ONLINE POLITICAL INFORMATION: FACEBOOK ADS, ELECTORATE SATURATION, AND ELECTORAL ACCOUNTABILITY IN MEXICO*

JOSÉ RAMÓN ENRÍQUEZ† HORACIO LARREGUY‡ JOHN MARSHALL§

ALBERTO SIMPSER¶

OCTOBER 2019

While inexpensive digital technologies like Facebook can spread misinformation, they could also enhance electoral accountability. We experimentally study voter responses to incumbent performance information disseminated through social media and how these responses vary with information campaign *saturation*—the share of an electorate targeted. We evaluate a non-partisan NGO's campaign that used Facebook ads to inform Mexican citizens about the extent of irregularities in audited government expenditures in their municipality prior to the 2018 general elections. The information campaign was randomized to target 0%, 20%, or 80% of a municipality's electorate. Around 15% of targeted voters watched at least part of the Facebook video ad. We find that incumbent parties which engaged in negligible irregularities received around 5 percentage points more votes among citizens directly targeted by the ads. This effect was twice as large in the 80% saturation municipalities than the 20% saturation municipalities, while the higher saturation condition also generated substantial spillovers to non-targeted voters within the same municipality. Social interactions between voters, rather than responses by politicians or media outlets, appear to drive both the direct and spillover effects. Information campaign saturation may thus constitute one of the keys to explaining the relatively large impact of information disseminated by broadcast and digital mass media on voting behavior and electoral accountability.

*We thank Daron Acemoglu, David Broockman, Bruce Buena de Mesquita, Chappell Lawson, Craig McIntosh, Elias Papaioannou, Maria Petrova, David Yang, and participants at the CESS of Oxford in Chile, EGAP, Harvard, ITAM, MIT GOV/LAB, the NYU-CESS Conference, UVA, and Wisconsin for their feedback and comments in various stages of the project. We are extremely grateful to Claudia Cervantes Pérez, Ricardo de la Teja, and Christian Aarón Ramírez García for excellent research assistance and GeoPoll for their survey implementation. This research was implemented in collaboration with Borde Político, and was supported by the Dean's Competitive Fund for Promising Scholarship of the Faculty of Arts and Sciences at Harvard University and the Foundations of Human Behavior Initiative. The project was approved by the Columbia University Human Research Protection office (AAAR8909), the Harvard Committee on the Use of Human Subjects (18-0743), and ITAM's Institutional Review Board. The pre-analysis plan was pre-registered at the AEA RCT Registry and EGAP, and is publicly available at https://www.socialscienceregistry.org/trials/3135 and http://egap.org/registration/4806.

†Department of Government, Harvard University. Email: jrenriquez@g.harvard.edu.

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

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

¶Department of Political Science, Instituto Tecnológico Autónomo de México. Email: alberto.simpser@itam.mx.

1 Introduction

Advances in digital communication technologies have created new opportunities for targeting information toward large numbers of citizens at limited cost. These advances have been particularly dramatic in developing contexts, where the use of internet and social media platforms—which are primarily accessed via cell phones—is rapidly catching up to levels in developed countries and has grown by more than 50% within the last 5 years (Poushter, Bishop and Chwe 2018). The growing availability of these technologies is revolutionizing access to politically-relevant information and democratizing who can provide such information (e.g. Acemoglu, Hassan and Tahoun 2017; Enikolopov, Petrova and Makarin 2018; Fergusson and Molina 2019; Howard 2010; Lynch 2011; Manacorda and Tesei 2018; Miner 2015).[1]

While the potential for partisan actors to manipulate or distract masses of voters with fake news and government propaganda are critical concerns for electoral accountability, the digital revolution also presents unprecedented opportunities to increase electoral accountability. By disseminating credible information about government performance in office, without needing to rely on under-resourced traditional media outlets that are often vulnerable to political capture (Anderson and McLaren 2012; Baron 2006; Besley and Prat 2006), non-partisan actors can enhance voter selection and control of elected representatives (e.g. Barro 1973; Fearon 1999; Ferejohn 1986; Rogoff 1990). This potential for facilitating electoral accountability—the focus of our article—is especially important in developing contexts where politician malfeasance and low-quality public goods provision are especially common (e.g. Khemani et al. 2016).

