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¶

May 2017

Information pertaining to incumbent performance can play a central role in electoral accountability. In light of mixed findings, we examine the extent to which performance benchmarks and common knowledge moderate the effects of information dissemination. These theoretical mechanisms are tested in the context of a field experiment providing voters with audit reports documenting mayoral malfeasance before the 2015 Mexican municipal elections. We find that neither benchmarking against the performance of mayors from other parties within the state nor accompanying leaflet delivery with loudspeakers announcing the leaflets' delivery within the community significantly altered the effect of information on incumbent party vote share. The ineffectiveness of benchmarking reflects voters not updating their beliefs from the benchmark, while the loudspeaker created common knowledge without facilitating coordination.

*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 Politics, New York University. Email: eric.arias@nyu.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 electoral 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 amassed is mixed. On one hand, various studies show that the revelation of good performance is indeed rewarded, while bad performance is punished (Banerjee et al. 2011; Ferraz and Finan 2008; Humphreys and Weinstein 2012), at least when it exceeds voters' expectations (Arias, Larreguy, Marshall and Querubín 2017). On the other hand, other studies observe little effect of providing information, or even that bad performance is rewarded (Adida et al. 2016; Chong et al. 2015; de Figueiredo, Hidalgo and Kasahara 2013).

In this article, we test two explanations that could help account for such varying results: performance benchmarks and common knowledge. Performance benchmarks provide voters with a point of reference to evaluate their incumbent politician's performance against, e.g. other incumbents working in similar contexts or previously working in the same context (Besley 2006; Holmstrom 1999). Previous studies may have failed to find an effect of information provision on electoral behavior because voters don't know how to interpret the information provided. To the extent that incumbent performance information suggests that an incumbent outperforms comparable challengers, voters may favorably update their beliefs about their incumbent and subsequently reward them or their political party. Common knowledge has the potential to induce tacit or explicit coordination among voters by establishing common knowledge or by stimulating coordinated action through discussion within communities (e.g. Arias, Balán, Larreguy, Marshall and Querubín 2017; Chwe 2000; Jackson 2010). This could strengthen electoral accountability by magnifying

responses to information through mechanisms often attributed to mass media (Adena et al. 2015; Arias 2016; Yanagizawa-Drott 2014).

Our experiment examines how the effect of providing Mexican voters with incumbent performance information varies with the *content* and *form* of the information disseminated. In a complementary paper, where the empirical analysis pools all treatments conditions analyzed separately in this paper, we demonstrated that revealing often severe levels of mayoral malfeasance increased the incumbent party's vote share on average (Arias, Larreguy, Marshall and Querubín 2017). We then showed that this in large part reflects voters' prior beliefs: our treatment particularly increased incumbent support among voters that were informed of low malfeasance, voters who already believed that their incumbent party was highly malfeasant, and voters who favorably updated their posterior beliefs. In this article, we extend this study by further exploring how variants on the information provided affect voters' posterior beliefs and ultimately their voting behavior. Specifically, in addition to reporting the share of funds earmarked for social infrastructure projects benefiting the poor that were spent on either unauthorized projects or projects that did not actually benefit the target population, we examine additional treatment conditions that: (1) benchmarked municipal mayors against the average performance of mayors from other parties within the state; and (2) accompanied the 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 benchmarked information or a loudspeaker, while the loudspeaker increased common knowledge.

However, we find little evidence that either variant of the standard information treatment translated into substantially different voting behavior. Contrary to theoretical expectations based on the information provided, benchmarking information against other parties within the state is neither more likely to decrease the vote share for the incumbent party nor more likely to decrease the vote share for the incumbent party when challengers perform comparatively better. Analysis of our survey data suggest that voters either did not understand the benchmark or did not believe

it to be relevant. Similarly, we find no evidence to suggest that public dissemination magnified, or otherwise moderated, the effect of providing leaflets. Despite increasing common knowledge, the loudspeaker did not significantly enhance voter coordination. We conclude by discussing the implications of these findings.

2 Background

This section provides a brief overview of our empirical context (see Arias, Larreguy, Marshall and Querubín 2017; Chong et al. 2015; Larreguy, Marshall and Snyder 2017 for greater detail), before outlining our theoretical framework and hypotheses.

2.1 Audits of Mexican municipal mayors

Mexico's federal system comprises of 31 states (and the Federal District of Mexico City), which in turn contain c.2,500 municipalities and c.67,000 electoral precincts. Within municipalities—the focus of this study—mayors serve non-renewable terms generally lasting for three years.[1] Although mayors could not seek re-election, voters have consistently been shown to hold parties responsible for the actions of individual politicians in Mexico's party-centric system (e.g. Chong et al. 2015; Larreguy, Marshall and Snyder 2017).

Mayors are responsible for delivering basic public services and managing local infrastructure, and account for 20% of total government spending. A key component of this is the Municipal Fund for Social Infrastructure (FISM), which represents 24% of the average mayor's budget. FISM funds are federal transfers mandated exclusively for infrastructure projects—such as investments in water supply, drainage, electrification, health infrastructure, education infrastructure, housing, and roads—that benefit impoverished citizens. Although these often represent major projects, voters

[1]Mayors will be able to seek re-election for the first time in 2018 in most states. Municipal elections are generally held alongside state legislative elections, and in some cases federal legislative and presidential elections.

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. 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 share of FISM funds in the municipal budget, previous performance, factors increasing the risk of mismanagement, and recency of the last audit (see Auditoría Superior de la Federación 2014). Audits are not announced until the year after spending occurred, and reports are released in the February two calendar years after spending occurred.

ASF reports detail two key dimensions of mayoral malfeasance that form the basis of 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 (i.e. non-social infrastructure) projects. Such violations are similar to the corrupt practices documented in Ferraz and Finan (2008), often including procurement violations and electorally-targeted projects. The 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. There is considerable heterogeneity in such malfeasance across municipalities.

