beliefs that the challenger is malfeasant.

Furthermore, in line with Bayesian updating, we expected to observe the following heterogeneous effects that adjust for how the information relates to voters' prior beliefs:

H2. *The effect of comparative malfeasance information on posterior beliefs about the incumbent's (challenger's) malfeasance, relative to only providing local malfeasance information, will be decreasing (increasing) in the difference between reported challenger (incumbent) party malfeasance and voters' prior expectations of such malfeasance.*

If hypotheses H1 and H2 hold, we then expect belief updating to in turn affect voting behavior. Although the information provided may not be important to all voters, some voters are likely to vote in line with changes in their beliefs about expected malfeasance—a key issue for voters in Mexico and many developing countries. Based on the content of the information that voters would receive, we conjectured the following hypotheses pertaining to the average electoral effects of providing comparative performance information:

H3. *Relative to local malfeasance information, comparative malfeasance information will, on average, decrease the incumbent party's vote share.*

H4. *The effect of comparative malfeasance information on incumbent party vote share, relative to only providing local malfeasance information, will be increasing in the difference between reported challenger party malfeasance and voters' prior expectation of such malfeasance.*

Previous non-experimental studies in predominantly developed contexts provide some evidence consistent with the importance of such cross-sectional benchmarks. Kayser and Peress (2012) show that media benchmarking of national economic performance relative to the international economy enables OECD voters to filter out common shocks and hold incumbents more electorally accountable for domestic than international components of economic growth. Similarly, Aytaç (2018) finds that incumbent vote shares in democracies increase with economic growth relative to both previous domestic and contemporaneous international growth rates. Focusing on U.S.

governors, Besley and Case (1995) also provide evidence that relative performance evaluation by voters induces incumbents to engage in yardstick competition vis-à-vis neighboring incumbents when setting tax rates. These studies assume that performance signals are widely observed by voters and are uncorrelated with other factors driving incumbent support. Our approach instead compares voters randomly assigned to receive different information, and explicitly considers how these relate to voters' prior beliefs.

2.2 Public dissemination

While we expect the content of the information provided to affect voters' posterior beliefs ahead of elections, and ultimately their voting behavior, the mode of information delivery may be equally important in influencing such behavior. In particular, we consider the potential role of public forms of dissemination, where information is provided in a manner designed to ensure that voters not only receive the information themselves but also know that other voters in their community also received the same information. In game theoretic jargon, this is referred to as a public signal.

Providing a public signal in addition to privately providing the same information to all voters could affect voter behavior through several coordination mechanisms. First, the common knowledge induced by a public signal could tacitly coordinate voter behavior. For example, Morris and Shin (2002) use a model based on a Keynesian beauty contest to demonstrate that a public signal can coordinate behavior around that signal when there exist complementarities to individuals acting as they expect others to. In our electoral context, this could reflect voters becoming more likely to vote in accordance with their posterior beliefs about malfeasance because they believe that others will also do so, which will entail sending a clearer message to politicians that malfeasance will not be tolerated (e.g. Lohmann 1993). As Morris and Shin (2002) highlight, such coordinated action could occur—and be socially suboptimal—even among voters that would prefer to vote on the basis of other issues, or do not update their beliefs.

Second, a public signal could similarly increase communication between voters that helps them solve the problem of coordinating on the best candidate by establishing the expectation that other voters are also willing to act on the information provided (Chwe 2000). In this instance, this could entail coordinating on a costly but beneficial action in response to the information's content. Communication could also explicitly induce coordination, e.g. where voters meet in response to the information and debate until an explicit agreement pertaining to a common response is reached.

Third, even without inducing tacit or explicit coordination, the public signal could enhance discussion and engagement with the information that increases the probability that the information is internalized and that beliefs are consolidated. In this respect, a public signal could instigate information diffusion within a social network (e.g. Alatas et al. 2016; Jackson 2010), and induce a seemingly collective response because any given individual becomes more likely to respond privately. Ultimately, each mechanism likely implies that the public signal will amplify the effect of the information's content.[5]

Applied to our empirical setting, we expected that—relative to a private mode of information provision—a public mode of information dissemination would increase the magnitude of responses to the information provided. We thus hypothesized that:

H5. *The magnitude of all (average and heterogeneous) effects of providing information on the incumbent party's vote share (i.e. H3 and H4) are greater when the information is delivered through public than private dissemination mechanisms.*

Since public signals do not necessarily require that individuals differentially update their beliefs about incumbent malfeasance, such amplification effects may not apply to hypotheses H1 and H2.

Various previous studies, particularly those focusing on the media, suggest that a public signal can meaningfully alter citizen behavior. In the case of electoral accountability, the effects of performance information disseminated through the media are often larger in magnitude than the effects of providing similar information to individuals privately (e.g. Banerjee et al. 2011; Chong et al. 2015; Ferraz and Finan 2008; Larreguy, Marshall and Snyder 2018). Similarly, modern

[5] Exceptions to this hypothesis may arise where public signals induce voter coordination around other potentially competing objectives. For example, coordination around relative posterior beliefs might induce different responses than coordination around the direction of updating.

technologies such as cellphones and social media appear to have coordinated protest participation across Africa (Manacorda and Tesei 2016), and in France (Larson et al. 2017) and Russia (Enikolopov, Makarin and Petrova 2016). While public information dissemination may have contributed positively to social welfare in such cases, radio access may also have helped coordinate anti-Semitic acts and electoral support for the Nazis (Adena et al. 2015) and the Rwandan genocide (Yanagizawa-Drott 2014). While such evidence is suggestive, no study of which we are aware has yet identified the differential effect of providing information publicly, rather than privately, to voters.

3 Mayoral malfeasance in Mexico

Mexico's federal system is divided into 31 states and the Federal District of Mexico City. The states, in turn, contain almost 2,500 municipalities, which receive state and federal transfers but often possess a limited capacity to raise their own revenues. Decentralization reforms in the 1990s empowered municipal governments to play an increasingly important role in public service and local infrastructure provision (Wellenstein, Núñez and Andrés 2006). Municipal governments are led by mayors who preside over 20% of total government spending. Although mayors were able to run for re-election for the first time in 2018 in most states, mayors were previously elected to non-renewable terms typically lasting for three years.

3.1 Independent federal audits of municipal spending

Mayors are responsible for delivering basic public services and managing local infrastructure. A major source of funds for investments in local infrastructure is the Municipal Fund for Social Infrastructure (FISM), which represents 24% of the average mayor's annual budget, and often substantially more than revenues raised by municipal taxes. FISM funds are direct federal transfers mandated exclusively for infrastructure projects—such as investments in water supply, drainage, electrification, health infrastructure, education infrastructure, housing, and roads—designed to improve public service delivery in localities defined by the federal government as impoverished.[6] Within these restrictions, mayors have the discretion to choose where and what types of projects are pursued. Although these often represent major projects, voters are poorly informed about mayoral responsibility for such provision (Chong et al. 2015).

FISM transfers are subject to independent audits by the Federal Auditor's Office (ASF). The ASF has constitutional authority to audit the spending, accounting, and management of federal funds, and is perceived to be neutral, autonomous, and professional (De La O and Martel García 2015). It can impose fines, recommend economic sanctions, and file or recommend criminal prosecution on the basis of its reports. Each year, the ASF audits around 150 municipalities, selected on the basis of the share of FISM funds in the municipal budget, previous performance, factors increasing the risk of mismanagement, and the recency of the last audit (see Auditoría Superior de la Federación 2014). Audits are not announced or conducted until the year after spending occurred, and reports are presented to Congress in the February two calendar years after audited spending occurred. The reports are publicly available online at asf.gob.mx.

ASF reports examine FISM administration across a variety of dimensions, but we focus on two key dimensions that form the basis of the information provided by our experiment: (1) the share of FISM funds spent on projects that did not benefit the poor; and (2) the share of FISM funds spent on unauthorized projects. Projects not benefiting the poor represent social infrastructure investments completed in localities that are not classified as impoverished. Unauthorized projects are non-social infrastructure projects, which in practice are often similar to the corrupt practices documented in Ferraz and Finan (2008), e.g. procurement violations and electorally-targeted projects. We refer to both violations as malfeasance.

Unfortunately, municipal malfeasance is not uncommon. ASF reports released between 2007 and 2015 document that, on average, 8% of audited funds were spent on projects that did not benefit the poor, while 6% were spent on unauthorized projects. Given the magnitude of the program, these sums are not trivial. Furthermore, malfeasance is often concentrated in particular municipalities,

[6]According to their marginalization index, the National Population Council (CONAPO) identified that 22.7% of citizens were living in impoverished localities in 2010.

and instances of FISM violations can be egregious. For example, nine mayors across the state of Tabasco diverted FISM toward the 2012 electoral campaigns of their parties' candidates,[7] the mayor of Oaxaca de Juárez created a fake union to collect illegal payments,[8] and 12 projects followed irregular tenders in Altamira in 2014.[9]

While the results of FISM audits can be reported locally, and media coverage does influence electoral accountability in urban areas (Larreguy, Marshall and Snyder 2018), voters—as in many developing contexts (Keefer 2007; Pande 2011)—are generally poorly informed about mayoral use of FISM funds. Chong et al. (2015) note that only around 10% of voters are aware of the FISM program. However, although voters are not generally aware of the program, dissatisfaction with services is high; 53% of voters are unsatisfied with service provision and 42% believe the municipal government to be dishonest. Knowledge of mayoral performance in other municipalities is also likely to be low.

3.2 Municipal electoral competition

Electoral competition in most Mexican municipalities is between two of the country's three largest parties. In most parts of the country, the populist PRI competes against either the relatively urban right-wing National Action Party (PAN) or the PRI's more rural left-wing offshoot Party of the Democratic Revolution (PRD).[10] The two dominant parties in a given municipality increasingly form coalitions with smaller parties, and in some cases one of the other large parties, as part of their electoral ticket. This further solidifies the *de facto* two-party competition reflected in the 2.5 effective number of party coalitions in the average municipal election.

Municipal election campaigns are generally oriented around political parties rather than specific candidates for various reasons. First, voters are much better informed about parties than

[7] *Tabasco Hoy*, "Pagaron pobres campañas 2012," March 6th 2014.

[8] BBM Noticias, "ASF: desvió Ugartchechea 370.9 mdp," October 21, 2013, link here.

[9] Centro Noticias Tamaulipas, "ASF detectó anomalías del FISMDF en Altamira," March 18th 2016; link here.

[10] The National Regeneration Movement (MORENA) party stood for the first time in 2015, and obtained 9% of the federal legislative vote. However, MORENA's local presence was more limited.

individual politicians (e.g. Chong et al. 2015; Larreguy, Marshall and Snyder forthcoming). In our sample, for example, 80% of survey respondents correctly identify the party of their municipal incumbent; this substantially exceeds individual politician name recognition at all levels of government. Second, voters may recognize that Mexico's main parties continue to use distinct candidate selection mechanisms that select candidates with similar characteristics over time (Langston 2003). Third, voters have consistently been shown to hold parties responsible for the actions of individual politicians (e.g. Chong et al. 2015; Larreguy, Marshall and Snyder 2018; Marshall 2018). Consequently, despite the fact that mayors could not themselves seek re-election, there are good reasons to believe that voters will hold their party responsible for their actions in office.

4 Experimental design

We conducted a field experiment around the June 7, 2015 Mexican municipal elections to test the theoretical predictions enumerated in section 2. The following subsection describes sample selection, treatment conditions, outcome variables, and our estimation strategy.

4.1 Sample selection

The experiment was conducted across 26 municipalities in the central states of Guanajuato, Estado de México, San Luis Potosí, and Querétaro. Of the municipalities in these states where an audit report was released in February 2015, the 26 were chosen to minimize safety risks to our implementation team, to match the distribution of incumbent parties across municipalities in these four states, and to maximize variation in reported malfeasance across municipalities subject to the constraint that at least one of our two measures of reported malfeasance was at least two percentage points higher or lower than the state average among audited opposition parties.[11] Figure 1 shows

[11]The municipalities of Aquismón and Villa Victoria were replaced by Atlacomulco, Temoaya, and additional precincts in Tlalnepantla de Baz because our team immediately received threats upon entering these municipalities. Combined with our block randomization design (see below), the risk of bias is likely to be minimal because replacement is uncorrelated with treatment due to the teams leaving before leaflets were delivered.

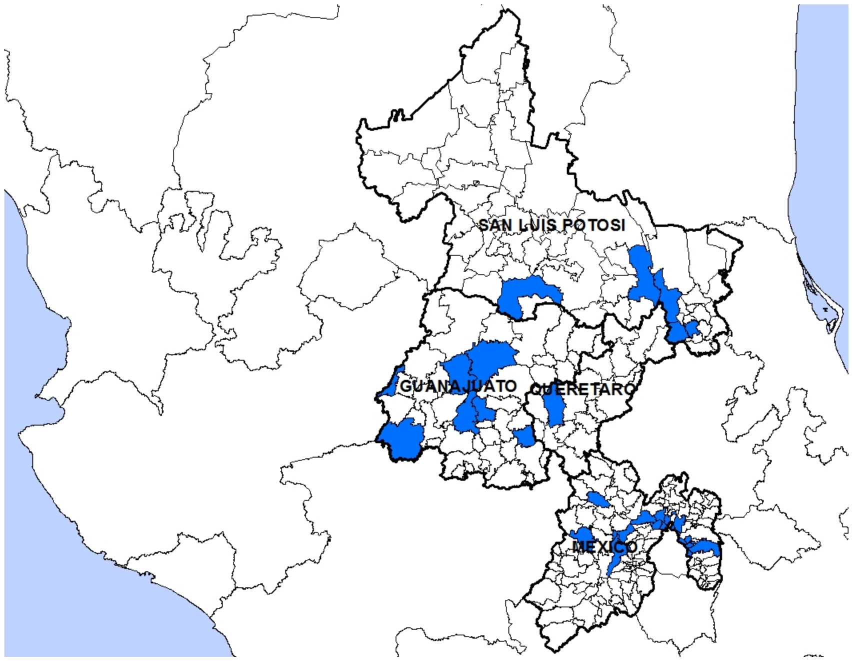

Figure 1: The municipalities in our experimental sample

the location of these municipalities.