A key feature of online communication technologies like Facebook and Twitter, as well as some broadcast media, is the capacity for information campaigns to reach large numbers of voters within a given electoral unit. Focusing on the potential benefits of the digital revolution, we argue that information campaign *saturation*—which we define as the share of an electorate with direct access to a campaign—can induce and amplify the impacts of online information campaigns on electoral accountability by facilitating interactions between voters.[2] Even among voters with access to information, campaign saturation is likely to: (i) increase the likelihood of belief updating in response to incumbent performance information, by increasing information diffusion between voters or inducing campaign responses and media reporting that in turn increases the probability of voters engaging with the information made available to them; and/or (ii) coordinate voting on the basis of incumbent performance in office—rather than competing influences on vote choices, such

[1] See Zhuravskaya, Petrova and Enikolopov (forthcoming) for an overview of this literature.

[2] Common alternative conceptions of saturation address the number of times that a given piece of information is received or the point at which information could be absorbed no further. However, to fix terminology, this article focuses on saturation in terms of its coverage across voters within a given electorate. In this sense, we adhere to the conception of saturation as the degree to which something is absorbed (as in Baird et al. 2018).

as clientelism or shared identity—by generating explicit communication and agreements (Chwe 2000; Larson 2017) and common knowledge between voters (Morris and Shin 2002). We thus expect the magnitude of any effect of information campaigns documenting incumbent performance in office on electoral support for the incumbent to increase in the campaign's saturation.

This article leverages a field experiment to estimate the effects of a large-scale online information campaign, and variation in its electorate-level saturation, on electoral accountability during the 2018 Mexican municipal elections. In particular, we evaluate the impact of a non-partisan campaign by Borde Político—a Mexican NGO that seeks to promote government transparency using digital tools—that used Facebook ads to inform voters of federal transfers to municipalities intended for social infrastructure projects benefiting the poor and the share of the municipal incumbent's expenditures under this program that were subject to irregularities. Mexico's independent Federal Auditor's Office (ASF) defines irregularities as funds spent on unauthorized projects or social infrastructure projects that do not benefit the law's intended recipients, and often constitute corruption (Chong et al. 2015; Larreguy, Marshall and Snyder 2019). This information was extracted from the ASF's publicly-available audit reports and disseminated via 26-second paid-for video ads in the week preceding the election. Corruption was a highly salient issue in the 2018 Mexican elections, in which anti-establishment presidential candidate Andrés Manuel López Obrador and his MORENA party won by a landslide across the federal and local levels.

In collaboration with Borde Político, we randomized whether their Facebook ad campaign targeted 0%, 20%, or 80% of the electorate in the 128 municipalities comprising our sample. Within the treated municipalities, we then randomized the targeting of Facebook ads across multiples of 5 segments (groups of contiguous electoral precincts) in accordance with the 20% and 80% saturation levels. Consequently, 4 (1) of every 5 segments within high (low) saturation municipalities were directly targeted with Facebook ads. This randomized saturation design identifies (i) the direct effect of access to the campaign within a given segment, (ii) the indirect—or "spillover"—effect of the campaign in untreated segments of treated municipalities, and (iii) how either segment-level effect varies with municipal saturation (see Baird et al. 2018). According to Facebook's ad campaign data, the ads ultimately reached 2.7 million unique Facebook users (appearing 3 times per person, on average) and resulted in around 15% of the targeted adult population—or about 20% of targeted Facebook users—watching at least 3 seconds of the ad.[3] Engagement with the campaign was broadly proportionate with the level of access prescribed by the campaign saturation level. A parallel panel survey shows that voters comprehended and retained the information provided by the

[3]Our Facebook analytics data can only distinguish whether users watched the ad at all, for at least 3 seconds, for at least 10 seconds, or entirely. Since viewers were informed of the ad's topic area at the outset and the level of irregularities in their municipality was reported just over halfway through the ad, and the ad allowed Facebook users to click through to access a Facebook page that showed the level of irregularities, we consider watching at least 3 seconds of an ad as the most appropriate measure of the campaign's reach.

video.