2.2 Hypotheses

There are good reasons to believe that information will affect voting behavior. The simple two-party learning model in Arias, Larreguy, Marshall and Querubín (2017) characterizes voters as Bayesians seeking to select the least malfeasant politician on the basis of the information available. While the effect of providing information is ambiguous and depends on how the information relates to voter prior beliefs, the vote share of the incumbent party increases in the extent to which the information causes voters to update their posterior beliefs about the incumbent's malfeasance favorably and causes voters to reduce their uncertainty about such beliefs. More specifically, the

model predicts that the effect of information provision on incumbent party vote share is decreasing in the favorability of prior beliefs, the level of malfeasance reported, the extent of unfavorable belief updating, and decreasing (increasing) in the precision of prior beliefs when the information induces favorable (unfavorable) voter updating.

Information about challengers feeds directly into this framework. Voters re-elect the incumbent to the extent that they believe—conditional on their partisan attachments—that the challenger will be less malfeasant than the incumbent. Consequently, providing benchmarked information detailing the performance of both incumbent and challenger that causes voters to update more favorably about the challenger than the incumbent is likely to induce voters, on the margin, to vote for the challenger. Based on our initial assumption that voters possess relatively similar prior beliefs about the malfeasance of incumbent and challenger, for the typical malfeasance revelations in our experimental sample—where the majority of voters receive information indicating less malfeasance among challengers—we pre-registered the following hypothesis:[2]

H1. *Comparative malfeasance information, on average, decreases the incumbent party's vote share more than just incumbent malfeasance information.*

Furthermore, we expected this anti-incumbent party effect to be particularly large where voters learn of especially high incumbent malfeasance—in both absolute and relative terms:

H2. *The differential effect on incumbent party vote share of providing comparative malfeasance information, relative to providing only incumbent malfeasance information, is increasing in the level of reported incumbent malfeasance, and the difference relative to the reported malfeasance of the challenger.*

Although our guiding model of voting behavior does not explicitly incorporate voter interactions, public delivery mechanisms could also play an important role in electoral accountability.

[2]We also pre-registered a hypothesis distinguishing between good and bad news, based on how the information was perceived by control respondents upon being shown the treatment at the end of our post-election survey. However, this measure of updating may only be reliable in relative rather than absolute terms. This reduces our confidence in the (absolute) cutoffs for defining good and bad news.

Public modes of transmission—through which voters become aware that other voters have also received a given piece of information—could produce powerful effects by inducing explicit or tacit voter coordination based on their common knowledge (e.g. Adena et al. 2015; Arias 2016; Arias, Balán, Larreguy, Marshall and Querubín 2017; Yanagizawa-Drott 2014). Explicit discussion of the information may result in voters engaging heavily with the information received, and in turn consolidating their beliefs around such information. This could even result in an agreement to harmonize vote choices. Alternatively, tacit coordination only relies on the delivery of information engendering the (higher-order) belief that others also received the information, and will likely act on such information (Arias, Balán, Larreguy, Marshall and Querubín 2017). Such coordination may in part explain the large effects of local media found in Mexico (Larreguy, Marshall and Snyder 2017). Both such cases of coordinated behavior induce shifts that could not be achieved by providing the same information using private modes of information transmission. We thus pre-registered the following hypothesis:

H3. *The magnitude of all (average and heterogeneous) effects of providing information on the incumbent party's vote share are greater when the information is delivered through a public mechanism.*

3 Experimental design

We conducted a field experiment around the June 7th, 2015 Mexican municipal elections to test these hypotheses. The experiment was conducted in rural and urban precincts across 26 municipalities in the states of Guanajuato, Estado de México, San Luis Potosí, and Querétaro. Municipalities were selected to assure the safety of our implementation team 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 lower (or, more often, higher) than the state average among audited opposition parties. Up to one third of precincts in a given municipal-

ity were then chosen, where priority was given to small but accessible rural precincts and small urban precincts minimizing the number of neighboring precincts included in the experimental sample. The resulting sample of 678 precincts broadly represents Mexico sociodemographically (see Arias, Larreguy, Marshall and Querubín 2017). Arias, Larreguy, Marshall and Querubín (2017) provide specific details of sample selection procedures.

Partnering with the Mexican NGO Borde Político, our baseline treatment disseminated leaflets documenting the results of the ASF audits. As the example from Salamanca, Guanjuato in Figure 1 illustrates, the leaflet explained that FISM funds were intended for social infrastructure projects benefiting the poor, before reporting the total amount of funds (54.1 million pesos) received by the municipal government and the share of those funds (0%) spent on either (but not both) projects not actually benefiting the poor (as in this example) or unauthorized projects (example in Arias, Larreguy, Marshall and Querubín 2017). 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. 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 reached 57% of households in the average precinct. Compliance with our delivery protocols was generally very good.[3]

To investigate how the content and mode of information provision impacts voters, we varied the leaflet along two dimensions corresponding to the hypotheses above. First, to identify the impact of providing voters with a benchmark against which to compare their incumbent party's malfeasance, we also delivered *benchmark* leaflets. In contrast with the *local* leaflet in the second panel of Figure 1, the *benchmark* 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. This local information shows 0% in this example, while the benchmark shows 16%. Figure 2 documents the distribution of malfeasant spending in our sample, showing

[3]Some leaflets were delivered to voters outside the precinct, or adverse weather conditions and poor road conditions prevented us from reaching a precinct. We preserve the randomization by estimating intent to treat effects.

Figure 1: Example of a benchmark information leaflet in Salamanca, Guanajuato

that the average precinct was informed of 21% malfeasant spending within their own municipality and 9% in municipalities within their state governed by other parties.

Second, to vary the extent to which the distribution of the leaflets is common knowledge among voters within the precinct, 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 was accompanied by a powerful portable loudspeaker carried on the back of a team member.[4] 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.