Within each municipality, up to one third of electoral precincts—Mexico's lowest geographical level of electoral aggregation, containing around 1,250 voters on average —were then selected for our experimental sample. We oversampled precincts from municipalities with high or low levels of incumbent malfeasance and stark contrasts with other parties. Priority was given to small but accessible rural precincts and small urban precincts minimizing the number of neighboring precincts included in the experimental sample, thereby reducing the risk of cross-precinct spillovers.[12] Appendix Table A1 (p. 9) shows that the resulting sample of 678 precincts is broadly representative of Mexico sociodemographically.

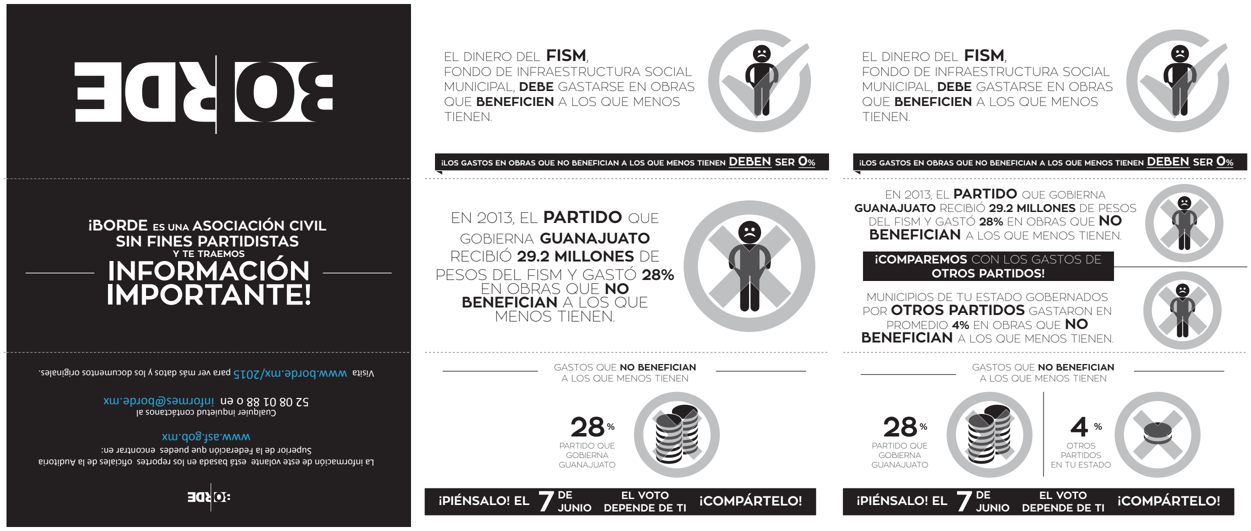

Figure 2: Example comparing local and comparative leaflets from Guanajuato, Guanajuato

4.2 Treatment conditions

Partnering with Mexican NGO Borde Político, our baseline treatment disseminated leaflets documenting the results of the ASF audits. As the example from Guanajuato, Guanjuato—shown in Figure 2—illustrates, the leaflet explained that FISM funds were intended for social infrastructure projects benefiting the poor, before reporting the total amount of funds (29.2 million pesos) received by the municipal government and the share of those funds (28%) spent on either (but not both) projects not actually benefiting the poor (as in this example) or unauthorized projects (see the example in Appendix Figure A1 (p. 7)). The front of the triptych notes that Borde Político is a non-partisan NGO and that the information provided can be accessed on the ASF website. The leaflet refers to the incumbent governments without explicitly naming the incumbent parties, and were designed in black and white colors, to minimize association with any particular political

[12]In urban areas, we restricted our sample to precincts with at most 1,750 registered voters and designed an algorithm to minimize the number of neighboring precincts. This entailed identifying all neighboring precincts that were eligible for our sample and iteratively removing the precinct with the most in-sample neighbors until we reached the specified number of precincts for that municipality.

party.[13]

Up to 200 leaflets were delivered to households in each precinct either by hand, or by placing them in a mail box or pinning them to doors where there was no mail box. We were able to reach 57% of households in the average precinct. Leaflets were delivered over the month before the election, and compliance with our delivery protocols was generally very good.[14] Leaflet delivery locations were logged by our enumerators so that our post-election survey team could interview only leaflet recipients in treated precincts.

To investigate how the content and mode of information provision affects voters, we varied our intervention along two dimensions corresponding to the hypotheses generated by the theoretical considerations outlined above. First, to identify the effect of providing voters with a benchmark against which to compare their incumbent party's malfeasance, we also delivered comparative leaflets. In contrast with the *local* leaflet in the middle panel of Figure 2, the *comparative* leaflet in the third panel also provided information about the mean outcome among all audited municipalities within the same state that were governed by a different political party. Given that a within-municipality comparison was not always recent, available, or from a different political party, a spatial comparison offered the most electorally-relevant contrast with a municipality's incumbent party.[15] In this example, the local information shows 28% while the benchmark shows 4%.

Figure 3 documents the distribution of malfeasant spending in our sample. Importantly, the average precinct was informed of 21% malfeasant spending within their municipality and 9% in municipalities within their state governed by other parties. As the figure illustrates, only 7 of our 26 municipalities learned that their incumbent's malfeasance fell below the average across incumbents from other parties audited within the same state. Although control group voters generally already viewed the incumbent as somewhat more malfeasant than challenger parties, as shown in Appendix

[13]These efforts to minimize perceived bias were largely successful, given that voters generally recognized the leaflets as non-partisan (see Appendix Table A11 (p. 20)).

[14]A few leaflets were delivered to voters outside the precinct, while adverse weather conditions and poor road conditions prevented us from reaching one precinct. We preserve the randomization by estimating intent to treat effects.

[15]We chose the statewide average because it should be more informative and less politically contentious than selecting a single municipality for comparison.

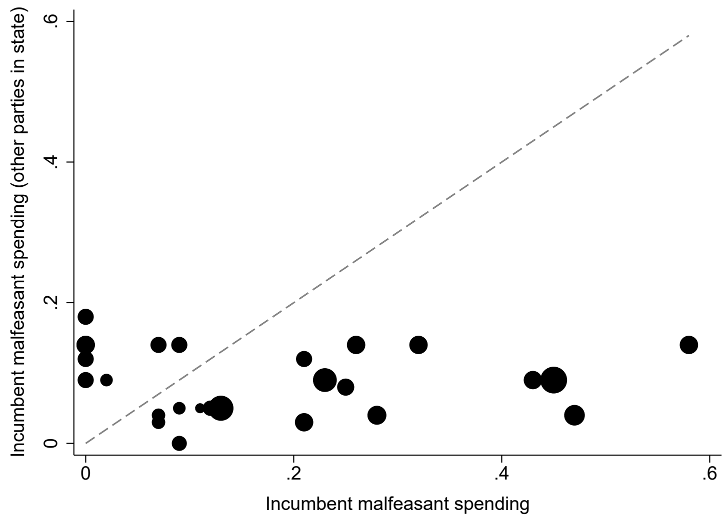

Figure 3: Precincts by share of malfeasant spending in our sample

Notes: The dotted line is the 45^o line. Each point is one of our 26 municipalities. The size of points corresponds to the number of precincts in our sample from that 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

Figure A3 (p. 8), these malfeasance indicators are likely to accentuate this difference.

Second, we varied whether the leaflet was delivered in a private or public manner. For the *public* mode of delivery, door-to-door delivery of the leaflet—our *private* mode of delivery—was augmented by a powerful portable loudspeaker carried on the back of a team member.[16] Akin to the vehicles commonly driving around before Mexican elections blaring campaign messages, a single *perifonista* walked through the streets of each precinct alongside other team members distributing leaflets while playing a 30-second message on loop. The message informed voters that their neighbors would also receive information concerning the malfeasance of their municipal mayor, and encouraged them to share and discuss the information provided.

[16]We purchased these modified rucksack loudspeakers from a vendor in Mexico City that also serves political campaigns similarly seeking to broadcast their message. See Appendix Figure A2 (p.8).

Treatments were randomly assigned within 100 blocks (stratified by rural/urban within a municipality) containing six or seven similar precincts according to the 2×2 factorial design with a pure control shown in Table 1.[17] Each block contained one precinct receiving each treatment condition, as well as two or three control precincts (depending on precinct availability). Block randomization ensures that all respondents within a block are subject to the same electoral race and receive the same information pertaining to their mayor, and can substantially increase statistical power. Appendix Table A2 (p. 11) demonstrates that pre-determined precinct- and individual-level covariates are well-balanced across treatment conditions.

4.3 Measurement of key variables

We examine two main classes of outcomes. First, we collected precinct-level electoral returns from state electoral institutes to measure incumbent party vote share as a proportion of registered voters.[18] Second, we surveyed ten voters per precinct in the weeks after the election in all treated precincts and in one control precinct from each block. We use this survey, which only visited households where leaflets were delivered, to measure posterior beliefs about incumbent and challenger party malfeasance and voter coordination.[19]

Although financial constraints prevented us from conducting a baseline survey, we measure the direction and extent of updating by showing respondents the leaflet at the end of our survey and eliciting beliefs about incumbent and challenger malfeasance before and after seeing the leaflet. We then construct a municipality-level measure of how voters updated their beliefs in response to the information, using the average change in beliefs upon seeing the leaflet among respondents in control precincts within the municipality. Using post-election surveys in this way requires: (1) that control group respondents are similar to treatment group respondents, (2) that control respon-

[17] Blocks were created using the R package `blockTools`, which sequentially creates the most similar blocks possible, based on 23 social, economic, demographic, and political variables.

[18] We obtain similar results using vote share as a proportion of turnout, but prefer the registered voters denominator because turnout could also be affected by treatment.

[19] In control precincts, enumerators were instructed to survey respondents using the same protocols that would have occurred had the precinct been treated.

dent beliefs are consistent across the month between the intervention and the post-election survey, and (3) that control group respondents internalized the information similarly to those in treated precincts. Appendix section A.5 (pp. 10-13) provides extensive support for these assumptions.

4.4 Estimation

The goal of this article is to identify how our experimental variation in the content of leaflet information and the mode of leaflet delivery affected voters. Accordingly, we estimate the following baseline regressions:

$$Y_{pbm} = \boldsymbol{\beta}\mathbf{T}_{pbm} + \eta_{bm} + \varepsilon_{pbm}, \tag{1}$$

where Y_{pbm} is an outcome for precinct p within block b of municipality m; we add an i subscript for individual-level survey responses. The coefficient vector $\boldsymbol{\beta}$ estimates the effect of our information treatment conditions relative to control precincts, while the block fixed effects η_{bm} adjust for the differential treatment probabilities across blocks arising from different block sizes. Since the *differential* effect of our treatment variants is the main focus of this article, we test for differences in treatment effects across treatments as well as drop the control group to explicitly estimate differences between comparative and local, or public and private, treatment conditions. Standard errors are clustered by treatment-municipality. We depart from our pre-analysis plan by weighting precinct-level observations by the share of voters within the precinct to whom we delivered a leaflet.[20] This is likely to increase precision by downweighting precincts where treatment delivery was more limited, although Appendix Tables A9-A10 (pp. 18-19) show that unweighted regressions produce similar results.

To examine how the effects of the treatments vary with reported levels of malfeasance and, most importantly, how voters updated their beliefs in response to the information and the content

[20] In control precincts, we use the share of leaflets delivered to the average treated precinct within their block.

of the information provided, we also estimate heterogeneous effects of the form:

$$Y_{pbm} = \beta \mathbf{T}_{pbm} + \gamma(\mathbf{T}_{pbm} \times X_{bm}) + \eta_{bm} + \varepsilon_{pbm}, \tag{2}$$

where X_{bm} is a block- or municipality-level moderator described below.

5 Results

We first demonstrate that the treatments were indeed received and internalized by voters, before turning to our main findings identifying the effects of comparative performance information and public dissemination.

5.1 Manipulation tests

To verify that our treatment reached voters and produced the intended effects, Table 2 reports several manipulation tests. Based on our post-treatment survey, columns (1)-(4) first confirm that voters assigned to receive any treatment were significantly more likely to remember the leaflet and correctly recall the issue discussed in the leaflet. Comparative and public treatments are slightly more effective in this regard, as shown by the coefficient equality tests at the foot of the table, although the differences are relatively small.

Column (5) shows that voters assigned to receive the comparative leaflets were significantly more likely than the local and control groups to recall receiving information about other parties in their state. Given the lack of cross-precinct spillovers documented in Appendix Section A.6 (pp. 13-15), the coefficient for precincts receiving only information about their own municipal government (local treatment) indicates fuzzy recall. Nevertheless, voters receiving the comparative information were two percentage points—or around 40%—more likely to recall receiving information about opposition incumbents in their state. The t test at the foot of the table indicates that this difference is statistically significant.