Our results show that this large-scale digital information campaign significantly affected polling station-level voting behavior. First, relative to pure control segments, the best-performing incumbent parties—those whose citizens were informed of zero or negligible levels of irregularities—increased their vote share among registered voters in the average segment that was directly targeted by Facebook ads by 4-5 percentage points, or almost half a standard deviation. The vote share of incumbent parties that presided over irregularities in the third quartile of the distribution was unaffected. Incumbent parties in the worst-performing quartile suffered a 1-2 percentage loss of votes in targeted segments, although this statistically insignificant effect may have been limited by risk averse voters becoming less uncertain about the incumbent party's type (Arias, Larreguy, Marshall and Querubín 2019) or the anti-incumbent wave of support for MORENA. Since a municipality's level of irregularities is not randomly assigned, we show that our results are robust to adjusting for treatment $\times$ covariate interactions and various codings of low and high malfeasance. While Borde Político's campaign increased turnout by around 1 percentage point across all treated segments, the changes in incumbent party vote share for the best-performing incumbent parties are not driven by aggregate shifts in voter mobilization.

Second, we further demonstrate that these effects in directly targeted segments are largely driven by municipalities that received the higher saturation information campaign. For the least malfeasant incumbent parties, the incumbent party's vote share increase by 5-6 percentage point in treated segments within the high saturation municipalities where 80% of voters were targeted. In contrast, 2-3 percentage point effect in treated segments within the 20% saturation condition is significantly smaller thanin the high saturation municipalities.[4] Since the level of engagement with the Facebook ads was similar across treated segments in the 20% and 80% saturation municipalities, these results imply a strong complementarity between access to Borde Político's Facebook information campaign and a substantial share of other Facebook users within the same municipality also having access to the campaign. Our estimates suggest, at least between our 20% and 80% targeted saturation levels, that this complementarity does not exhibit increasing returns in the saturation level.

We further show that the effectiveness of high saturation Facebook ad campaigns is driven by interactions between voters, rather than substantial persuasion among individuals viewing the ads in isolation or politician reactions and/or media reporting induced by Borde Político's ad campaign. Consistent with descriptive data indicating that information diffusion between voters and explicit and tacit voter coordination are common before Mexico elections, the vote share of the least malfeasant incumbent parties also increased by around 5 percentage points in untreated seg-

[4]The average effect across saturation levels is closer to that in high saturation municipalities because there are fewer directly treated segments in low saturation municipalities.

ments within high saturation municipalities, whereas we fail to detect such spillover effects on untreated segments in low saturation municipalities registering similar levels of malfeasance. This suggests that social interactions induced by the campaign account for most of the effect in directly-treated segments.[5] In contrast, we find no evidence to suggest that other potential amplification mechanisms—specifically, online political responses or media reporting—could account for the results. Although our design does not allow us to distinguish whether changes in voting behavior reflect individual belief updating induced by information sharing between voters or voter coordination (or both), these results indicate that subsequent social interactions are key mechanisms in enabling high saturation campaigns to amplify the effects of mass online information campaigns.

This study makes several main contributions. First, we highlight the substantial potential for electoral impact of information disseminated via social media platforms during election campaigns. Our estimates suggest that Facebook may play a greater role in persuading voters to change who to vote for than in convincing citizens to turn out. Although we observe small increases in turnout comparable to Facebook's own "get out the vote" campaigns (Bond et al. 2012; Jones et al. 2017), our evaluation of Borde Político's provision of non-partisan information pertaining to incumbent performance further shows that mass campaigns can far more substantially influence vote choice. These large effects align with recent evidence that Facebook polarizes political attitudes (Allcott et al. 2019) and that micro-targeted ads significantly increased self-reported support for Donald Trump in the 2016 U.S. presidential election (Liberini et al. 2018). Our results also align with a non-experimental literature examining the effects of the internet more generally, which finds growing political impacts of the internet over time (Campante, Durante and Sobbrio 2017) and that access to the internet helps voters to hold their governments to account (Miner 2015). However, our design overcomes the formidable challenges in isolating the effects of social media content *and* penetration that observational studies have struggled to distinguish (Zhuravskaya, Petrova and Enikolopov forthcoming).

Second, our findings nevertheless counterbalance concerns that have rightly been raised about how fake news disseminated via social media may have shaped vote choices in the U.S. (e.g. Allcott and Gentzkow 2017) and developing democracies like Brazil, India, and Nigeria.[6] In contrast, we provide the first evidence that factual and non-partisan information disseminated via Facebook's low-cost ads can also promote electoral accountability. Like radio-based anti-vote buying campaigns that reduced support for India's more clientelistic parties (Vasudevan 2019) and election-related information on the radio that increased electoral competition in the U.S. (Panagopoulos and Green 2008), the increased support that we observe for incumbent parties that are ostensibly less corrupt will most likely increase voter welfare in Mexico. In light of growing efforts to regulate

[5] We find no indication that the substantial indirect effects reflect inaccuracy in Facebook's spatial targeting of ads.