Treatments were randomly assigned within 100 blocks (stratified by rural/urban within a municipality) according to the factorial design with a pure control shown in Table 1. Each block

[4]We purchased these modified rucksack loudspeakers from a vendor in Mexico City that also serves political campaigns similarly seeking to broadcast their message.

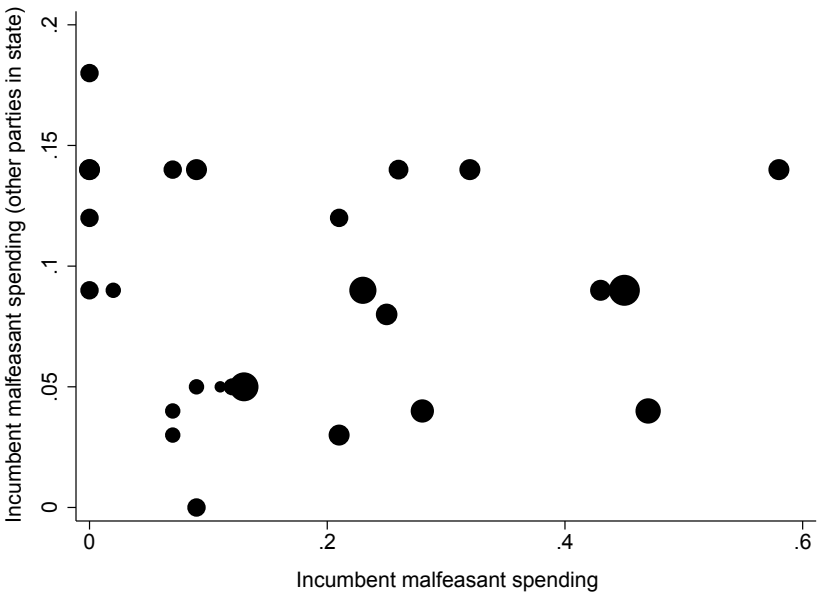

Figure 2: Precincts by share of malfeasant spending in our sample

Notes: 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
Benchmark		100 precincts	100 precincts

contained one precinct receiving each treatment condition, as well as 2 or 3 control precincts (depending on precinct availability). Block randomization ensures that all respondents within a block face the same race and receive information pertaining to the same mayor, and can substantially increase statistical power. Table A1 in the Appendix shows that pre-determined covariates are well-balanced across treatment conditions.

We examine two classes of outcomes. First, we collected precinct-level electoral returns to measure incumbent party vote share. Second, we conducted a survey 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. We also use beliefs in control precincts to proxy for the level and certainty of prior (i.e. pre-treatment delivery) beliefs; we measure the direction and extent of updating by showing respondents the leaflet at the end of our survey and calculating the change in beliefs among control respondents. Arias, Larreguy, Marshall and Querubín (2017) extensively validate this approach.

Whereas Arias, Larreguy, Marshall and Querubín (2017) pool across all information treatment conditions, the goal of this paper is to identify how the variation in the content and form of the leaflets affect voters. Accordingly, we estimate the following baseline regressions specified in our pre-analysis plan:

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

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 β provides estimates of the effect of our information treatment conditions relative to the control, while the randomization block fixed effects η_{bm} adjust for the differential treatment probabilities across blocks arising from different block sizes. Standard errors are clustered by treatment-municipality. The only departure from our pre-analysis plan is that we weight precinct-level observations by the share of voters within the precinct to whom we delivered a leaflet.[5] This permits more precise estimates, although unweighted regressions produce similar results.

To examine how the effects of the treatments vary with prior beliefs and the content of the information provided, we also estimate heterogeneous effects of the form:

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

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

where X_{bm} is a block- or municipality-level moderator (the lower-order term is subsumed by η_{bm}).

4 Results

We first demonstrate that the treatments indeed reached voters, before turning to our main findings identifying the effects of benchmarked information and public dissemination. We also discuss the interpretation of our null findings.

4.1 Manipulation tests

To verify that our treatment reached voters and produced the intended effects, we first assess several manipulation tests in Table 2. Based on our post-treatment survey, Columns (1) and (2) of Table 2 first confirm that voters receiving any treatment were significantly more likely to remember the leaflet and correctly recall the issue discussed in the leaflet. Public and benchmark treatments are slightly more effective in this regard, although the differences are relatively small.

Column (3) shows that voters who received both the local and benchmark treatments were significantly more likely to recall receiving information about other parties in their state than control voters.[6] The effect in precincts receiving only information about their own municipal government indicates fuzzy recall. Nevertheless, voters receiving the benchmarked information were 2 percentage points—or around 40%—more likely to recall receiving information about opposition incumbents in other states. Although relatively small, the t test at the foot of the table indicates that this difference is statistically significant.

Furthermore, Columns (4) and (5) show that the public dissemination treatment also elicited the expected responses. Column (4) confirms that voters receiving the private information treatment were as likely as control respondents to recall a loudspeaker. However, voters in precincts subject

[6]The dependent variable in Column (3) is a dummy variable that indicates whether the respondent reports that the leaflet also included information on other parties in the state.