Table 2: Treatment manipulation checks

	Remember leaflet		Correctly remember content		Remember comparative content	Remember loud speaker	Share of community received
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Local treatment	0.234***		0.128***		0.046***		
	(0.023)		(0.019)		(0.010)		
Comparative treatment	0.260***		0.148***		0.066***		
	(0.023)		(0.020)		(0.010)		
Private treatment		0.231***		0.130***		0.006	0.483***
		(0.023)		(0.020)		(0.007)	(0.062)
Public treatment		0.263***		0.145***		0.057***	0.566***
		(0.024)		(0.020)		(0.008)	(0.064)
Outcome range	{0,1}	{0,1}	{0,1}	{0,1}	{0,1}	{0,1}	{1,2,3,4,5}
Control outcome mean	0.09	0.09	0.06	0.06	0.03	0.03	1.45
Control outcome std. dev.	0.28	0.28	0.25	0.25	0.18	0.16	1.01
Test: same treatment effect (p value)	0.08	0.06	0.04	0.19	0.03	0.00	0.06
Observations	4,958	4,958	4,958	4,958	4,958	4,958	4,929

Notes: All specifications include block fixed effects, and are estimated using OLS. Standard errors clustered by municipality-treatment are in parentheses. "Remember leaflet" is an indicator for whether the respondents recalls receiving the leaflet. "Correctly remember content is an indicator for respondents that correctly recall the issue discussed in the leaflet (i.e. unauthorized spending or not spending on the poor, from among four options). "Remembers comparative content" is an indicator for whether the respondent reports that the leaflet also included information on other parties in the state. "Remember loud speaker" is an indicator for whether the respondent recalled listening to a loudspeaker with a recording accompanying the leaflet distribution. "Share of community received" measures on a five-point scale the fraction of community members that the respondent believes received a leaflet—1=very few, 2=less than half,... and 5=almost everyone. * denotes $p < 0.1$, ** denotes $p < 0.05$, *** denotes $p < 0.01$.

The public dissemination treatment also elicited the expected responses. Column (6) shows that voters in precincts receiving the private information treatment were as likely as control respondents to recall a loudspeaker. However, voters in precincts subject to public dissemination were six percentage points more likely to correctly recall that leaflet delivery was accompanied by a loudspeaker. Moreover, the test at the bottom of column (7) indicates that voters in such precincts were also significantly more likely to believe that a large fraction of their community received the leaflets. This was a central message of the loudspeaker script and is a key mechanism through which we expected public dissemination could generate coordination.

5.2 The effect of providing comparative malfeasance information

To test H1 and H2, we first examine the differential effect of providing comparative performance information on voters' posterior beliefs about incumbent and challenger malfeasance. In both cases, voters were asked to rate party malfeasance (on the dimension about which they received information) on a five-point scale where high values represent high perceived malfeasance. While the incumbent party can always be matched to a particular party, the challenger is not always well defined. Accordingly, we consider three possible definitions of challenger: the party receiving the second largest vote share at the previous municipal election in 2012; a respondent's second most preferred party; and the average across whichever of the PAN, PRD, and PRI were not in office before the 2015 election. We focus on the first definition in the main text, and provide similar results using the second and third definitions in the Appendix.

Table 3 first explores how the provision of information induced voters to update their posterior beliefs about the incumbent party. Column (1) detects no statistically significant average treatment effect of either the local or comparative treatments, relative to the control group. Moreover, and contrary to H1, the t tests at the foot of panel A and the point estimates in panel B provide no evidence that comparative information induced more unfavorable updating than local information. In fact, our design is powered to precisely estimate this null effect: the 95% confidence interval—$(-0.077, 0.095)$—implied by the estimate in panel B indicates that we can reject even a 0.07

Table 3: Effect of local and comparative information treatments on voters' posterior beliefs about incumbent party malfeasance

	Perceived incumbent party malfeasance (very low - very high)		
	(1)	(2)	(3)
Panel A: Control group as baseline			
Local treatment	0.000	0.041	-0.090
	(0.045)	(0.095)	(0.056)
Comparative treatment	-0.003	-0.027	-0.094*
	(0.046)	(0.093)	(0.053)
Local treatment × Incumbent malfeasant spending		-0.051	
		(0.237)	
Comparative treatment × Incumbent malfeasant spending		-0.117	
		(0.234)	
Local treatment × Challenger malfeasant spending		-0.328	
		(0.892)	
Comparative treatment × Challenger malfeasant spending		0.557	
		(0.925)	
Local treatment × Unfavorable incumbent updating			0.131**
			(0.064)
Comparative treatment × Unfavorable incumbent updating			0.155**
			(0.071)
Local treatment × Unfavorable challenger updating			-0.040
			(0.073)
Comparative treatment × Unfavorable challenger updating			-0.071
			(0.087)
Control outcome mean	-0.14	-0.14	-0.14
Control outcome std. dev.	1.48	1.48	1.48
Interaction 1 mean		0.21	0.91
Interaction 1 std. dev.		0.17	1.00
Interaction 2 mean		0.09	0.71
Interaction 2 std. dev.		0.05	0.93
Test: same treatment effect (p value)	0.94	0.53	0.96
Test: same interaction (1) effect (p value)		0.76	0.66
Test: same interaction (2) effect (p value)		0.37	0.62
Observations	4,624	4,624	4,624
Panel B: Local treatment group as baseline			
Comparative treatment	0.009	-0.071	0.002
	(0.044)	(0.110)	(0.061)
Comparative treatment × Incumbent malfeasant spending		-0.053	
		(0.214)	
Comparative treatment × Challenger malfeasant spending		1.023	
		(1.002)	
Comparative treatment × Unfavorable incumbent updating			0.017
			(0.061)
Comparative treatment × Unfavorable challenger updating			-0.014
			(0.069)
Local treatment outcome mean	-0.10	-0.10	-0.10
Local treatment outcome std. dev.	1.49	1.49	1.49
Interaction 1 mean		0.21	0.90
Interaction 1 std. dev.		0.17	1.01
Interaction 2 mean		0.09	0.68
Interaction 2 std. dev.		0.04	0.92
Observations	3,555	3,555	3,555
Outcome range	{-2,-1,0,1,2}	{-2,-1,0,1,2}	{-2,-1,0,1,2}

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 $p < 0.01$.

standard deviation increase in perceived incumbent malfeasance due to the differential effect of providing a cross-sectional performance benchmark. This is surprising given that voters assigned to this treatment learned of relatively low malfeasance levels for challenger parties, and were thus likely to attribute the often high levels of reported incumbent malfeasance to the incumbent party's particularly high underlying level of malfeasance, rather than to a common shock affecting all parties similarly.

We further find little evidence that the effects of providing comparative performance information vary with either the benchmarked component of the information's content or voter belief updating about challengers. While the estimates in column (2) show that neither the effect of local nor comparative information significantly varied by the amount of malfeasance reported for the incumbent and challenger parties, column (3) shows that—consistent with our Bayesian framework—treated voters in municipalities that updated most unfavorably about the incumbent from the information provided were significantly more likely to believe that incumbent parties were more malfeasant.[21] However, contrary to H2, the extent of updating about the challenger—which allows voters to more accurately filter out common shocks—did not significantly moderate the influence of the comparative treatment or differentially do so relative to the local information treatment. Indeed, while the interaction coefficient in panel B is negative (consistent with our theory), it is negligible and far from being statistically significant. In particular, the 95% confidence interval implied by the interaction with respect to unfavorable challengers updating in column (3) of panel B does not contain standardized differential effects of more than 0.1 standard deviations.

We find slightly stronger support for hypotheses H1 and H2 in the case of challenger parties. Consistent with H1, the negative coefficient in column (1) of panel B in Table 4 indicates that the comparative treatment induced more favorable updating about the challenger, relative to incumbent-only information. However, the point estimate is statistically insignificant and represents less than 0.05 standard deviations of the outcome; the 95% confidence interval—$(-0.102, 0.016)$—does not include negative effects exceeding 0.11 standard deviations of the chal-

[21] This result is not mechanical, since municipality-level measures of unfavorable updating are based only on responses from voters in control precincts upon receiving the leaflet.

Table 4: Effect of local and comparative information treatments on voters' posterior beliefs about challenger party malfeasance, where the challenger is the party that received the second-largest vote share in the last municipal election

	Perceived challenger party malfeasance (very low - very high)		
	(1)	(2)	(3)
Panel A: Control group as baseline			
Local treatment	0.027	-0.064	-0.011
	(0.042)	(0.103)	(0.044)
Comparative treatment	-0.042	-0.029	-0.080*
	(0.043)	(0.091)	(0.042)
Local treatment × Incumbent malfeasant spending		0.217	
		(0.242)	
Comparative treatment × Incumbent malfeasant spending		-0.173	
		(0.245)	
Local treatment × Challenger malfeasant spending		0.488	
		(0.962)	
Comparative treatment × Challenger malfeasant spending		0.265	
		(1.007)	
Local treatment × Unfavorable incumbent updating			0.014
			(0.075)
Comparative treatment × Unfavorable incumbent updating			-0.051
			(0.074)
Local treatment × Unfavorable challenger updating			0.031
			(0.079)
Comparative treatment × Unfavorable challenger updating			0.118
			(0.087)
Control outcome mean	-0.30	-0.30	-0.30
Control outcome std. dev.	1.36	1.36	1.36
Test: same treatment effect (p value)	0.08	0.68	0.09
Test: same interaction (1) effect (p value)		0.78	0.14
Test: same interaction (2) effect (p value)		0.07	0.29
Observations	4,958	4,958	4,958
Panel B: Local treatment group as baseline			
Comparative treatment	-0.059	0.044	-0.064
	(0.041)	(0.085)	(0.040)
Comparative treatment × Incumbent malfeasant spending		-0.427*	
		(0.227)	
Comparative treatment × Challenger malfeasant spending		-0.131	
		(0.872)	
Comparative treatment × Unfavorable incumbent updating			-0.084
			(0.065)
Comparative treatment × Unfavorable challenger updating			0.117
			(0.069)
Local treatment outcome mean	-0.22	-0.22	-0.22
Local treatment outcome std. dev.	1.36	1.36	1.36
Observations	3,819	3,819	3,819
Outcome range	{-2,-1,0,1,2}	{-2,-1,0,1,2}	{-2,-1,0,1,2}

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 $p < 0.01$.

lenger posterior belief. Consistent with Bayesian learning, the interaction of the comparative treatment with our measure of unfavorable challenger updating is positive, though also not statistically significant. A similar pattern emerges regarding H2: the interactions between the comparative treatment and incumbent malfeasant spending (column (2) in panel B) and incumbent unfavorable updating (column (3) in panel B) are also negative, though the latter is not statistically significant. Appendix Tables A7 and A8 (pp. 16-17) report similar patterns across our other definitions of the challenger. In sum, while the signs of the coefficients are consistent with the predictions in hypotheses H1 and H2, the estimates are generally small and statistically insignificant. These results suggest that any effect of benchmarked performance information on posterior beliefs about challengers was limited.

We next examine the effects of providing comparative information on incumbent vote share as a share of registered voters. H3 and H4 hypothesized that providing voters with a benchmark—especially one that contrasts the incumbent party's malfeasance with that of challenger parties in office elsewhere in the state (see Figure 3)—would elicit stronger sanctioning of the incumbent. However, perhaps unsurprisingly in the light of the preceding results with respect to beliefs about malfeasance, the findings in Table 5 indicate that comparative malfeasance information did not have a differential effect on electoral outcomes.

In particular, column (1) first shows that comparative information did not differentially affect voting behavior on average. The positive treatment effect on incumbent vote share, relative to the control group, may initially seem surprising. However, as noted above, Arias et al. (2018) show that while the malfeasance reports did not affect the level of posterior beliefs *on average*, the reports did reduce voter uncertainty about incumbent party malfeasance and elicited responses from incumbent and challenger parties that were likely to have differentially benefited the incumbent party and increased its vote share. More importantly for this article's focus on benchmarking incumbent performance, the coefficient in column (1) of panel B demonstrates that the effects of local and comparative performance information are statistically indistinguishable. Furthermore, this null effect is fairly precisely estimated: the 95% confidence interval—$(-0.009, 0.019)$—does

Table 5: Effect of local and comparative information treatments on incumbent party vote share

	Incumbent party vote share (share of registered voters)		
	(1)	(2)	(3)
Panel A: Control group as baseline			
Local treatment	0.012***	0.031***	0.022***
	(0.004)	(0.010)	(0.004)
Comparative treatment	0.015***	0.019*	0.015***
	(0.005)	(0.011)	(0.005)
Local treatment × Incumbent malfeasant spending		-0.051**	
		(0.025)	
Comparative treatment × Incumbent malfeasant spending		-0.036*	
		(0.021)	
Local treatment × Challenger malfeasant spending		-0.096	
		(0.080)	
Comparative treatment × Challenger malfeasant spending		0.044	
		(0.088)	
Local treatment × Unfavorable incumbent updating			-0.012*
			(0.006)
Comparative treatment × Unfavorable incumbent updating			-0.009
			(0.006)
Local treatment × Unfavorable challenger updating			0.000
			(0.006)
Comparative treatment × Unfavorable challenger updating			0.008
			(0.007)
Control outcome mean	0.19	0.19	0.19
Control outcome std. dev.	0.07	0.07	0.07
Interaction 1 mean		0.21	0.91
Interaction 1 std. dev.		0.17	1.00
Interaction 2 mean		0.09	0.71
Interaction 2 std. dev.		0.04	0.95
Test: same treatment effect (p value)	0.61	0.42	0.31
Test: same interaction (1) effect (p value)		0.68	0.81
Test: same interaction (2) effect (p value)		0.26	0.42
Observations	675	675	651
Panel B: Local treatment group as baseline			
Comparative treatment	0.005	-0.010	-0.004
	(0.007)	(0.017)	(0.007)
Comparative treatment × Incumbent malfeasant spending		0.007	
		(0.038)	
Comparative treatment × Challenger malfeasant spending		0.142	
		(0.136)	
Comparative treatment × Unfavorable incumbent updating			-0.000
			(0.010)
Comparative treatment × Unfavorable challenger updating			0.010
			(0.010)
Local treatment outcome mean	0.21	0.21	0.21
Local treatment outcome std. dev.	0.08	0.08	0.08
Interaction 1 mean		0.22	0.92
Interaction 1 std. dev.		0.17	1.00
Interaction 2 mean		0.09	0.72
Interaction 2 std. dev.		0.04	0.94
Observations	398	398	382
Outcome range	[0.03,0.40]	[0.03,0.40]	[0.03,0.40]

Notes: All specifications include block fixed effects, weight 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 column (3) reflects 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 $p < 0.01$.

not include negative effects exceeding 0.11 standard deviations of the incumbent party vote share. In contrast with hypothesis H3, which was based on the substantially lower malfeasance reported about the average incumbent from a different party, this suggests that both types of information affected voter behavior similarly, on average.