[6] India's 2019 election was even described as its "WhatsApp election" by the *Financial Times* ("India: the WhatsApp election," May 5, 2019).

social media during elections (e.g. in India and Turkey), in response to widely-circulated fake news, our findings show that democratizing control of content provision can also support electoral accountability.

Third, we demonstrate that the impact of incumbent performance information disseminated via social media is, to a substantial degree, causally driven by the saturation of the information campaign at the electorate level. While Adida et al. (forthcoming) also experimentally varied the saturation of an accountability campaign at the electorate level in Benin, their greatest level of saturated targeting (15%) is substantially lower than in Borde Político's campaign.[7] Moreover, Adida et al. (forthcoming) find that saturation principally amplified the effect of civics training, rather than the effect of the incumbent performance information provided (via video) alongside such in-person training. The stark differences that we find across the 20% and 80% information campaign saturation levels suggest that saturationcould can account for the notable heterogeneity in the treatment effects of informational interventions (see Dunning et al. forthcoming). In particular, low saturation information campaigns that reached few other voters have had little impact on treated individuals' vote choices (Adida et al. forthcoming; Boas, Hidalgo and Melo 2019; Dunning et al. forthcoming; Humphreys and Weinstein 2012; Lierl and Holmlund forthcoming). Conversely, interventions that instead made information accessible to large shares of the electorate, via the media (e.g. Banerjee et al. 2011; Ferraz and Finan 2008; Larreguy, Marshall and Snyder 2019; Marshall 2019*b*) or concentrated leafleting (Arias, Larreguy, Marshall and Querubín 2019; Cruz et al. 2019), report greater electoral rewards and punishment on the basis of incumbent performance.[8] Our results thus provide causal evidence that saturation—a defining characteristic of broadcast and print, as well as digital, media—may be a key driver of the larger impacts of information delivered by mass media. Other recent studies similarly suggest that network-level saturation of incentives to protest increase sustained protest participation (Bursztyn et al. 2019).

Fourth, in providing evidence suggesting that interactions between voters drive saturation's effects, we further highlight how social interactions can amplify the effects of information on electoral accountability. This finding complements recent evidence suggesting that information campaigns have diffused within communities or households to influence vote choices (Bhandari, Larreguy and Marshall 2019) and turnout (Fafchamps, Vaz and Vicente Forthcoming; Nickerson 2008), coordinated voters around better candidates (Arias, Balán, Larreguy, Marshall and Querubín 2019), and produced larger effects when widespread campaigns were common knowledge (George,

[7] Buntaine et al. (2018) have also experimentally varied the *village* level saturation of a similar accountability campaign in Uganda, but do not vary saturation at the *electorate* level we focus on and—perhaps unsurprisingly—fail to detect differential saturation effects.

[8] Appendix section A.1 thoroughly describes how prior studies vary in terms of information campaign saturation. However, none of these studies exogenously varied high degrees of saturation. The main exception to the correlation between information campaign saturation and effect magnitude is Bhandari, Larreguy and Marshall (2019), where the information diffusion that resulted from a very low scale campaign was substantial in rural Senegal.

Gupta and Neggers 2018). Beyond elections, communication between citizens also appears to have stimulated protest and collective action (Acemoglu, Hassan and Tahoun 2017; Enikolopov, Petrova and Makarin 2018; Fergusson and Molina 2019; García-Jimeno, Iglesias and Yildirim 2018; Manacorda and Tesei 2018; Pierskalla and Hollenbach 2013; Steinert-Threlkeld 2017). In contrast with these studies, our randomized saturation design enables us to identify spillover and saturation effects to show that social effects are particularly prominent in high saturation campaigns.