Table 2: Treatment manipulation checks

	Remember leaflet (1)	Correctly remember content (2)	Remember benchmark states (3)	Remember loud speaker (4)	Share of community received (5)
Private local treatment	0.223***	0.120***			
	(0.026)	(0.021)			
Public local treatment	0.244***	0.135***			
	(0.029)	(0.022)			
Private benchmark treatment	0.239***	0.140***			
	(0.027)	(0.023)			
Public benchmark treatment	0.283***	0.157***			
	(0.023)	(0.021)			
Local treatment			0.046***		
			(0.010)		
Benchmark treatment			0.066***		
			(0.010)		
Private treatment				0.006	0.483***
				(0.007)	(0.062)
Public treatment				0.057***	0.566***
				(0.008)	(0.064)
Outcome range	{0,1}	{0,1}	{0,1}	{0,1}	{1,2,3,4,5}
Control outcome mean	0.09	0.06	0.03	0.03	1.45
Control outcome std. dev.	0.28	0.25	0.18	0.16	1.01
Test: same treatment effect (p value)	0.00	0.06	0.03	0.00	0.06
R^2	0.11	0.10	0.05	0.05	0.09
Observations	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. * denotes $p < 0.1$, ** denotes $p < 0.05$, *** denotes $p < 0.01$.

to public dissemination were six percentage points more likely to correctly recall that the leaflets were delivered accompanied by a loudspeaker.[7] Moreover, the tests at the bottom of Column (5) indicate that voters in such precincts were also significantly more likely to believe that a large

[7]The dependent variable in Column (4) is a dummy variable that indicates whether the respondent recalled listening to a loudspeaker with a recording accompanying the leaflet distribution.

fraction of their community received the leaflets—a key message of the loudspeaker.[8]

4.2 The effect of providing benchmarked malfeasance information

We first examine the effects of providing benchmarked information on our two main outcomes: incumbent vote share as a share of turnout and incumbent vote share as a share of registered voters. We 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 2)—would elicit stronger responses to our information treatment by creating a starker comparison. However, the results in Table 3 indicate that our benchmarked malfeasance information treatment did not accentuate the effect of providing local incumbent malfeasance information.

First, Column (1) shows that benchmarked information did not differentially affect voting behavior on average. The positive estimates are consistent with those in Arias, Larreguy, Marshall and Querubín (2017), who show that the malfeasance reports did not affect the level of posterior beliefs *on average*, but did reduce voter uncertainty about incumbent party malfeasance and elicited responses from incumbent and challenger parties that could have differentially benefited the incumbent party. However, the tests at the foot of Column (1) demonstrate that the effects of local and benchmarked information are statistically indistinguishable. In contrast with hypothesis H1, this suggests that both types of information equally affected voter behavior on average. If anything, the larger effect associated with providing a benchmark on incumbent vote share is surprising, given that the vast majority of treatments informed voters that the challenger was performing better.

Second, we also find little evidence to suggest that the effect of providing benchmarked information systematically varies with the malfeasance information reported (H2). In the case of simply providing local information, Column (2) shows that—as expected—voter reward incum-

[8]The dependent variable in Column 3 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.

Table 3: Effect of local and comparative information treatments on incumbent party vote share

	Incumbent party vote share			
	(1)	(2)	(3)	(4)
Panel A: Incumbent party vote share (share of turnout)				
Local treatment	0.022***	0.041***	0.042***	0.033***
	(0.007)	(0.010)	(0.010)	(0.007)
Comparative treatment	0.031***	0.042***	0.037***	0.031***
	(0.008)	(0.011)	(0.009)	(0.008)
Local treatment × Incumbent malfeasant spending		-0.091*	-0.093**	
		(0.046)	(0.046)	
Comparative treatment × Incumbent malfeasant spending		-0.052		
		(0.036)		
Comparative treatment × Difference in malfeasant spending (incumbent - challenger)			-0.055	
			(0.036)	
Local treatment × Unfavorable incumbent updating				-0.015**
				(0.007)
Comparative treatment × Unfavorable incumbent updating				-0.005
				(0.006)
Outcome range	[0.07,0.85]	[0.07,0.85]	[0.07,0.85]	[0.07,0.85]
Control outcome mean	0.38	0.38	0.38	0.38
Control outcome std. dev.	0.12	0.12	0.12	0.12
Test: same treatment effect (p value)	0.39	0.97	0.73	0.83
Minimum detectable absolute difference	0.020	0.029	0.025	0.018
Test: same interaction effect (p value)		0.56	0.56	0.32
Minimum detectable absolute difference		0.127	0.126	0.019
R^2	0.61	0.62	0.62	0.61
Panel B: Incumbent party vote share (share of registered voters)				
Local treatment	0.012***	0.022***	0.022***	0.019***
	(0.004)	(0.006)	(0.006)	(0.004)
Comparative treatment	0.015***	0.023***	0.020***	0.016***
	(0.005)	(0.007)	(0.005)	(0.005)
Local treatment × Incumbent malfeasant spending		-0.048*	-0.050*	
		(0.026)	(0.026)	
Comparative treatment × Incumbent malfeasant spending		-0.037*		
		(0.020)		
Comparative treatment × Difference in malfeasant spending (incumbent - challenger)			-0.040**	
			(0.020)	
Local treatment × Unfavorable incumbent updating				-0.010***
				(0.003)
Comparative treatment × Unfavorable incumbent updating				-0.004
				(0.003)
Control outcome mean	0.19	0.19	0.19	0.19
Control outcome std. dev.	0.07	0.07	0.07	0.07
Test: same treatment effect (p value)	0.61	0.92	0.78	0.60
Minimum detectable treatment difference	0.012	0.019	0.016	0.014
Test: same interaction effect (p value)		0.78	0.78	0.28
Minimum detectable interaction difference		0.071	0.069	0.010
R^2	0.64	0.64	0.64	0.64
Outcome range	[0.03,0.40]	[0.03,0.40]	[0.03,0.40]	[0.03,0.40]
Observations	675	675	675	651
Interaction mean		0.21	0.22	0.85
Interaction std. dev.		0.17	0.17	1.07

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 minimum detectable absolute difference reflects the difference in effects that could be detected at the 95% level, based on the standard error of the test. Standard errors clustered by municipality-treatment are in parentheses. * denotes $p < 0.1$, ** denotes $p < 0.05$, *** denotes $p < 0.01$.