Consistent with Bayesian learning, columns (2) and (3) further show that the electoral reward for the incumbent is lower in municipalities in which greater incumbent malfeasance was reported, and in which voters updated more unfavorably about the incumbent based on the information reported. However, we find no evidence of a *differential* impact of providing comparative information. In particular, and contrary to H4, column (3) reports no differential effect of our comparative treatment based on the extent of unfavorable updating about the challenger. While the point estimate for the interaction term is positive (as predicted), the effect is relatively small and statistically insignificant.[22] The interaction terms in panel B indicate that, for standard deviation increases in challenger malfeasant spending and unfavorable challenger updating respectively, the 95% confidence intervals imply that the differential effect of providing comparative performance information does not contain incumbent vote share increases exceeding 1.6 and 2.8 percentage points.

The lack of a meaningful differential effect associated with providing comparative information could reflect several possibilities. First, since voters in the control group already believed the main local challenger to be less malfeasant, information about challenger parties may have already corresponded with voters' prior beliefs. Second, voters may have simply failed to comprehend the benchmark component of the treatment. Third, the benchmark itself may not have been relevant. For example, Marshall (2018) finds that Mexican voters benchmark local homicides against prior incumbent parties from the same municipality, but do not benchmark their incumbent's performance against neighboring municipalities.

While it is difficult to disentangle the reasons behind null findings (Lieberman, Posner and Tsai 2014), our survey data can help separate between these explanations by examining whether the comparative treatments differentially affected voter beliefs. First, the estimates in Table 3 show no

[22]Unreported estimates for turnout also suggest that local and comparative information impacted turnout similarly.

differential updating about the incumbent party's malfeasance across local and comparative treatments. This indicates that voters understood the information provided about the incumbent, but—consistent with a learning model where performance elsewhere is orthogonal to performance in a voter's municipality or where challenger information already conformed with prior expectations—adding a benchmark did not adjust how voters' updated about their incumbent party.

Second, the limited updating about challengers shown in Table 4 indicates that voters primarily updated from the information provided about the incumbent party. This suggests that voters either struggled to comprehend the comparative component of the information, which some of our enumerators highlighted as they conducted the survey, or voters did not believe that the malfeasance of parties in other municipalities represents a good proxy for how such parties would perform in their municipality. In contrast, the fact that voters did not differentially update about the challenger when information differed from prior beliefs suggests that malfeasance indicators conforming with prior expectations does not explain this intervention's limited effects. It is thus possible that, had they received what they considered more appropriate comparative information, voters would have updated differentially about the incumbent and consequently changed their voting behavior. Unfortunately, we cannot distinguish between these potential explanations, and thus assess the potential effect of such a treatment variant.

5.3 Limited amplifying effect of public information dissemination

While the effects of our basic incumbent performance information are not altered by further providing comparative performance information, it is possible that public dissemination may more effectively stimulate voter responses. To investigate this, we examine whether the effects of information on voting behavior—that revealing incumbent party malfeasance is most likely to increase incumbent party vote share where ASF reports reveal lower levels of malfeasance and where voters updated most favorably about the incumbent from the information provided—are amplified when the provision of performance information is accompanied by a loudspeaker announcing its dissemination.

Table 6: Effect of private and public information treatments on incumbent party vote share

	Incumbent party vote share (share of registered voters)		
	(1)	(2)	(3)
Panel A: Control group as baseline			
Private treatment	0.020***	0.041***	0.023***
	(0.005)	(0.013)	(0.005)
Public treatment	0.007	0.009	0.013***
	(0.004)	(0.014)	(0.004)
Private treatment × Incumbent malfeasant spending		-0.065***	
		(0.021)	
Public treatment × Incumbent malfeasant spending		-0.021	
		(0.022)	
Private treatment × Challenger malfeasant spending		-0.079	
		(0.103)	
Public treatment × Challenger malfeasant spending		0.028	
		(0.096)	
Private treatment × Unfavorable incumbent updating			-0.015**
			(0.006)
Public treatment × Unfavorable incumbent updating			-0.005
			(0.004)
Private treatment × Unfavorable challenger updating			0.011
			(0.007)
Public treatment × Unfavorable challenger updating			-0.003
			(0.005)
Control outcome mean	0.19	0.19	0.19
Control outcome std. dev.	0.07	0.07	0.07
Interaction 1 mean		0.21	0.91
Interaction 1 std. dev.		0.17	1.00
Interaction 2 mean		0.09	0.71
Interaction 2 std. dev.		0.04	0.95
Test: same treatment effect (p value)	0.09	0.18	0.22
Test: same interaction (1) effect (p value)		0.18	0.15
Test: same interaction (2) effect (p value)		0.52	0.09
Observations	675	675	651
Panel B: Private treatment group as baseline			
Public treatment	-0.012	-0.027	-0.009
	(0.008)	(0.025)	(0.009)
Public treatment × Incumbent malfeasant spending		0.040	
		(0.034)	
Public treatment × Challenger malfeasant spending		0.059	
		(0.175)	
Public treatment × Unfavorable incumbent updating			0.010
			(0.007)
Public treatment × Unfavorable challenger updating			-0.016*
			(0.008)
Private treatment outcome mean	0.21	0.21	0.21
Private treatment outcome std. dev.	0.08	0.08	0.08
Interaction 1 mean		0.22	0.92
Interaction 1 std. dev.		0.17	1.00
Interaction 2 mean		0.09	0.72
Interaction 2 std. dev.		0.04	0.94
Observations	398	398	382
Outcome range	[0.03,0.40]	[0.03,0.40]	[0.03,0.40]

Notes: All specifications include block fixed effects, weight 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 Column (3) 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 $p < 0.01$.

However, despite being more likely to recall a loud speaker and believe that a large fraction of the community received the leaflets (see Table 2), public dissemination produced similar—if not weaker—responses from voters. Column (1) of Table 6 reports a smaller increase in incumbent vote share associated with public dissemination, relative to private treatment dissemination. In contrast with the expectation that public dissemination would amplify the positive average treatment effect, the 95% confidence implied by the estimate in column (1) of panel B—$(-0.028, 0.004)$—excludes an increase in incumbent vote share of 0.4 percentage points (or 0.05 standard deviations) or more. While this could reflect a differentially less sanguine response to generally high levels of malfeasance, columns (2) and (3) also document no significant difference in the differential slopes with respect to the level of incumbent and challenger reported malfeasance and belief updating. In sum, these results provide little evidence to support hypothesis H5.

The limited effect of adding a loudspeaker contrasts with the large effects of information disseminated by the media in similar contexts (Banerjee et al. 2011; Ferraz and Finan 2008; Larreguy, Marshall and Snyder 2018; Marshall 2018). One potential explanation for the limited voter response is the greater capacity of broadcast media to foster either explicit or tacit coordination through common knowledge (Adena et al. 2015; Yanagizawa-Drott 2014). Indeed, even though—as intended—the public treatment generated greater common knowledge, we find little evidence that this translated into greater voter coordination: columns (1)-(3) in Table 7 report no significant increase in discussion of the leaflet, vote coordination on the basis of the leaflet, or changes in voting behavior on the basis of discussions of the leaflet between the public and private forms of information dissemination.

Another possibility is that the leaflet became common knowledge even with only private dissemination. There is certainly some evidence that the leaflet itself induced significant coordination: the estimates in Table 7 show that both the private and public leaflet increased discussion and coordination. However, the varied and relatively small scale of such responses, as well as the relatively low recall rates in Table 2, suggest that it is very unlikely that private dissemination had already produced maximal common knowledge.

Table 7: Effect of variants of information treatment on social transmission

	Social discussion of leaflet (1)	Discussion created vote coordination (2)	Discussion of leaflet changed vote (3)
Private information treatment	0.111***	0.022***	0.028***
	(0.015)	(0.008)	(0.007)
Public information treatment	0.125***	0.030***	0.030***
	(0.014)	(0.008)	(0.008)
Outcome range	{0,1}	{0,1}	{0,1}
Control outcome mean	0.05	0.02	0.02
Control outcome std. dev.	0.23	0.13	0.12
Test: same treatment effect (p value)	0.22	0.26	0.82
Observations	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. "Social discussion of leaflet" is an indicator for respondents that reported discussing the contents of the leaflets with other members of the community. "Discussion created vote coordination" is an indicator for people in the community coordinating to vote for the same party as result of discussing the leaflet. "Discussion of leaflet changed vote" captures that discussions about the leaflet with other members of the community changed the respondent's vote choice. * denotes $p < 0.1$, ** denotes $p < 0.05$, *** denotes $p < 0.01$.

A third possibility is that the loudspeaker led voters to perceive our intervention as being more partisan. Since political parties frequently use these loudspeakers as part of their campaigns, this could have led respondents to discount the information in the leaflets, which could explain the somewhat weaker effects. However, Appendix Table A11 (p. 20) reports no evidence that the public treatment increased voter perceptions that the leaflet was delivered by the municipal incumbent, municipal challengers, or the PAN, PRD, or PRI parties.

6 Conclusion

This article examines how comparative performance information and public dissemination moderate the effect of an NGO information campaign in the context of Mexican municipal elections. Leveraging a large-scale field experiment varying the provision of performance benchmarks and public dissemination by loudspeaker, we find little evidence of a differential effect on either voter belief updating or voting behavior. First, while voters were significantly more likely to recall receiving information about other parties in their state than control voters, treated voters did not differentially update their beliefs about the incumbent. Second, while the loudspeaker differentially increased common knowledge that the leaflets were delivered among treated voters, these voters were not more likely to coordinate their behavior around the treatment information. In light of the widely-recognized problem of publication bias, we believe that the null findings of our well-powered field experiment have important implications for understanding political behavior, future research, and campaign design.

First, the provision of a cross-sectional benchmark could reduce voter comprehension of the information and might prove irrelevant to voters. It is then essential for information campaigns to provide benchmarks that are both easy to comprehend and relevant to voters. To that end, it is important to start by eliciting which comparative information voters deem relevant to assessing the relative performance of their incumbents. While only contemporaneous information regarding mayors from other parties within the state was available in our case due to the infrequency of ASF audits, voters might regard other comparative information, e.g. from previous incumbents in their municipalities, as more relevant.[23] Additionally, extensively piloting is of great relevance to determining the most effective way of depicting comparisons.

Second, public dissemination through devices such as loudspeakers might be insufficient to produce additional coordination beyond the small levels of coordination induced by leaflets. To the extent that the large effects of the media on political outcomes (e.g. Adena et al. 2015; DellaVigna and Kaplan 2007; Enikolopov, Petrova and Zhuravskaya 2011; Larreguy, Marshall and Snyder 2018; Marshall 2018; Snyder and Strömberg 2010; Yanagizawa-Drott 2014) reflect coordination, our findings indicate that a loudspeaker cannot achieve this. However, further research is still needed to understand how to get a critical mass of voters to coordinate around the treatment information. Future research could thus also examine the coordination potential of candidate meetings (Bidwell, Casey and Glennerster, 2018), public meetings, social media, or large-scale messaging

[23]The ASF had not recently audited many of the municipalities in our sample, which prevented us from using such comparative information.

using modern communication technologies. In doing so, researchers might usefully assess whether making common knowledge that the treatment information was delivered, e.g. via directly communicating the extent to which other voters also got the information, is a more effective way of inducing coordination. Our findings nevertheless demonstrate that leaflets can be effective, underscoring the importance of clear information about incumbent performance in office.

References

Adena, Maja, Ruben Enikolopov, Maria Petrova, Veronica Santarosa and Ekaterina Zhuravskaya. 2015. "Radio and the Rise of The Nazis in Prewar Germany." *Quarterly Journal of Economics* 130(4):1885–1939.

Alatas, Vivi, Abhijit Banerjee, Arun G Chandrasekhar, Rema Hanna and Benjamin A Olken. 2016. "Network Structure and the Aggregation of Information: Theory and Evidence from Indonesia." *American Economic Review* 106(7):1663–1704.

Arias, Eric. forthcoming. "How Does Media Influence Social Norms? A Field Experiment on the Role of Common Knowledge." *Political Science Research and Methods* .

Arias, Eric, Horacio A. Larreguy, John Marshall and Pablo Querubín. 2018. "Priors rule: When do malfeasance revelations help or hurt incumbent parties?" Working paper.
URL: *https://goo.gl/rq9upp*

Auditoría Superior de la Federación. 2014. "Informe del Resultado de la Fiscalización Superior de la Cuenta Pública 2012." Audit Summary Report.
URL: *https://goo.gl/fR5nmt*

Aytaç, Selim Erdem. 2018. "Relative Economic Performance and the Incumbent Vote: A Reference Point Theory." *Journal of Politics* 80(1):16–29.