Finally, our finding that a non-partisan NGO campaign can influence vote choices relates to a broader partisan persuasion literature. Across various contexts, campaign ads of political parties have proven effective at winning votes (Da Silveira and De Mello 2011; Gerber et al. 2011; Larreguy, Marshall and Snyder 2018; Spenkuch and Toniatti 2018) and debates between candidates have increased support for the best-performing candidates (Bidwell, Casey and Glennerster 2019; Bowles and Larreguy 2018; Platas Izama and Raffler 2018). Similarly, news content has persuaded voters to switch parties (Adena et al. 2015; DellaVigna and Kaplan 2007; DellaVigna, Enikolopov, Mironova, Petrova and Zhuravskaya 2014; Enikolopov, Petrova and Zhuravskaya 2011; Gentzkow, Shapiro and Sinkinson 2011; Martin and Yurukoglu 2017). We observe larger reduced form effects of non-partisan independent audit information than most of these studies do for partisan mass media, suggesting that maintaining a role for transparency-oriented NGOs may be critical in improving electoral accountability in developing contexts.

This article is structured as follows. Section 2 discusses potential mechanisms for a saturation effect. Sections 3 and 4 describe the context and experimental evaluation of Borde Político's Facebook ad campaign. Section 5 then describes the campaign's reach, before we report our main results in section 6 and explore mechanisms in section 7. Section 8 concludes.

2 How information campaign saturation could affect electoral accountability

Theories of electoral accountability posit that information about an incumbent's performance in office can help voters to better select and control elected representatives. First, incumbent performance information can mitigate adverse selection problems by helping voters to identify politicians likely to perform competently or pursue policies aligned with their interests (Fearon 1999; Rogoff 1990). Second, incumbent performance information can help voters to replace shirking or corrupt politicians, which in turn establishes incentives for future incumbents (Barro 1973; Ferejohn 1986). These theories ultimately predict that providing relevant and novel incumbent performance information will induce voters to sanction and reward politicians at the ballot box, especially where information deviates from voters' prior beliefs (see e.g. Arias, Larreguy, Marshall and Querubín

2019; Banerjee et al. 2011; Cruz et al. 2019; Kendall, Nannicini and Trebbi 2015) or serves to coordinate voters (e.g. Arias, Balán, Larreguy, Marshall and Querubín 2019; Larson 2017; Morris and Shin 2002).

In light of the mixed evidence that information campaigns help voters reward (sanction) better (worse) performing incumbents discussed earlier (see Appendix section A.1 for further details), and the growing degree to which political information disseminated through social media is reaching substantial shares of the electorate, we consider the extent to which the *saturation* of an information campaign influences electoral accountability. The proportion of voters within a given electoral unit that can access incumbent performance information could induce or amplify the effects of providing such information on electoral accountability through two primary mechanisms: information diffusion and voter coordination. While the former mechanism captures an increase in the likelihood that voters update their beliefs, the latter captures an increase in the likelihood of coordinated voting behavior for any given posterior belief about the incumbent's type or effort.

Most straightforwardly, saturation could amplify an information campaign's effect among those with access to the campaign by increasing the probability that information ultimately reaches—and is internalized by—its targets. While voters often ignore, or only cursorily view, pamphlets, text messages, broadcast media programming, or online content providing unsolicited politically-relevant information (e.g. Dunning et al. forthcoming), saturation could increase engagement through several channels. First, where citizens discuss political issues with each other, campaign saturation is likely to increase exposure to information about incumbent performance. This could result from voters directly sharing their information with others (e.g. Alatas et al. 2016; García-Jimeno, Iglesias and Yildirim 2018), indirectly transmitting the information through social networks or digital interactions (e.g. Alt et al. 2019; Bhandari, Larreguy and Marshall 2019; Buntaine et al. 2018), or encouraging others to acquire political information (e.g. DellaVigna, List, Malmendier and Rao 2014; Marshall 2019*a*).

Second, information campaigns conducted at high levels of saturation may induce responses from political parties or media outlets that could retransmit the information to large audiences, and thereby generate greater belief updating. Saturation could increase the likelihood that parties and media outlets become aware of the information or learn that it is of interest to the voters or consumers that they seek to attract. Consequently, to the extent to which an information campaign disseminates relevant content capable of influencing a citizen's capacity to hold politicians accountable, greater saturation is expected to amplify the campaign's effects on treated targets by increasing their probability of engaging with the information.