bent parties less when the leaflet reveals greater malfeasance. If anything, voters are less likely to respond to benchmarked information in such a Bayesian fashion. Although the differences between the interaction coefficients are not statistically significant, Column (2) indicates that voters sanction incumbent parties less for higher rates of malfeasance when presented with benchmarked malfeasance information, while Column (3) shows similar results when we instead interact the benchmark treatment with the difference in malfeasance spending between the incumbent and the average challenger elsewhere in the state. Column (4) also suggests a weaker response to the comparative information treatment by our measure of voter updating of posterior beliefs about the incumbent party among control respondents upon receiving the leaflet at the end of our survey.[9]

The lack of an additional effect associated with providing benchmarked information could reflect several possibilities. First, voters in the control group already believed the main local challenger to be less malfeasant: on our five-point scale (-2 to 2), mean perceived malfeasance about the incumbent (-0.10) was significantly higher than mean perceived malfeasance about the party that placed second in the last election (-0.26) in the control groups. The information that we provided may thus have induced limited updating because it coincided with voters' prior beliefs. Second, voters may have simply failed to comprehend the benchmark component of the treatment. Third, the benchmark itself may simply not have been relevant. For example, Marshall (2017) 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 (e.g. Lieberman, Posner and Tsai 2014), our survey data can help separate between these explanations by examining whether the benchmark treatments differentially affected voter beliefs. First, the estimates in Table A2 of the Appendix show no differential updating about the incumbent party's malfeasance across local and comparative treatments. This indicates that voters understood the information provided about

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

the incumbent, but—consistent with a learning model where performance elsewhere is orthogonal to performance in a voter's own municipality—adding a benchmark did not adjust how voters' updated about their incumbent party.

Second, and more surprisingly, given that the local treatment did not contain information about challengers, Tables A3-A5 show that the benchmarked information consistently does not induce voters to differentially update about challengers. This is robust across three definitions of the challenger in a given municipality. These results cast doubt on the possibility that differential priors explain the null finding because receiving quite different information about the challenger relative to the incumbent does not cause voters to update differently. Rather, the results suggest that voters only responded to the information provided about the incumbent party. This suggests that voters either struggled to comprehend the benchmarked component of the information, which some of our enumerators highlighted as they conducted the survey, or did not believe that the malfeasance of parties in other municipalities represents a good proxy for how such parties would perform in their municipality. Unfortunately, we cannot distinguish between these potential explanations.

4.3 Limited additional impact of public information dissemination

While our information treatment's effects are not affected by providing a performance benchmark, it is possible that the blunter instrument of public transmission may more effectively stimulate voter responses. To investigate this, we examine whether the main findings in Arias, Larreguy, Marshall and Querubín (2017) are magnified when the information is accompanied by a public dissemination device. Panels A and B in Table 4 report the results, reiterating the findings from Arias, Larreguy, Marshall and Querubín (2017) that revealing incumbent party malfeasance is most likely to increase incumbent party vote share where voters possessed unfavorable and imprecise prior beliefs about the incumbent, the ASF report reveals lower levels of malfeasance, and voters updated most favorably about the incumbent from the information.

However, despite increasing common awareness of our leaflets (see Table 2), public dissemi-

Table 4: Effect of private and public information treatments on incumbent party vote share

	Incumbent party vote share				
	(1)	(2)	(3)	(4)	(5)
Panel A: Incumbent party vote share (share of turnout)					
Private information treatment	0.035***	0.032***	-0.007	0.060***	0.040***
	(0.009)	(0.008)	(0.062)	(0.013)	(0.009)
Public information treatment	0.017**	0.015*	0.194**	0.023	0.023**
	(0.008)	(0.008)	(0.075)	(0.014)	(0.009)
Private × Incumbent malfeasance prior		0.010			
		(0.008)			
Public × Incumbent malfeasance prior		0.009			
		(0.008)			
Private × Incumbent prior precision			0.012		
			(0.020)		
Public × Incumbent prior precision			-0.055**		
			(0.023)		
Private × Incumbent malfeasant spending				-0.117***	
				(0.035)	
Public × Incumbent malfeasant spending				-0.029	
				(0.045)	
Private × Unfavorable incumbent updating					-0.010
					(0.006)
Public × Unfavorable incumbent updating					-0.010
					(0.007)
Outcome range	[0.07,0.85]	[0.07,0.85]	[0.07,0.85]	[0.07,0.85]	
Control outcome mean	0.38	0.38	0.38	0.38	0.38
Control outcome std. dev.	0.12	0.12	0.12	0.12	0.12
Test: same treatment effect (p value)	0.14	0.18	0.00	0.10	0.22
Minimum detectable treatment difference	0.023	0.024	0.117	0.044	0.027
Test: same interaction effect (p value)		0.90	0.00	0.16	0.98
Minimum detectable interaction difference		0.021	0.038	0.121	0.019
Panel B: Incumbent party vote share (share of registered voters)					
Private information treatment	0.020***	0.018***	-0.028	0.033***	0.024***
	(0.005)	(0.005)	(0.037)	(0.008)	(0.006)
Public information treatment	0.007	0.006	0.058	0.012	0.011**
	(0.004)	(0.004)	(0.040)	(0.007)	(0.005)
Private × Incumbent malfeasance prior		0.009**			
		(0.004)			
Public × Incumbent malfeasance prior		0.005			
		(0.004)			
Private × Incumbent prior precision			0.014		
			(0.012)		
Public × Incumbent prior precision			-0.016		
			(0.012)		
Private × Incumbent malfeasant spending				-0.063***	
				(0.021)	
Public × Incumbent malfeasant spending				-0.023	
				(0.021)	
Private × Unfavorable incumbent updating					-0.008**
					(0.003)
Public × Unfavorable incumbent updating					-0.006*
					(0.003)
Outcome range	[0.03,0.40]	[0.03,0.40]	[0.03,0.40]	[0.03,0.40]	
Control outcome mean	0.19	0.19	0.19	0.19	0.19
Control outcome std. dev.	0.07	0.07	0.07	0.07	0.07
Test: same treatment effect (p value)	0.09	0.15	0.01	0.10	0.15
Minimum detectable treatment difference	0.014	0.015	0.065	0.025	0.017
Test: same interaction effect (p value)		0.51	0.00	0.19	0.67
Minimum detectable interaction difference		0.011	0.020	0.060	0.010
Observations	675	651	651	675	651
Interaction mean		-0.05	3.24	0.22	0.85
Interaction std. dev.		0.90	0.34	0.17	1.07