Banerjee, Abhijit V., Selvan Kumar, Rohini Pande and Felix Su. 2011. "Do Informed Voters Make Better Choices? Experimental Evidence from Urban India." Working paper.
URL: *goo.gl/Xps2hn*

Barro, Robert J. 1973. "The Control of Politicians: An Economic Model." *Public Choice* 14(1):19–42.

Besley, Timothy and Anne Case. 1995. "Incumbent Behavior: Vote-Seeking, Tax-Setting, and Yardstick Competition." *American Economic Review* 85(1):25–45.

Besley, Timothy J. 2006. *Principled Agents? The Political Economy of Good Government*. Oxford, UK: The Lindahl Lectures, Oxford University Press.

Bidwell, Kelly, Katherine Casey and Rachel Glennerster. 2018. "Debates: Voting and Expenditure Responses to Political Communication." Working Paper.
URL: *goo.gl/PRHxfh*

Bishop, Christopher M. 2006. *Pattern recognition and machine learning*. Springer.

Boas, Taylor, F. Daniel Hidalgo and Marcus A. Melo. forthcoming. "Norms versus Action: Why Voters Fail to Sanction Malfeasance in Brazil." *American Journal of Political Science* .

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.

Chwe, Michael Suk-Young. 2000. "Communication and coordination in social networks." *Review of Economic Studies* 67(1):1–16.

de Figueiredo, Miguel F.P., F. Daniel Hidalgo and Yuri Kasahara. 2014. "When Do Voters Punish Corrupt Politicians? Experimental Evidence from Brazil." Working paper.
URL: *goo.gl/rWXBbm*

De La O, Ana L. and Fernando Martel García. 2015. "Can Intrastate Accountability Reduce Local Capture? Results from a Field Experiment in Mexico." Working Paper.
URL: *goo.gl/UgxDio*

DellaVigna, Stefano and Ethan Kaplan. 2007. "The Fox News effect: Media bias and voting." *Quarterly Journal of Economics* 122(3):1187–1234.

Dunning, Thad, Guy Grossman, Macartan Humphreys, Susan Hyde and Craig McIntosh. forthcoming. *Metaketa I: The Limits of Electoral Accountability*. Cambridge University Press.

Enikolopov, Ruben, Aleksey Makarin and Maria Petrova. 2016. "Social Media and Protest Participation: Evidence from Russia." Working Paper.
URL: *goo.gl/dbTL9F*

Enikolopov, Ruben, Maria Petrova and Ekaterina Zhuravskaya. 2011. "Media and political persuasion: Evidence from Russia." *American Economic Review* 101(7):3253–3285.

Fearon, James D. 1999. Electoral Accountability and the Control of Politicians: Selecting Good Types versus Sanctioning Poor Performance. In *Democracy, Accountability, and Representation*, ed. 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.

Holmstrom, Bengt. 1982. "Moral hazard in teams." *Bell Journal of Economics* pp. 324–340.

Humphreys, Macartan and Jeremy Weinstein. 2012. "Policing Politicians: Citizen Empowerment and Political Accountability in Uganda Preliminary Analysis." Working paper.
URL: *goo.gl/EiVitc*

Jackson, Matthew O. 2010. *Social and Economic Networks*. Princeton University Press.

Kayser, Mark Andreas and Michael Peress. 2012. "Benchmarking across Borders: Electoral Accountability and the Necessity of Comparison." *American Political Science Review* 106(3):661684.

Keefer, Phillip. 2007. Seeing and Believing: Political Obstacles to Better Service Delivery. In *The Politics of Service Delivery in Democracies: Better Access for the Poor*, ed. Shantayanan Devarajan and Ingrid Widlund. Stockholm: Expert Group on Development Issues pp. 42–55.

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." World Bank Policy Research Report 106337.
URL: *https://goo.gl/5n58cG*

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 A., John Marshall and Jr. Snyder, James M. 2018. "Publicizing Malfeasance: How Local Media Faciliates Electoral Sanctioning of Mayors in Mexico." Working paper.
URL: *goo.gl/mecCDd*

Larreguy, Horacio A., John Marshall and Jr. Snyder, James M. forthcoming. "Leveling the playing field: How equalizing access to campaign advertising helps locally non-dominant parties in consolidating democracies." *Journal of the European Economic Association* .

Larson, Jennifer M., Jonathan Nagler, Jonathan Ronen and Joshua A. Tucker. 2017. "Social Networks and Protest Participation: Evidence from 130 Million Twitter Users." Working paper.
URL: *https://goo.gl/YPTmtV*

Lieberman, Evan S., Daniel N. Posner and Lily L. Tsai. 2014. "Does information lead to more active citizenship? Evidence from an education intervention in rural Kenya." *World Development* 60:69–83.

Lohmann, Susanne. 1993. "A Signaling Model of Informative and Manipulative Political Action." *American Political Science Review* 87(2):319–333.

Manacorda, Marco and Andrea Tesei. 2016. "Liberation technology: mobile phones and political mobilization in Africa." Working paper.
URL: *goo.gl/xAy1zX*

Manin, Bernard, Adam Przeworski and Susan C. Stokes. 1999. Elections and Representation. In *Democracy, Accountability, and Representation*, ed. Bernard Manin, Adam Przeworski and Susan C. Stokes. Cambridge, MA: Cambridge University Press pp. 29–54.

Marshall, John. 2018. "Political information cycles: When do voters sanction incumbent parties for high homicide rates?" Working paper.
URL: *goo.gl/6uFYNR*

Meyer, Margaret A. and John Vickers. 1997. "Performance comparisons and dynamic incentives." *Journal of Political Economy* 105(3):547–581.

Morris, Stephen and Hyun Song Shin. 2002. "Social value of public information." *American Economic Review* 92(5):1521–1534.

Nickerson, David W. 2008. "Is voting contagious? Evidence from two field experiments." *American Political Science Review* 102(1):49–57.

Pande, Rohini. 2011. "Can informed voters enforce better governance? Experiments in low-income democracies." *Annual Review of Economics* 3(1):215–237.

Rogoff, Kenneth. 1990. "Equilibrium Political Budget Cycles." *American Economic Review* 80(1):21–36.

Snyder, James M., Jr. and David Strömberg. 2010. "Press Coverage and Political Accountability." *Journal of Political Economy* 118(2):355–408.

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*, ed. The World Bank. Mexico City: The World Bank pp. 167–222.

Yanagizawa-Drott, David. 2014. "Propaganda and Conflict: Evidence from the Rwandan Genocide." *Quarterly Journal of Economics* 129(4):1947–1994.

A Supporting Information for "Does the content and mode of delivery of information matter for electoral accountability? Evidence from a field experiment in Mexico"

Contents

A.1 Formal derivation of the effect of providing incumbent-only performance information and comparative performance information

To formally establish the basis for our hypotheses regarding the effects of providing comparative performance information, we compare the Bayesian inferences that voters draw from incumbent-only and benchmarked signals using a Normal learning framework. In our model, a voter seeks to learn about the unobservable malfeasance of incumbent party I, m_I, and the unobservable malfeasance of "neighboring" incumbent party/parties N, m_N, in the presence of an unobserved common shock m_C that equally impacts the observable performance of both I and N. We assume that our representative voter's prior beliefs over these quantities are given by $N(\theta_I, 1/p_I)$, $N(\theta_N, 1/p_N)$, and $N(\theta_C, 1/p_C)$ respectively. For simplicity, we assume that these prior beliefs are independently distributed.

A.1.1 Incumbent-only malfeasance information

We first consider the simpler case where a voter receives any given realization of the incumbent-only malfeasance signal, $\hat{s}_I$, drawn from signal distribution $N(m_I + m_C, 1/\rho_I)$, where the signal's precision ρ_I is known. This indicator is a noisy signal of the combined effects of the underlying malfeasance of the incumbent party in a voter's own municipality and the common shock. The following proposition establishes voters' posterior inferences about I's malfeasance and the common shock:

Proposition 1. *(Incumbent-only performance information) Upon receiving realized signal $\hat{s}_I$, a voter's posterior expectation of incumbent party I's malfeasance is $w_I(\hat{s}_I - \theta_C) + (1 - w_I)\theta_I$ and of the common shock is $w_C(\hat{s}_I - \theta_I) + (1 - w_C)\theta_C$, where w_I and w_C are weights (defined within the proof) that both increase with ρ_I and respectively increase in p_C and p_I.*

Proof: We first define $\mathbf{m} = [m_I, m_C]'$, $\boldsymbol{\mu} = [\theta_I, \theta_C]'$, $\Lambda^{-1} = \begin{bmatrix} 1/p_I & 0 \\ 0 & 1/p_C \end{bmatrix}$, $\mathbf{A} = [1,1]$, and $\mathbf{L}^{-1} = [1/\rho_I]$. Applying a standard multivariate updating result (e.g. Bishop 2006:93) implies that posterior beliefs are distributed according to:

$$\mathbf{m}|\hat{s}_I \sim N\bigg((\Lambda + \mathbf{A}'\mathbf{L}\mathbf{A})^{-1}(\mathbf{A}'\mathbf{L}\hat{s}_I + \Lambda\boldsymbol{\mu}), (\Lambda + \mathbf{A}'\mathbf{L}\mathbf{A})^{-1} \bigg), \tag{A1}$$

where the application of matrix operations to the model in hand implies:

$$\begin{aligned} (\Lambda + \mathbf{A}'\mathbf{L}\mathbf{A})^{-1} &= \begin{bmatrix} p_I + \rho_I & \rho_I \\ \rho_I & p_C + \rho_I \end{bmatrix}^{-1} \\ &= \frac{1}{p_I p_C + p_I\rho_I + p_C\rho_I} \begin{bmatrix} p_C + \rho_I & -\rho_I \\ -\rho_I & p_I + \rho_I \end{bmatrix} := \Sigma, \qquad \text{(A2)} \\ (\mathbf{A}'\mathbf{L}\hat{s}_I + \Lambda\boldsymbol{\mu}) &= \begin{bmatrix} \rho_I\hat{s}_I + p_I\theta_I \\ \rho_I\hat{s}_I + p_C\theta_C \end{bmatrix}. \qquad \text{(A3)} \end{aligned}$$

Combining these results yields probability distribution:

$$p(\mathbf{m}|\hat{s}_I) \sim N\left(\begin{bmatrix} w_I(\hat{s}_I - \theta_C) + (1-w_I)\theta_I \\ w_C(\hat{s}_I - \theta_I) + (1-w_C)\theta_C \end{bmatrix}, \Sigma\right), \tag{A4}$$

where $w_I := \frac{p_C \rho_I}{p_I p_C + p_I \rho_I + p_C \rho_I}$ and $w_C := \frac{p_I \rho_I}{p_I p_C + p_I \rho_I + p_C \rho_I}$. ■

This result shows that incumbent performance information influences voter beliefs to the extent that the common shock-adjusted signal $(\hat{s}_I - \theta_C)$ differs from the voter's prior belief θ_I. Since the common shock is also uncertain, voters have limited capacity to update about the value of this shock, and thus rely on their prior belief θ_C. Relative to receiving no information about incumbent performance, and thus retaining the prior belief θ_I, a voter upwardly (downwardly) updates their expectation of I's malfeasance when $\theta_I < (>) w_I(\hat{s}_I - \theta_C) + (1-w_I)\theta_I \iff \theta_I < (>) \hat{s}_I - \theta_C$. Intuitively, this implies that, after netting out prior expectations of the common shock, voters update unfavorably about the incumbent party when the signal exceeds the voter's prior expectation. The same expression pertains to evaluating the posterior belief regarding the expected *difference* in I's malfeasance (or, more generally, "quality") relative to N's malfeasance—a common assumption in models of vote choice, which seems appropriate to our model to the extent that incumbent parties in other municipalities within the same state approximate challenger parties within our voter's own municipality.

A.1.2 Benchmarked malfeasance information

We now consider the more demanding case where a voter receives benchmarked signal, $\hat{s}_N$, *in addition to* the incumbent malfeasance signal $\hat{s}_I$ previously analyzed. We similarly assume that $\hat{s}_N$ is drawn from signal distribution $N(m_N + m_C, 1/\rho_N)$, where the signal's precision ρ_N is also known. The presence of a second signal enables a voter to draw more precise inferences by filtering out more precisely estimated common shocks, as well as learn more about the performance of incumbent parties in other municipalities that—to the extent that parties are believed to be correlated across municipalities—is informative about local challenger parties.