Greater saturation could also solve coordination problems that discourage voters from selecting certain types of politicians or responding to information provided at a low scale. This is a common challenge in developing contexts, where clientelistic equilibria (e.g. Adida et al. forthcoming;

Arias, Balán, Larreguy, Marshall and Querubín 2019) or equilibria benefiting political elites (e.g. Myerson 1999) are hard to break away from because voters do not believe that others will also pursue the costly behavior required for change. A coordination stimulus would likely coordinate voters around welfare-improving equilibria, where better politicians are selected.[9] Accountability campaigns focused on valence issues like mayoral malfeasance may make such issues focal and thereby facilitate voter coordination around less corrupt candidates.

High saturation information provision could induce voter coordination through at least two channels. The first channel involves increasing the probability of explicit coordination, whereby information provision sparks voters to get together—either in person or via modern communication technologies—and reach implicit understandings or explicit agreements to synchronize their behavior through communication (Arias, Balán, Larreguy, Marshall and Querubín 2019; Chwe 2000; Larson 2017). Such explicit communication could emerge directly in response to any discussion of the information provided, or because providing information indirectly stimulates political agreement (e.g. arising from a greater general interest in politics or issues of government performance) that would not have occurred otherwise. Arias, Balán, Larreguy, Marshall and Querubín (2019) provide evidence indicating that an information campaign in Mexico induced explicit coordination that influenced vote choices. Beyond voting behavior, Acemoglu, Hassan and Tahoun (2017) and Lynch (2011) further argue that social media similarly facilitated protests during the Arab Spring, while Enikolopov, Petrova and Makarin (2018) and Fergusson and Molina (2019) respectively provide more concrete evidence of this from Russia and across the globe. Manacorda and Tesei (2018) show that cell phones can play a similar role in Africa.

A second, and related, tacit coordination channel relies instead exclusively on common knowledge, such that—even without explicit interpersonal communication—a high saturation information campaign may lead voters to believe that many other voters also received the same information. Such public signals can facilitate coordinated behavior when citizens seek to match the actions of others (Morris and Shin 2002), e.g. by voting for the candidate or party that voters believe to be better or signaling their dissatisfaction with the political system. Consistent with a significant role for common knowledge, George, Gupta and Neggers (2018) find that Indian voters are more likely to vote against candidates accused of serious crimes when they are informed that many other voters received the same SMS.

Taken together, the preceding information diffusion and voter coordination mechanisms generally predict that any impact of information dissemination is likely to be induced or amplified by greater campaign saturation. We test this hypothesis in the context of a large-scale accountability campaign undertaken before Mexico's 2018 elections.

[9] This is not necessarily the case, as revelations of electoral fraud discouraging voter turnout illustrate (Chong et al. 2015; McCann and Domínguez 1998; Simpser 2012).

3 Mayoral malfeasance and accountability in Mexico

Mexico's c. 2,500 municipal governments are led by mayors typically elected to three-year terms, which became renewable for the first time in most states in 2018. These governments are responsible for delivering basic public services and managing local infrastructure, which can—if used effectively—play an important role in poverty alleviation and local development (Rodríguez-Castelán, Cadena and Moreno 2018). However, municipal accountability remains limited, and corruption is still common. While the constitutional reform to permit re-election was partly designed to mitigate this, it is unlikely to be sufficient in a context where voters are poorly informed about government performance (Chong et al. 2015).

3.1 Independent audits of municipal spending

A key source of funding for mayors is the Municipal Fund for Social Infrastructure (FISM). These direct federal transfers represent around a quarter of the average municipality's budget and are mandated exclusively for infrastructure projects that benefit (i) localities deemed to be marginalized by the National Population Council (CONAPO), (ii) citizens in extreme poverty, or (iii) priority zones.[10] Eligible projects include investments in the water supply, drainage, electrification, health infrastructure, education infrastructure, housing, and roads.

The use of FISM transfers is audited in around 200 municipalities each year by Mexico's independent Federal Auditor's Office (ASF). ASF's audits are announced after spending has occurred, and address the spending, accounting, and management of FISM funds from the previous fiscal year. Municipalities are selected by the ASF on the basis of the importance of FISM transfers to the municipal budget, historical performance, factors that raise the likelihood of irregularities in the management of funds, and whether the municipality has recently been audited (including concurrent federal audits of other programs) (see Auditoría Superior de la Federación 2014). The large municipalities comprising most of the country's population have now received multiple audits since the systematic audits began in 2004.