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 Columns (2), (3), and (5) reflect the lack of data on prior beliefs about the incumbent party in Apaseo el Alto. The minimum detectable absolute difference reflects the difference in effects that could be detected at the 95% level, based on the standard error of the test. Standard errors clustered by municipality-treatment are in parentheses. * denotes $p < 0.1$, ** denotes $p < 0.05$, *** denotes $p < 0.01$.

nation produced similar—if not weaker—behavioral responses. Column (1) reports a significantly smaller increase in incumbent vote share associated with public dissemination. While this could reflect a less sanguine response to generally high levels of malfeasance, Columns (4) and (5) also document weaker albeit statistically indistinguishable slopes with respect to the level of malfeasance reported and belief updating. Only in the case of the interaction with the precision of voter prior beliefs in Column (3) do we find any evidence that the public treatment elicited a stronger response. In sum, these results provide little evidence to support hypothesis H3.

The limited impact of adding a loudspeaker contrasts with the large effects of media found in similar contexts (Ferraz and Finan 2008; Larreguy, Marshall and Snyder 2017; Marshall 2017). One potential explanation for this 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). Even though the public treatment created greater common knowledge, we find little evidence that this translated into coordination: Columns (1)-(3) in Table 5 find 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 private and public forms of information dissemination.[10]

Second, it is nevertheless possible that the leaflet became common knowledge even with private transmission. There is certainly some evidence that the leaflet itself induced significant coordination: the estimates in Table 5 show that both the private and public leaflet increased discussion and coordination. Moreover, Arias, Balán, Larreguy, Marshall and Querubín (2017) find that the information created greater coordination in precincts with more connected social networks, and that this helped coordinate voters around the candidate that voters believed would be less malfeasant. However, the varied and small scale of such responses, as well as the relatively low recall rates

[10]The dependent variable in Column 1 is a dummy variable for whether the respondent discussed the contents of the leaflets with other members of the community. The dependent variable in Column 2 is a dummy variable for whether, following discussions about the leaflet, people in the community coordinated to vote for the same party. Finally, the dependent variable in Column 3 is a dummy for whether discussions about the leaflet with other members of the community changed the respondent's vote choice.

Table 5: 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
Minimum detectable treatment difference	0.022	0.013	0.012
R^2	0.08	0.07	0.06
Observations	4,958	4,958	4,958

Notes: All specifications include block fixed effects and are estimated using OLS. The minimum detectable absolute difference reflects the difference in effects that could be detected at the 95% level, based on the standard error of the test. Standard errors clustered by municipality-treatment are in parentheses. * denotes $p < 0.1$, ** denotes $p < 0.05$, *** denotes $p < 0.01$.

in Table 2, suggest that it is unlikely that the private transmission had already produced maximal common knowledge.

A third possibility is that the loudspeakers led voters to perceive our intervention as being more partisan. 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, Table A6 reports no evidence that the public treatment increased voter perceptions that the leaflet was delivered by the municipal incumbent, municipal challengers, PAN, PRD, or PRI.

5 Conclusion

This article examines how performance benchmarking and common knowledge moderate the effect of an NGO information campaign in the context of Mexican municipal elections. Leveraging experimental variation in the provision of benchmarks and public dissemination by loudspeaker, we find little evidence that voters differentially updated from the benchmark or coordinated around common knowledge.

The results have various implications for campaign design and future research. First, it is essential for information campaigns to ensure that benchmarks are both easy to comprehend and relevant to voters. One way to achieve this is to extensively pilot the information, which may include experimenting with a variety of benchmarks and approaches to depicting such comparisons. Second, loudspeakers appear to be insufficient to produce additional coordination beyond the small levels of coordination induced by leaflets without a public coordinating device. 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 2017; Marshall 2017; 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 the mechanisms driving the large effects of the media. Future research could also examine the coordination potential of candidate meetings, public meetings, loudspeakers attached to cars, etc. Third, our findings nevertheless demonstrate that leaflets can be effective, underscoring the importance of clear information about incumbent performance.

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.

Adida, Claire, Jessica Gottlieb, Eric Kramon and Gwyneth McClendon. 2016. "Breaking the Clientelistic Voting Equilibrium: The Joint Importance of Salience and Coordination." Working paper.

Arias, Eric. 2016. "How Does Media Influence Social Norms? A Field Experiment on the Role of Common Knowledge." Working paper.

Arias, Eric, Horacio A. Larreguy, John Marshall and Pablo Querubín. 2017. "Priors rule: When do malfeasance revelations help or hurt incumbent parties?" Working paper.

Arias, Eric, Pablo Balán, Horacio A. Larreguy, John Marshall and Pablo Querubín. 2017. "Malfeasance Revelations, Coordination Through Social Networks and Electoral Sanctioning: Experimental Evidence from Mexico." Working paper.

Auditoría Superior de la Federación. 2014. "Informe del Resultado de la Fiscalización Superior de la Cuenta Pública 2012." Audit Summary Report.

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

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

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. 2013. "When Do Voters Punish Corrupt Politicians? Experimental Evidence from Brazil." Working paper.

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

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.

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. 1999. "Managerial Incentive Problems: A Dynamic Perspective." *Review of Economic Studies* 66(1):169–182.