Extending our first proposition, our main proposition establishes voters' posterior beliefs following the provision of comparative performance information:

Proposition 2. *(Comparative performance information) Upon receiving realized signals $\hat{s}_I$ and $\hat{s}_N$, a voter's posterior expectation of incumbent party I's malfeasance is $w_{I,s}\hat{s}_I - w_{I,C}\theta_C - w_{I,\Delta}(\hat{s}_N - \theta_N) + w_{I,I}\theta_I$, of the neighboring incumbent party N's malfeasance is $w_{N,s}\hat{s}_N - w_{N,C}\theta_C - w_{N,\Delta}(\hat{s}_I - \theta_I) + w_{N,N}\theta_N$, and of the common shock is $w_{C,I}(\hat{s}_I - \theta_I) + w_{C,N}(\hat{s}_N - \theta_N) + w_{C,C}\theta_C$, where the weights are defined within the proof.*

Proof: We first define $\hat{\mathbf{s}} = [\hat{s}_I, \hat{s}_N]'$, $\mathbf{m} = [m_I, m_N, m_C]'$, $\boldsymbol{\mu} = [\theta_I, \theta_N, \theta_C]'$, $\Lambda^{-1} = \begin{bmatrix} 1/p_I & 0 & 0 \\ 0 & 1/p_C & 0 \\ 0 & 0 & 1/p_C \end{bmatrix}$, $\mathbf{A} = \begin{bmatrix} 1 & 0 & 1 \\ 0 & 1 & 1 \end{bmatrix}$, and $\mathbf{L}^{-1} = \begin{bmatrix} 1/\rho_I & 0 \\ 0 & 1/\rho_N \end{bmatrix}$. We then apply the same theorem as in the previous proof, where the application of matrix operations to the model in hand implies:

$$\begin{aligned}
(\Lambda+\mathbf{A}'\mathbf{L}\mathbf{A})^{-1} &= \begin{bmatrix} p_I+\rho_I & 0 & \rho_I \\ 0 & p_N+\rho_N & \rho_N \\ \rho_I & \rho_N & p_C+\rho_I+\rho_N \end{bmatrix}^{-1} \\
&= \frac{1}{p_I\rho_I(p_N+\rho_N)+p_N\rho_N(p_I+\rho_I)+p_C(p_I+\rho_I)(p_N+\rho_N)} \\
&\quad \times \begin{bmatrix} (p_N+\rho_N)(p_C+\rho_I)+p_N\rho_N & \rho_I\rho_N & -\rho_I(p_N+\rho_N) \\ \rho_I\rho_N & (p_I+\rho_I)(p_C+\rho_N)+p_I\rho_I & -\rho_N(p_I+\rho_I) \\ -\rho_I(p_N+\rho_N) & -\rho_N(p_I+\rho_I) & (p_I+\rho_I)(p_N+\rho_N) \end{bmatrix} \\
&:= \Sigma_B, \qquad \text{(A5)} \\
(\mathbf{A}'\mathbf{L}\hat{\mathbf{s}}+\Lambda\boldsymbol{\mu}) &= \begin{bmatrix} \rho_I\hat{s}_I+p_I\theta_I \\ \rho_N\hat{s}_N+p_N\theta_N \\ \rho_I\hat{s}_I+\rho_N\hat{s}_N+p_C\theta_C \end{bmatrix}. \qquad \text{(A6)}
\end{aligned}$$

Combining these results yields the probability distribution:

$$p(\mathbf{m}|\hat{s}_I,\hat{s}_N) \sim N\left(\begin{bmatrix} w_{I,s}\hat{s}_I - w_{I,C}\theta_C - w_{I,\Delta}(\hat{s}_N-\theta_N)+w_{I,I}\theta_I \\ w_{N,s}\hat{s}_N - w_{N,C}\theta_C - w_{N,\Delta}(\hat{s}_I-\theta_I)+w_{N,N}\theta_N \\ w_{C,I}(\hat{s}_I-\theta_I)+w_{C,N}(\hat{s}_N-\theta_N)+w_{C,C}\theta_C \end{bmatrix}, \Sigma_B\right), \qquad \text{(A7)}$$

where the weights are given by $w_{I,s} := \frac{\rho_I(p_Np_C+p_C\rho_N+p_N\rho_N)}{D}$, $w_{I,C} := \frac{p_C\rho_I(p_N+\rho_N)}{D}$, $w_{I,\Delta} := \frac{p_N\rho_I\rho_N}{D}$, $w_{I,I} := \frac{p_I(p_Np_C+p_C\rho_N+p_N\rho_N+p_N\rho_I+\rho_I\rho_N)}{D}$, $w_{N,s} := \frac{\rho_N(p_Ip_C+p_C\rho_I+p_I\rho_I)}{D}$, $w_{N,C} := \frac{p_C\rho_N(p_I+\rho_I)}{D}$, $w_{N,\Delta} := \frac{p_I\rho_I\rho_N}{D}$, $w_{N,N} := \frac{p_N(p_Ip_C+p_C\rho_I+p_I\rho_I+p_I\rho_N+\rho_I\rho_N)}{D}$, $w_{C,I} := \frac{p_I\rho_I(p_N+\rho_N)}{D}$, $w_{C,N} := \frac{p_N\rho_N(p_I+\rho_I)}{D}$, and $w_{C,C} := \frac{p_C(p_I+\rho_I)(p_N+\rho_N)}{D}$, where $D := [p_I\rho_I(p_N+\rho_N)+p_N\rho_N(p_I+\rho_I)+p_C(p_I+\rho_I)(p_N+\rho_N)]^{-1}$ and all weights are positive. ■

This proposition clearly illustrates that voter posterior beliefs about the level of incumbent malfeasance increase with the extent to which indicators of incumbent malfeasance exceed expectations that now explicitly adjust for updated beliefs about the common shock. Specifically, like incumbent-only information, the expected level of incumbent malfeasance increases in the difference between the signal and the prior expectation of the common shock, i.e. $\hat{s}_I-\theta_C$. However, a Bayesian voter now also uses the benchmarked signal to further account for the possibility that high incumbent malfeasance could reflect a high realization of the common shock, i.e. $\hat{s}_N-\theta_N$. Relative to receiving no information, benchmarked performance information will induce upward (downward) updating when: $\theta_I < (>) w_{I,s}\hat{s}_I - w_{I,C}\theta_C - w_{I,\Delta}(\hat{s}_N-\theta_N)+w_{I,I}\theta_I$. This will hold when malfeasance indicators, adjusted for updated expectations of the common shock, exceed prior expectations of malfeasance. The same logic applies to evaluations of neighboring incumbent parties. The voter's belief about the common shock itself, $w_{I,C}\theta_C+w_{I,\Delta}(\hat{s}_N-\theta_N)$, is intuitively increasing in the extent to which the signal exceeds the voter's prior expectations.

Combining the two propositions, benchmarked information induces a more unfavorable (favorable) posterior expectation of incumbent party malfeasance than an incumbent-only signal when:

$$w_I(\hat{s}_I-\theta_C)+(1-w_I)\theta_I < (>) w_{I,s}\hat{s}_I - w_{I,C}\theta_C - w_{I,\Delta}(\hat{s}_N-\theta_N)+w_{I,I}\theta_I. \qquad \text{(A8)}$$

Where the weights attached to the signal and prior beliefs do not drastically differ (and thus $\hat{s}_I$, θ_C, and θ_I cancel out), this expression demonstrates that voters will generally update unfavorably when $\hat{s}_N - \theta_N < 0$, i.e. when neighboring incumbent parties perform better than expected. This reflects the second signal inducing the voter to infer that there was a smaller common shock, due to such better performance, and thus becoming more likely to attribute underlying malfeasance to any high signal realization. The condition that the weights do not drastically differ implies that the additional precision imparted by the second signal does not substantially increase the weight attached to the signal vis-à-vis prior beliefs.

When it comes to vote choice, voters may instead rely on a *relative* comparison between local incumbent parties and incumbent parties elsewhere. This relative comparison contrasts with updating about beliefs about the level of incumbent malfeasance, because the signal about the challenger now serves the function of both updating about common shocks and updating about levels of challenger malfeasance. In the case where vote choice reflects a preference for the less malfeasant party, benchmarked information induces a larger difference in expected malfeasance between local incumbent and neighboring (and, thus, challenger) incumbent parties relative to incumbent-only information when:

$$w_{I,s}^* \hat{s}_I + w_{I,I}\theta_I - w_{N,s}^* \hat{s}_N - w_{N,N}\theta_N > w_I(\hat{s}_I - \theta_C - \theta_N) + (1 - w_I)(\theta_I - \theta_N), \tag{A9}$$

where the common shock is identically accounted for when comparing posterior beliefs about the incumbent and neighboring incumbents (but adjusts the weighting coefficients to account for extracting the common shock). When the weight coefficients on comparable terms are similar in magnitude, this condition is positive when $\hat{s}_N - (\theta_N + \theta_C) < 0$. To the extent that the additional signal decreases the weight attached to prior beliefs, support for the incumbent will also increase in $\hat{s}_I - \theta_I$.

A.1.3 Empirical implications

With respect to absolute posterior beliefs, these performance metrics imply the following relationships:

- The effect of benchmarked information v. incumbent-only information on malfeasance beliefs is positive when, approximately, $\theta_N > \hat{s}_N$, and is decreasing in $(\hat{s}_N - \theta_N)$.
- The effect of incumbent-only information v. no information on malfeasance beliefs is positive when $\hat{s}_I > \theta_I + \theta_C$, and is increasing in $(\hat{s}_I - \theta_I)$.
- The effect of benchmarked information v. no information on malfeasance beliefs is positive when, approximately, $\hat{s}_I > \theta_I + \mathbb{E}[m|\hat{s}_I, \hat{s}_N]$, and is increasing in $(\hat{s}_I - \theta_I)$ and decreasing in $(\hat{s}_N - \theta_N)$.

Where relevant, relationships are approximate because we assume that the weights do not meaningfully differ.

With respect to relative comparisons between incumbent and “neighboring” incumbents (a proxy for challengers, empirically), which argue approximates vote choices, these performance metrics imply the following relationships:

- The effect of benchmark v. incumbent-only information on incumbent vote share is negative when, approximately, $\theta_N + \theta_C > \hat{s}_N$, and this ATE is thus increasing in $(\hat{s}_N - \theta_N)$.
- The effect of incumbent-only information v. control on incumbent vote share is negative when $\hat{s}_I > \theta_I + \theta_C$, and is decreasing in $(\hat{s}_I - \theta_I)$.
- The effect of benchmark information v. control on incumbent vote share is negative when, approximately, $\hat{s}_I > \hat{s}_N$, and is decreasing in $(\hat{s}_I - \hat{s}_N)$.

A.2 Additional treatment graphics

Figure A1: Example comparing local and comparative leaflets from Ecatepec, Mexico

Figure A2: Loudspeaker accompanying leaflet delivery in the public treatment

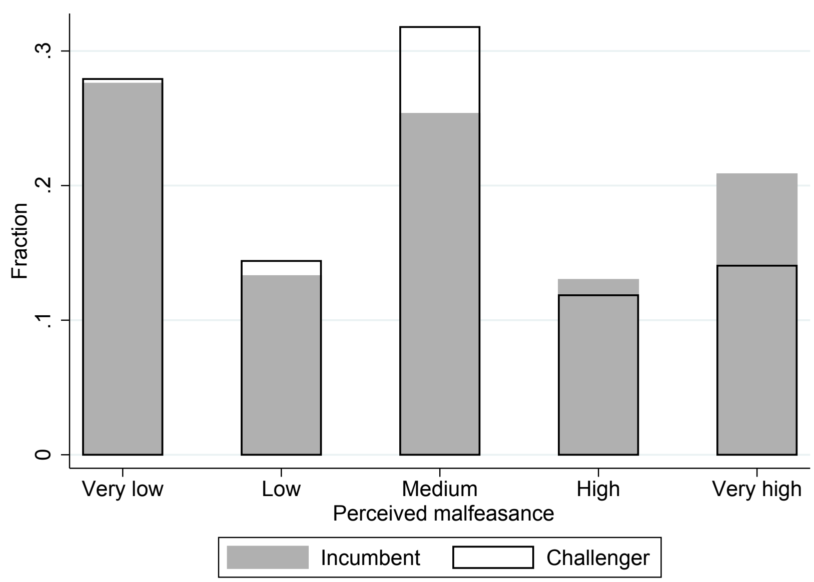

Figure A3: Perceptions of incumbent and challenger malfeasance in *control* precincts

Note: The challenger is defined as the party receiving the second largest vote share at the previous municipal election in 2012.

A.3 Deviations from pre-analysis plan

In our pre-analysis plan we mis-wrote our hypothesis regarding the differential effect of comparative performance information on beliefs on incumbent malfeasance. In H22 we wrote that the

Table A1: Precinct-level comparison of Census 2010 characteristics between our sample and the nation

Variable	Experimental Sample Observations	Mean	Std. dev.	All Precincts in Mexico Observations	Mean	Std. dev.
Population	678	1,633.18	997.00	66,740	1,683.20	1,878.04
Share working age	678	0.63	0.06	66,685	0.63	0.06
Average children per woman	678	2.49	0.58	66,740	2.50	0.62
Share indigenous speakers	678	0.05	0.15	66,682	0.06	0.19
Average years of schooling	678	7.98	2.39	66,740	8.27	2.47
Share economically active	678	0.38	0.07	66,685	0.39	0.07
Average occupants per room	678	1.16	0.28	66,740	1.11	0.35
Share of homes with water, drainage, and electricity	678	0.41	0.29	66,681	0.41	0.27
Shares of homes with a television	678	0.91	0.14	66,681	0.90	0.15
Share of homes with internet	678	0.16	0.19	66,681	0.19	0.20

Note: All variables are unweighted.

comparative treatment would "on average, have a weaker effect on perceptions of corruption or lack of interest on marginalized populations of the incumbent than the local treatment". As section A.1 illustrates, this is a mistake as our Bayesian framework suggests that comparative performance information will induce more unfavorable incumbent updating where challenger party malfeasance falls below expectations—a likely condition to hold in our particular experimental context. Consistent with this, in the pre-analysis plan we hypothesized that the comparative treatment would have a more negative effect on the incumbent's vote share than the local treatment (see H9).

The other deviation is that we weight precinct-level observations by the share of voters within the precinct to whom we delivered a leaflet. Although we did not pre-register weighting scheme, we ultimately believed that—in light of significant variation in the fraction of a precinct that 200 leaflets could reach—it was important to attach greater weight to precincts in which a larger fraction of the voting population received a leaflet and thus for which precinct-level electoral results are a better signal of voter's behavior. However, as Tables A9-A10 show, all of our point estimates are similar if we run unweighted regressions.

A.4 Summary statistics

Table A1 uses 2010 Census characteristics to compare our sample of 678 precincts to the national distribution. As noted in the main text, our sample is broadly nationally representative with respect to these indicators.

A.5 Validation of measures of voters' prior beliefs

We provide evidence to support our claim that post-treatment beliefs and updating in the control precincts proxy for pre-treatment prior beliefs and updating in the treated precincts within the same municipality. To do so, we show that the three key assumptions—(1) that control group respondents are similar to treatment group respondents, (2) that control group respondent beliefs are consistent across the month between the intervention and the post-election survey, and (3) that control group respondents internalized the information similarly to those in treated precincts—are plausible in the context of this study.

First, our randomization ensures that treated and control precincts are identical in expectation. The balance over individual-level characteristics observed in Table A2 is particularly important because it indicates that our treatment did not affect the willingness of different types of voters to participate in the endline survey. Moreover, our blocking strategy ensures substantial within-block similarity in practice: block fixed effects account for 60% of the variation in precinct-level incumbent vote share and 29% of the variation in individual-level beliefs within our samples.

Second, we examine whether the election outcome itself influenced beliefs between the dissemination of the treatment and the post-election survey. Table A3 shows that the 2015 *municipal*-level election outcomes are generally uncorrelated with the level of beliefs about incumbent party malfeasance among respondents in the control group, conditioning on the municipal incumbent party's vote share in the previous election—a pre-treatment proxy for prior beliefs in the control group. The exception is in column (4), where the municipal incumbent party's vote share is positively correlated with the precision of prior beliefs in the control group. However, the magnitude is small: a 70 percentage point increase in vote share is required to increase the precision of beliefs in the control group by a standard deviation. Moreover, the election outcome itself is not significantly correlated with belief precision in the control group. The results suggest that the intervening election outcomes themselves did not substantially influence voter beliefs (and thus violate our second assumption). This is not surprising, since electoral expectations were likely to be relatively fixed in advance and the scale of our intervention was specifically designed not to influence electoral outcomes.

Third, and more generally, the 2012 Mexican Panel Survey shows that voter assessments of politicians are relatively persistent in the months prior to the election. Voters' opinions of the presidential candidates before and after the election—three months apart, in contrast to the 3–4 weeks apart we examine—exhibit a 0.4 correlation.

Fourth, if the information is indeed novel to the control group, then the control group should update its beliefs substantially more than the treatment group after being shown the leaflet at the end of the post-election survey. Table A4 shows that control respondents perceive their incumbent to be more malfeasant when shown a leaflet revealing high levels of malfeasance for the first time at the end of the post-election survey. Control respondents thus seem to react similarly to treated respondents, suggesting that treated respondents likely possessed similar prior beliefs and that control group respondents responded similarly to reading the leaflet during the survey to how treated respondents responded to its delivery in the field.

Finally, we use data from a similar randomized intervention to ours conducted around the October 2016 Brazilian municipal elections by Boas, Hidalgo and Melo (forthcoming). This study informed voters about the local government's use of funds and about educational performance in the municipality. Critical for our purposes, their study collected voters' beliefs on local governments' performance at both baseline and endline, which allows us to look directly at the extent to which endline beliefs of respondents in control units are valid proxies for the prior beliefs of respondents in treated units. Our own analysis of the Brazilian data (additional details available upon request) reveals that:

1. Correlation of baseline priors for treatment and control is large and positive (0.86), which is perhaps not surprising, given that treatment was randomly assigned.

2. The correlation between the control group at baseline and endline is 0.86. Survey responses are noisy, and thus we would not expect a perfect serial correlation even absent any treat-

Table A2: Balance across 40 precinct-level variables and 8 individual-level variables

	Control		Private local treatment		Public local treatment		Private comparative treatment		Public comparative treatment			Test: all treatment effects
	Mean	Std. dev.	Coefficient	Std. error	Coefficient	Std. error	Coefficient	Std. error	Coefficient	Std. error	Observations	=0 (p value)
Precinct-level covariates												
Area	10.020	[19.19]	-1.603	(1.47)	0.441	(1.83)	-1.751	(1.145)	-1.435	(1.299)	675	0.54
Population	1,372.550	[783.43]	21.385	(64.888)	15.329	(70.627)	-29.826	(75.527)	41.928	(62.579)	675	0.96
Population density	6,126.540	[7512.33]	186.065	(267.087)	120.194	(552.855)	-160.258	(500.067)	-516.456	(330.143)	675	0.33
Distance from municipal centroid	7,645.410	[6889.7]	242.372	(511.366)	526.728	(389.86)	1094.903**	(428.663)	641.529	(649.42)	675	0.11
Number of households	329.380	[174.73]	5.985	(15.58)	4.451	(16.339)	-6.479	(17.188)	9.382	(13.859)	675	0.95
Number of private dwellings	395.930	[214.92]	2.798	(18.574)	10.463	(19.573)	-8.505	(19.399)	2.465	(17.023)	675	0.97
Average occupants dwelling	4.100	[0.52]	0.010	(0.041)	0.025	(0.035)	0.003	(0.036)	0.028	(0.037)	675	0.89
Average occupants per room	1.150	[0.28]	0.025	(0.019)	-0.012	(0.019)	0.025	(0.02)	-0.001	(0.021)	675	0.36
Share of homes with 2+ rooms	0.660	[0.14]	-0.017*	(0.01)	0.023*	(0.012)	-0.004	(0.014)	-0.005	(0.011)	675	0.10
Share of homes with 3+ rooms	0.760	[0.13]	-0.020**	(0.01)	0.020*	(0.012)	-0.008	(0.014)	0.000	(0.012)	675	0.07
Average years of schooling	8.120	[2.47]	-0.198**	(0.091)	0.035	(0.127)	-0.216	(0.17)	-0.114	(0.097)	675	0.12
Share married	0.550	[0.05]	-0.003	(0.004)	0.003	(0.005)	-0.002	(0.006)	0.005	(0.004)	675	0.47
Share working age	0.630	[0.06]	-0.004	(0.004)	0.001	(0.004)	-0.004	(0.006)	-0.003	(0.004)	675	0.63
Share economically active	0.380	[0.07]	-0.001	(0.004)	0.000	(0.004)	-0.001	(0.004)	0.000	(0.005)	675	1.00
Share without health care	0.340	[0.12]	0.014	(0.012)	-0.008	(0.011)	0.021	(0.013)	0.019*	(0.011)	675	0.07
Share with state workers health care	0.040	[0.05]	-0.004	(0.003)	0.001	(0.005)	0.000	(0.004)	0.002	(0.003)	675	0.56
Share old	0.060	[0.03]	0.002	(0.003)	0.000	(0.003)	-0.001	(0.002)	0.003	(0.003)	675	0.55
Average children per woman	2.470	[0.62]	0.059*	(0.031)	0.043	(0.033)	0.081**	(0.04)	0.071**	(0.033)	675	0.02
Share of households with male head	0.770	[0.07]	-0.006	(0.006)	0.003	(0.007)	-0.007	(0.01)	-0.002	(0.007)	675	0.80
Share born out of state	0.270	[0.26]	0.001	(0.01)	0.014	(0.009)	0.008	(0.01)	0.013	(0.01)	675	0.52
Share indigenous speakers	0.060	[0.17]	-0.001	(0.013)	0.007	(0.007)	0.009*	(0.005)	0.015	(0.016)	675	0.50
Share of homes without a dirt floor	0.920	[0.11]	-0.007	(0.012)	0.001	(0.007)	0.003	(0.007)	-0.008	(0.018)	675	0.81
Share of homes with a toilet	0.890	[0.18]	0.000	(0.012)	0.009	(0.012)	0.003	(0.012)	0.007	(0.011)	675	0.88
Share of homes with water	0.840	[0.27]	0.032	(0.023)	0.018	(0.021)	0.003	(0.032)	-0.017	(0.027)	675	0.73
Share of homes with drainage	0.830	[0.24]	0.009	(0.012)	0.015	(0.02)	-0.001	(0.017)	-0.014	(0.022)	675	0.89
Share of homes with electricity	0.960	[0.09]	0.003	(0.006)	0.006	(0.007)	0.008*	(0.005)	0.001	(0.009)	675	0.24
Share of homes with water, drainage, and electricity	0.760	[0.31]	0.017	(0.019)	0.014	(0.022)	-0.013	(0.031)	-0.019	(0.026)	675	0.88
Share of homes with a washing machine	0.580	[0.26]	-0.001	(0.011)	0.009	(0.015)	-0.002	(0.014)	0.007	(0.011)	675	0.95
Share of homes with a landline telephone	0.420	[0.29]	-0.032***	(0.011)	-0.002	(0.013)	-0.032	(0.021)	-0.013	(0.012)	675	0.03
Share of homes with a radio	0.820	[0.1]	0.004	(0.006)	0.006	(0.007)	-0.011	(0.009)	-0.007	(0.008)	675	0.73
Share of homes with a fridge	0.750	[0.23]	-0.009	(0.013)	0.011	(0.014)	-0.018	(0.018)	0.006	(0.013)	675	0.55
Share of homes with a cell phone	0.550	[0.25]	-0.014	(0.013)	0.012	(0.013)	-0.001	(0.013)	0.008	(0.013)	675	0.78
Share of homes with a television	0.900	[0.15]	0.000	(0.007)	-0.005	(0.008)	-0.005	(0.008)	-0.006	(0.014)	675	0.87
Number of local media stations	2.320	[3.16]	0.055	(0.052)	0.008	(0.039)	0.043	(0.046)	0.105**	(0.048)	675	0.16
Share of homes with a car	0.390	[0.18]	-0.027*	(0.015)	-0.008	(0.017)	-0.010	(0.016)	-0.003	(0.011)	675	0.47
Share of homes with a computer	0.250	[0.24]	-0.021**	(0.01)	0.006	(0.013)	-0.015	(0.015)	-0.012	(0.009)	675	0.15
Share of homes with internet	0.170	[0.2]	-0.018*	(0.01)	-0.002	(0.012)	-0.011	(0.014)	-0.011	(0.008)	675	0.35
Turnout in 2012	0.630	[0.08]	0.006	(0.006)	0.007	(0.007)	0.013*	(0.008)	0.006	(0.005)	675	0.14
Incumbent vote party margin in 2012	-0.170	[0.13]	-0.031**	(0.015)	-0.009	(0.016)	-0.046	(0.028)	-0.018	(0.013)	675	0.13
Incumbent vote party share in 2012	0.420	[0.12]	0.017	(0.013)	0.004	(0.011)	0.032*	(0.019)	0.021*	(0.01)	675	0.10
Survey-level covariates												
Female	0.64	[0.48]	0.035	(0.022)	-0.009	(0.022)	0.013	(0.024)	0.044*	(0.023)	4,958	0.03
Age	44.42	[16.07]	-0.274	(0.697)	-0.986	(0.824)	-0.519	(0.667)	-0.309	(0.781)	4,869	0.81
Education	8.03	[4.14]	-0.277	(0.182)	0.11	(0.178)	-0.088	(0.183)	0.005	(0.171)	4,948	0.33
Ln(income)	1.14	[0.44]	-0.052**	(0.023)	0.008	(0.022)	-0.012	(0.021)	0.014	(0.024)	4,402	0.09
Employed	0.42	[0.49]	-0.024	(0.019)	0.015	(0.019)	0.002	(0.022)	-0.017	(0.021)	4,950	0.29
Turnout in 2012	0.63	[0.48]	-0.026	(0.019)	0.027	(0.018)	0.001	(0.018)	0.015	(0.017)	4,958	0.07
Voted for Incumbent in 2012	0.54	[0.5]	0.022	(0.03)	-0.016	(0.023)	-0.032	(0.026)	-0.001	(0.031)	3,122	0.23
Political knowledge index	2.4	[0.86]	0.004	(0.039)	0.024	(0.039)	0.047	(0.036)	-0.054	(0.037)	4,958	0.34

Notes: All specifications include block fixed effects, and precinct-level estimates are weight by the share of the precinct that was treated, and are estimated using OLS. Standard errors clustered by municipality-treatment are in parentheses. * denotes $p < 0.1$, ** denotes $p < 0.05$, *** denotes $p < 0.01$.

Table A3: Correlation between municipal-level election outcomes and prior beliefs in the control group

	Incumbent malfeasance prior	
	(1)	(2)
Municipal incumbent won election (2015)	-0.516	
	(0.382)	
Municipal incumbent vote share (2015)		-1.713
		(1.661)
Municipal incumbent vote share (2012)	3.307*	3.723**
	(1.690)	(1.767)
Constant	-1.198	-1.110
	(0.779)	(1.007)
Control outcome mean	-0.14	-0.14
Control outcome std. dev.	1.48	1.48
2015 election outcome mean	0.75	0.38
2015 election outcome std. dev.	0.44	0.08
Observations	1,038	1,038

Notes: Specifications are estimated using OLS. Standard errors clustered by municipality are in parentheses. * denotes $p < 0.1$, ** denotes $p < 0.05$, *** denotes $p < 0.01$.

ment, as other events between baseline and endline (i.e. the election and the preceding campaign) may change some people's preferences. So a positive correlation of around 0.9 is consistent with control group respondent beliefs being consistent across the month between the intervention and the post-election survey.

3. The correlation between prior beliefs of the treated group and the endline evaluations of the control group is 0.78, which suggests that the latter may be used as valid proxies for baseline responses of the treated.

Since this exercise was performed in the context of a different country and a different intervention, it is hard to assess the extent to which these correlations would be similar in the context of our experiment had we conducted a baseline survey. However, together with the evidence reported in Tables A3-A4, these results are encouraging regarding the use of our approach to proxy for voters' prior beliefs.

A.6 No evidence of cross-precinct spillovers

In this section, we test for whether control precincts were subject to information spillovers. Table A5 reports the effects of treatment spillovers from precincts in our experimental sample to neighboring precincts (any precinct that partially borders a precinct in our experimental sample) that were not in our experimental sample. Here, the unit of observation is the precinct-neighbor level; precincts are inversely weighted by the number of neighbors in the experimental sample.