Humphreys, Macartan and Jeremy Weinstein. 2012. "Policing Politicians: Citizen Empowerment and Political Accountability in Uganda Preliminary Analysis." Working paper.

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

Larreguy, Horacio A., John Marshall and Jr. Snyder, James M. 2017. "Publicizing Malfeasance: How Local Media Faciliates Electoral Sanctioning of Mayors in Mexico." NBER Working Paper 20697.

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.

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. 2017. "Political information cycles: When do voters sanction incumbent parties for high homicide rates?" Working paper.

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.

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

A Appendix

Contents

A.1 Additional results

Tables A1-A6 report additional results cited in the main paper.

Table A1: Balance across 40 precinct-level variables

	Control		Private local treatment		Public local treatment		Private benchmark treatment		Public benchmark treatment			Test: all treatment effects
	Mean	Std. dev.	Coefficient	Std. error	Coefficient	Std. error	Coefficient	Std. error	Coefficient	Std. error	Observations	=0 (p value)
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

Notes: All specifications include block fixed effects, 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 A2: 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)	(4)	(5)	(6)
Local treatment	0.002	-0.027	0.752*	0.029	0.004	-0.156**
	(0.044)	(0.051)	(0.378)	(0.075)	(0.058)	(0.066)
Comparative treatment	-0.002	-0.031	1.040***	-0.055	0.016	-0.172**
	(0.046)	(0.053)	(0.379)	(0.078)	(0.063)	(0.066)
Local treatment × Incumbent malfeasance prior		-0.246***				
		(0.048)				
Comparative treatment × Incumbent malfeasance prior		-0.305***				
		(0.045)				
Local treatment × Incumbent prior precision			-0.231*			
			(0.118)			
Comparative treatment × Incumbent prior precision			-0.321***			
			(0.120)			
Local treatment × Incumbent malfeasant spending				-0.308		
				(0.894)		
Comparative treatment × Incumbent malfeasant spending				0.595		
				(0.938)		
Local treatment × Difference in malfeasant spending					-0.017	
					(0.227)	
Comparative treatment × Difference in malfeasant spending					-0.147	
					(0.241)	
Local treatment × Unfavorable incumbent updating						0.170***
						(0.042)
Comparative treatment × Unfavorable incumbent updating						0.187***
						(0.041)
Outcome range	{-2,-1,0,1,2}	{-2,-1,0,1,2}	{-2,-1,0,1,2}	{-2,-1,0,1,2}	{-2,-1,0,1,2}	{-2,-1,0,1,2}
Control outcome mean	-0.14	-0.14	-0.14	-0.14	-0.14	-0.14
Control outcome std. dev.	1.48	1.48	1.48	1.48	1.48	1.48
Interaction mean		-0.08	3.24	0.21	0.21	0.89
Interaction std. dev.		0.89	0.37	0.17	0.17	1.07
Test: same treatment effect (p value)	0.93	0.94	0.47	0.37	0.83	0.79
Minimum detectable treatment difference	0.085	0.085	0.776	0.183	0.111	0.121
Test: same interaction effect (p value)						
Minimum detectable interaction difference		0.094	0.240	1.909	0.424	0.086
R^2	0.29	0.30	0.29	0.29	0.29	0.30
Observations	4,635	4,635	4,629	4,635	4,635	4,635

Notes: All specifications include block fixed effects, and are estimated using OLS. Lower-order interaction terms are absorbed by the block fixed effects. The minimum detectable absolute difference reflects the difference in effects that could be detected at the 95% level, based on the standard error of the test. Standard errors clustered by municipality-treatment are in parentheses. * denotes $p < 0.1$, ** denotes $p < 0.05$, *** denotes $p < 0.01$.

Table A3: 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)	(4)	(5)	(6)
Local information treatment	-0.005	-0.106**	-0.178	-0.034	-0.030	-0.122***
	(0.037)	(0.047)	(0.348)	(0.073)	(0.047)	(0.044)
Comparative information treatment	-0.054	-0.162***	-0.410	-0.110	-0.017	-0.211***
	(0.041)	(0.047)	(0.345)	(0.080)	(0.056)	(0.044)
Local × Challenger malfeasance prior		-0.555***				
		(0.077)				
Comparative × Challenger malfeasance prior		-0.619***				
		(0.072)				
Local × Challenger prior precision			0.056			
			(0.109)			
Comparative × Challenger prior precision			0.115			
			(0.110)			
Local × Challenger malfeasant spending				0.319		
				(0.788)		
Comparative × Challenger malfeasant spending				0.627		
				(0.911)		
Local × Difference in malfeasant spending (incumbent - challenger)					0.195	
					(0.196)	
Comparative × Difference in malfeasant spending (incumbent - challenger)					-0.289	
					(0.222)	
Local × Unfavorable challenger updating						0.178***
						(0.049)
Comparative × Unfavorable challenger updating						0.249***
						(0.050)
Outcome range	{-2,-1,0,1,2}	{-2,-1,0,1,2}	{-2,-1,0,1,2}	{-2,-1,0,1,2}	{-2,-1,0,1,2}	{-2,-1,0,1,2}
Control outcome mean	-0.19	-0.19	-0.19	-0.19	-0.19	-0.19
Control outcome std. dev.	1.30	1.30	1.30	1.30	1.30	1.30
Interaction mean		-0.08	3.22	0.21	0.21	0.89
Interaction std. dev.		0.87	0.39	0.17	0.17	1.05
Test: same treatment effect (p value)	0.26	0.19	0.49	0.40	0.83	0.03
Minimum detectable treatment difference	0.085	0.082	0.653	0.176	0.112	0.080
Test: same interaction effect (p value)						
Minimum detectable interaction difference		0.150	0.207	2.019	0.417	0.089
R^2	0.08	0.09	0.08	0.08	0.09	0.09
Observations	4,958	4,958	4,958	4,958	4,958	4,958

Notes: All specifications include block fixed effects, and are estimated using OLS. Lower-order interaction terms are absorbed by the block fixed effects. The minimum detectable absolute difference reflects the difference in effects that could be detected at the 95% level, based on the standard error of the test. Standard errors clustered by municipality-treatment are in parentheses. * denotes $p < 0.1$, ** denotes $p < 0.05$, *** denotes $p < 0.01$.