Table A4: Effect of showing voters the leaflet in the post-treatment survey

	Perceived incumbent party malfeasance (very low - very high)	
	(1)	(2)
Shown leaflet for first time	0.061*	-0.008
	(0.031)	(0.043)
× Incumbent malfeasant spending		0.329*
		(0.171)
Perceived incumbent party malfeasance (pre-leaflet)	-0.001	-0.002
	(0.041)	(0.041)
Outcome range	{-2,-1,0,1,2}	{-2,-1,0,1,2}
Control outcome mean	0.75	0.75
Control outcome std. dev.	1.07	1.07
Interaction range		[0,0.58]
Interaction mean		0.21
Interaction std. dev.		0.17
Observations	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 $p < 0.01$.

The positive interaction with the malfeasance level reported is exactly opposite to our findings and prediction from our theoretical framework. It is then hard to see how these results could reflect our information treatment. Table A6 shows that leaflet recall is unaffected by the share of treated neighbors among respondents in control precincts. In addition, columns (5) and (6) show that the increased political responses in treated precincts do not spill over into neighboring control precincts. These checks indicate that information from treated precincts did not influence beliefs in the control group in the three weeks between the treatment and the post-election survey, and thus violate our second assumption.

A.7 Additional results

Tables A7-A11 report additional results cited in the main text.

Table A5: Neighbor spillover effects of information treatment on incumbent party vote share

	Incumbent party vote share		
	(1)	(2)	(3)
Panel A: Incumbent party vote share (share of turnout)			
Neighbor information treatment	-0.001	-0.008**	-0.002
	(0.003)	(0.004)	(0.004)
× Incumbent malfeasant spending		0.028**	
		(0.011)	
× Neighbor unfavorable incumbent updating			0.001
			(0.003)
Outcome range	[0.05,0.89]	[0.05,0.89]	[0.05,0.89]
Control outcome mean	0.39	0.39	0.39
Control outcome std. dev.	0.12	0.12	0.12
Panel B: Incumbent party vote share (share of registered voters)			
Neighbor information treatment	-0.003*	-0.008***	-0.004
	(0.002)	(0.003)	(0.003)
× Incumbent malfeasant spending		0.022***	
		(0.007)	
× Neighbor unfavorable incumbent updating			0.001
			(0.001)
Outcome range	[0.03,0.46]	[0.03,0.46]	[0.03,0.46]
Control outcome mean	0.19	0.19	0.19
Control outcome std. dev.	0.06	0.06	0.06
Interaction range		[0,0.58]	[-0.6,2.7]
Interaction mean		0.24	0.97
Interaction std. dev.		0.19	1.05
Observations	2,297	2,297	2,263

Notes: The sample contains all precinct-neighboring precincts pairs for which the neighboring precinct (which partially shares a border with a precinct in the experimental sample) is included in the experimental sample, but the spillover precinct is not. Specifications include neighbor-level block fixed effects, weight by the share of the neighboring precinct that was treated divided by the number of precincts in the experimental sample that a precinct neighbors, and are estimated using OLS. Lower-order interaction terms are absorbed by the block fixed effects. The smaller sample in Column (3) 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 $p < 0.01$.

Table A6: Neighbor spillover of information treatment on self-reported engagement with leaflet and political responses in control precincts

	Remember leaflet (1)	Remember reading leaflet (2)	Correctly remember content (3)	Leaflet influenced vote (4)	Total incumbent activities (5)	Total challenger activities (6)
Share of treated neighbors	-0.014	-0.013	-0.017	0.007	-0.396*	-0.254
	(0.040)	(0.024)	(0.022)	(0.011)	(0.193)	(0.183)
Outcome range	{0,1}	{0,1}	{0,1}	{0,1}	{0,1,2,3,4,5}	{0,1,2,3,4,5}
Outcome mean	0.09	0.05	0.06	0.02	0.43	0.40
Outcome std. dev.	0.28	0.22	0.25	0.14	1.18	1.17
Share of treated neighbors mean	0.41	0.41	0.41	0.41	0.41	0.41
Share of treated neighbors std. dev.	0.42	0.42	0.42	0.42	0.42	0.42
Observations	1,139	1,139	1,139	1,139	1,139	1,139

Notes: The sample includes all control precincts within our experimental sample. All specifications are estimated using OLS. Standard errors clustered by municipality-treatment are in parentheses. * denotes $p < 0.1$, ** denotes $p < 0.05$, *** denotes $p < 0.01$.

Table A7: Effect of local and comparative information treatments on voters' posterior beliefs about challenger party malfeasance, where the challenger is each voter's second-choice party

	Perceived challenger party malfeasance (very low - very high)		
	(1)	(2)	(3)
Panel A: Control group as baseline			
Local treatment	-0.005	-0.100	-0.035
	(0.037)	(0.095)	(0.044)
Comparative treatment	-0.054	-0.050	-0.111**
	(0.041)	(0.095)	(0.044)
Local treatment × Challenger malfeasant spending		0.436	
		(0.777)	
Comparative treatment × Challenger malfeasant spending		0.571	
		(0.899)	
Local treatment × Incumbent malfeasant spending		0.258	
		(0.209)	
Comparative treatment × Incumbent malfeasant spending		-0.257	
		(0.215)	
Local treatment × Unfavorable challenger updating			0.016
			(0.082)
Comparative treatment × Unfavorable challenger updating			0.196**
			(0.086)
Local treatment × Unfavorable incumbent updating			0.019
			(0.058)
Comparative treatment × Unfavorable incumbent updating			-0.075
			(0.058)
Control outcome mean	-0.19	-0.19	-0.19
Control outcome std. dev.	1.30	1.30	1.30
Test: same treatment effect (p value)	0.26	0.68	0.07
Test: same interaction (1) effect (p value)		0.89	0.00
Test: same interaction (2) effect (p value)		0.02	0.07
Observations	4,958	4,958	4,958
Panel B: Local treatment group as baseline			
Comparative treatment	-0.059	0.044	-0.087**
	(0.041)	(0.085)	(0.039)
Comparative treatment × Challenger malfeasant spending		-0.131	
		(0.872)	
Comparative treatment × Incumbent malfeasant spending		-0.427*	
		(0.227)	
Comparative treatment × Unfavorable challenger updating			0.163**
			(0.066)
Comparative treatment × Unfavorable incumbent updating			-0.083
			(0.052)
Local treatment outcome mean	-0.22	-0.22	-0.22
Local treatment outcome std. dev.	1.36	1.36	1.36
Observations	3,819	3,819	3,819
Outcome range	{-2,-1,0,1,2}	{-2,-1,0,1,2}	{-2,-1,0,1,2}

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 $p < 0.01$.

Table A8: Effect of local and comparative information treatments on voters' posterior beliefs about challenger party malfeasance, where the challenger is the average posterior belief across the PAN, PRD, and PRI where they are not the municipal incumbent

	Perceived challenger party malfeasance (very low - very high)		
	(1)	(2)	(3)
Panel A: Control group as baseline			
Local treatment	0.032	-0.079	-0.009
	(0.037)	(0.087)	(0.037)
Comparative treatment	-0.014	-0.046	-0.044
	(0.039)	(0.084)	(0.036)
Local treatment × Challenger malfeasant spending		1.551**	
		(0.752)	
Comparative treatment × Challenger malfeasant spending		1.067	
		(0.816)	
Local treatment × Incumbent malfeasant spending		-0.137	
		(0.239)	
Comparative treatment × Incumbent malfeasant spending		-0.302	
		(0.240)	
Local treatment × Unfavorable challenger updating			0.109*
			(0.059)
Comparative treatment × Unfavorable challenger updating			0.150**
			(0.066)
Local treatment × Unfavorable incumbent updating			-0.049
			(0.062)
Comparative treatment × Unfavorable incumbent updating			-0.091
			(0.063)
Control outcome mean	-0.33	-0.33	-0.33
Control outcome std. dev.	1.20	1.20	1.20
Test: same treatment effect (p value)	0.17	0.56	0.38
Test: same interaction (1) effect (p value)		0.45	0.44
Test: same interaction (2) effect (p value)		0.34	0.34
Observations	4,958	4,958	4,958
Panel B: Local treatment group as baseline			
Comparative treatment	-0.038	0.035	-0.030
	(0.035)	(0.059)	(0.041)
Comparative treatment × Challenger malfeasant spending		-0.385	
		(0.691)	
Comparative treatment × Incumbent malfeasant spending		-0.180	
		(0.183)	
Comparative treatment × Unfavorable challenger updating			0.054
			(0.059)
Comparative treatment × Unfavorable incumbent updating			-0.051
			(0.047)
Local treatment outcome mean	-0.24	-0.24	-0.24
Local treatment outcome std. dev.	1.20	1.20	1.20
Observations	3,819	3,819	3,819
Outcome range	{-2,-1,0,1,2}	{-2,-1,0,1,2}	{-2,-1,0,1,2}

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 $p < 0.01$.

Table A9: Effect of local and comparative information treatments on incumbent party vote share, unweighted estimates

	Incumbent party vote share (share of registered voters)		
	(1)	(2)	(3)
Panel A: Control group as baseline			
Local treatment	0.004	0.015*	0.012***
	(0.003)	(0.008)	(0.003)
Comparative treatment	0.012***	0.017*	0.012***
	(0.003)	(0.009)	(0.004)
Local treatment × Incumbent malfeasant spending		-0.027	
		(0.022)	
Comparative treatment × Incumbent malfeasant spending		-0.032*	
		(0.018)	
Local treatment × Challenger malfeasant spending		-0.053	
		(0.061)	
Comparative treatment × Challenger malfeasant spending		0.018	
		(0.071)	
Local treatment × Unfavorable incumbent updating			-0.007
			(0.005)
Comparative treatment × Unfavorable incumbent updating			-0.009*
			(0.005)
Local treatment × Unfavorable challenger updating			-0.003
			(0.005)
Comparative treatment × Unfavorable challenger updating			0.009*
			(0.005)
Control outcome mean	0.19	0.19	0.20
Control outcome std. dev.	0.07	0.07	0.07
Test: same treatment effect (p value)	0.15	0.86	1.00
Test: same interaction (1) effect (p value)		0.86	0.78
Test: same interaction (2) effect (p value)		0.47	0.12
Observations	675	675	651
Panel B: Local treatment group as baseline			
Comparative treatment	0.007	0.003	0.000
	(0.005)	(0.013)	(0.006)
Comparative treatment × Incumbent malfeasant spending		-0.006	
		(0.033)	
Comparative treatment × Challenger malfeasant spending		0.067	
		(0.104)	
Comparative treatment × Unfavorable incumbent updating			-0.002
			(0.008)
Comparative treatment × Unfavorable challenger updating			0.012
			(0.008)
Local treatment outcome mean	0.20	0.20	0.20
Local treatment outcome std. dev.	0.07	0.07	0.07
Observations	398	398	382
Outcome range	[0.03,0.40]	[0.03,0.40]	[0.03,0.40]

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 column (3) 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 $p < 0.01$.

Table A10: Effect of private and public information treatments on incumbent party vote share, unweighted estimates

	Incumbent party vote share (share of registered voters)		
	(1)	(2)	(3)
Panel A: Control group as baseline			
Private treatment	0.011***	0.024**	0.014***
	(0.003)	(0.010)	(0.004)
Public treatment	0.005	0.008	0.010**
	(0.003)	(0.010)	(0.004)
Private treatment × Incumbent malfeasant spending		-0.048***	
		(0.018)	
Public treatment × Incumbent malfeasant spending		-0.011	
		(0.018)	
Private treatment × Challenger malfeasant spending		-0.023	
		(0.071)	
Public treatment × Challenger malfeasant spending		-0.012	
		(0.072)	
Private treatment × Unfavorable incumbent updating			-0.013***
			(0.004)
Public treatment × Unfavorable incumbent updating			-0.002
			(0.004)
Private treatment × Unfavorable challenger updating			0.011**
			(0.005)
Public treatment × Unfavorable challenger updating			-0.005
			(0.003)
Control outcome mean	0.19	0.19	0.20
Control outcome std. dev.	0.07	0.07	0.07
Test: same treatment effect (p value)	0.20	0.34	0.54
Test: same interaction (1) effect (p value)		0.13	0.07
Test: same interaction (2) effect (p value)		0.92	0.01
Observations	675	675	651
Panel B: Private treatment group as baseline			
Public treatment	-0.006	-0.015	-0.004
	(0.005)	(0.017)	(0.007)
Public treatment × Incumbent malfeasant spending		0.037	
		(0.026)	
Public treatment × Challenger malfeasant spending		0.008	
		(0.120)	
Public treatment × Unfavorable incumbent updating			0.011
			(0.006)
Public treatment × Unfavorable challenger updating			-0.017**
			(0.006)
Private treatment outcome mean	0.20	0.20	0.20
Private treatment outcome std. dev.	0.07	0.07	0.07
Observations	398	398	382
Outcome range	[0.03,0.40]	[0.03,0.40]	[0.03,0.40]

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 column (3) 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 $p < 0.01$.

Table A11: Effect of public treatment on belief about the leaflet's provenance

	Believe that the leaflet was disseminated by...				
	...municipal incumbent party	...municipal challenger party	...PAN	...PRD	...PRI
	(1)	(2)	(3)	(4)	(5)
Public treatment	0.009	-0.005	-0.007	-0.016	-0.002
	(0.019)	(0.013)	(0.013)	(0.012)	(0.017)
Outcome range	{0,1}	{0,1}	{0,1}	{0,1}	{0,1}
Outcome mean	0.26	0.16	0.14	0.12	0.17
Outcome std. dev.	0.44	0.36	0.35	0.33	0.38
Observations	3,659	3,659	3,659	3,659	3,659

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