Table A4: 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)	(4)	(5)	(6)
Local information treatment	0.027	-0.075	-0.527	-0.008	0.007	-0.086*
	(0.042)	(0.057)	(0.460)	(0.077)	(0.052)	(0.049)
Comparative information treatment	-0.042	-0.144**	-0.600	-0.069	-0.018	-0.165***
	(0.043)	(0.056)	(0.420)	(0.079)	(0.060)	(0.046)
Local × Challenger malfeasance prior		-0.370***				
		(0.054)				
Comparative × Challenger malfeasance prior		-0.391***				
		(0.062)				
Local × Challenger prior precision			0.177			
			(0.145)			
Comparative × Challenger prior precision			0.179			
			(0.135)			
Local × Challenger malfeasant spending				0.388		
				(0.960)		
Comparative × Challenger malfeasant spending				0.297		
				(1.024)		
Local × Difference in malfeasant spending (incumbent - challenger)					0.154	
					(0.228)	
Comparative × Difference in malfeasant spending (incumbent - challenger)					-0.186	
					(0.257)	
Local × Unfavorable challenger updating						0.151***
						(0.036)
Comparative × Unfavorable challenger updating						0.171***
						(0.043)
Outcome range	{-2,-1,0,1,2}	{-2,-1,0,1,2}	{-2,-1,0,1,2}	{-2,-1,0,1,2}	{-2,-1,0,1,2}	{-2,-1,0,1,2}
Control outcome mean	-0.30	-0.30	-0.31	-0.30	-0.30	-0.30
Control outcome std. dev.	1.36	1.36	1.37	1.36	1.36	1.36
Interaction mean		-0.08	3.23	0.21	0.21	0.89
Interaction std. dev.		0.87	0.38	0.17	0.17	1.05
Test: same treatment effect (p value)	0.08	0.07	0.85	0.30	0.57	0.05
Minimum detectable treatment difference	0.075	0.074	0.770	0.113	0.088	0.078
Test: same interaction effect (p value)						
Minimum detectable interaction difference		0.092	0.247	1.588	0.420	0.060
R^2	0.19	0.19	0.19	0.19	0.19	0.19
Observations	4,958	4,958	4,908	4,958	4,958	4,958

Notes: All specifications include block fixed effects, and are estimated using OLS. Lower-order interaction terms are absorbed by the block fixed effects. The smaller sample in Column (3) reflects a lack of data on prior beliefs about the challenger in Tamasopo. The minimum detectable absolute difference reflects the difference in effects that could be detected at the 95% level, based on the standard error of the test. Standard errors clustered by municipality-treatment are in parentheses. * denotes $p < 0.1$, ** denotes $p < 0.05$, *** denotes $p < 0.01$.

Table A5: 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)	(4)	(5)	(6)
Local information treatment	0.032	-0.051	-0.405	-0.115**	0.064	-0.075**
	(0.037)	(0.045)	(0.401)	(0.056)	(0.044)	(0.037)
Comparative information treatment	-0.014	-0.097**	-0.528	-0.121**	0.033	-0.124***
	(0.039)	(0.046)	(0.373)	(0.060)	(0.051)	(0.038)
Local × Challenger malfeasance prior		-0.280***				
		(0.052)				
Comparative × Challenger malfeasance prior		-0.294***				
		(0.056)				
Local × Challenger prior precision			0.136			
			(0.125)			
Comparative × Challenger prior precision			0.159			
			(0.118)			
Local × Challenger malfeasant spending				1.631**		
				(0.721)		
Comparative × Challenger malfeasant spending				1.191		
				(0.815)		
Local × Difference in malfeasant spending (incumbent - challenger)					-0.266	
					(0.226)	
Comparative × Difference in malfeasant spending (incumbent - challenger)					-0.377	
					(0.236)	
Local × Unfavorable challenger updating						0.140***
						(0.033)
Comparative × Unfavorable challenger updating						0.148***
						(0.037)
Outcome range	{-2,-1,0,1,2}	{-2,-1,0,1,2}	{-2,-1,0,1,2}	{-2,-1,0,1,2}	{-2,-1,0,1,2}	{-2,-1,0,1,2}
Control outcome mean	-0.33	-0.33	-0.33	-0.33	-0.33	-0.33
Control outcome std. dev.	1.20	1.20	1.20	1.20	1.20	1.20
Interaction mean		-0.08	3.22	0.21	0.21	0.89
Interaction std. dev.		0.87	0.39	0.17	0.17	1.05
Test: same treatment effect (p value)	0.17	0.18	0.61	0.90	0.50	0.19
Minimum detectable treatment difference	0.064	0.067	0.463	0.099	0.091	0.072
Test: same interaction effect (p value)						
Minimum detectable interaction difference		0.080	0.143	1.328	0.363	0.056
R^2	0.30	0.30	0.30	0.30	0.30	0.30
Observations	4,958	4,958	4,958	4,958	4,958	4,958

Notes: All specifications include block fixed effects, and are estimated using OLS. Lower-order interaction terms are absorbed by the block fixed effects. The minimum detectable absolute difference reflects the difference in effects that could be detected at the 95% level, based on the standard error of the test. Standard errors clustered by municipality-treatment are in parentheses. * denotes $p < 0.1$, ** denotes $p < 0.05$, *** denotes $p < 0.01$.

Table A6: 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
R^2	0.04	0.06	0.05	0.04	0.05
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 $p < 0.01$.