This article focuses on irregularities in the expenditure of FISM resources—the primary outcome of the ASF's audits. Irregularities typically entail funds that are spent on projects not benefiting the poor (based on the distribution criteria above) or spent on unauthorized projects that do not constitute social infrastructure projects (e.g. personal expenses and election campaigns). The audit reports indicate that such irregularities typically arise from failing to demonstrate that the project benefited its intended recipients, the transfer of funds to non-FISM bank accounts or contractors, or failures to produce documentation proving that expenses related to claimed projects. These ac-

[10] In 2010, the CONAPO defined the 79% of localities scoring high or very high on its marginalization index as eligible.

tions often reflect corruption in the form of kickbacks, preferential contracting, and embezzlement (Larreguy, Marshall and Snyder 2019). Between 2009 and 2018, the ASF determined that 17% of funds spent were subject to irregularities.[11]

The potential for voters to punish high levels of mayoral malfeasance and reward clean incumbents is limited by an electorate largely uninformed about the ASF's reports. Most voters are unaware of the resources available to mayors and even their responsibility to provide basic public services in the first place (Chong et al. 2015). ASF's reports are publicized in some media outlets and have been shown to influence voting behavior in large urban environments where there is a large audience for such information (Larreguy, Marshall and Snyder 2019). However, because coverage is not widespread and voter engagement with news programming varies, the dissemination of such information can significantly alter voters' beliefs and voting behavior. Consistent with this, Arias, Larreguy, Marshall and Querubín (2019) find that distributing the results of ASF reports via non-partisan leaflets caused voters to update their high expectations of incumbent party malfeasance, and vote accordingly. Chong et al. (2015) have also found that publicizing severe levels of unauthorized FISM spending can breed voter disengagement, with a particularly detrimental effect on supporters of challengers. This article complements these studies by investigating the extent to which the saturation of information provision facilitates electoral accountability.

3.2 Electoral context

Until recently, electoral competition in most Mexican municipalities was between two of the country's main three parties. In most parts of the country, the populist PRI competed 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). In 2014, ex-PRD leader Andrés Manuel López Obrador formed MORENA, a new left-wing and anti-corruption party which stood for the first time in 2015 and displaced the PRD in many areas. Although MORENA's local presence was initially limited, it swept the 2018 elections as López Obrador won the presidency by a landslide. MORENA's national success carried over to local elections as well, with MORENA claiming multiple governor offices and hundreds of mayoral offices across the country.

Municipal election campaigns are generally oriented around political parties, rather than specific candidates, for several reasons. First, given that mayoral consecutive re-election was only permitted for the first time in 2018, voters are much better informed about parties than individual politicians (e.g. Arias, Larreguy, Marshall and Querubín 2019; Chong et al. 2015; Larreguy,

[11] Given that other programs and non-federal transfers are not subject to such audits, mayoral malfeasance could be greater on other dimensions. Nevertheless, we expect malfeasance across areas to be correlated, and thus that information about FISM irregularities—which represent a substantial share of a municipality's budget—will be indicative of an administration's broader malfeasance.

Marshall and Snyder 2019). Second, voters may recognize that Mexico's main parties use distinct candidate selection mechanisms that select candidates with similar characteristics over time (Langston 2003). Consequently, voters have held parties responsible for the actions of individual politicians (e.g. Chong et al. 2015; de Figueiredo, Hidalgo and Kasahara 2014; Larreguy, Marshall and Snyder 2019; Marshall 2019*b*). Despite the fact that only 22% of mayors sought re-election in 2018, there are thus good reasons to believe that voters will make inferences about the party of the mayor whose audited expenditures are publicized and vote accordingly.

3.3 Political information and social media environment

While broadcast media outlets have traditionally been the primary source of political information in Mexico (e.g. Marshall 2019*b*), mobile technology and social media have created new opportunities for information dissemination. According to Hootsuite and We are Social (2018*a*,*b*), 65% of Mexicans accessed the internet in 2018, with the average respondent spending more than eight hours a day online—the 7th highest rate in the world. Moreover, 72% of adults own a smart phone—the primary means of accessing the internet in Mexico for most adults—and 64% of adults used social media in 2018; social media users reported spending an average of more than three hours a day using it. With almost all social media users using Facebook at least once a month, Mexico ranks 5th in the world in terms of active Facebook users. Due to it being free to use once internet access has been established, WhatsApp has become the messaging service of choice in major developing countries like Brazil, India, and Mexico. In Mexico, it is the most used and most downloaded mobile phone app.

Growing access to digital information has emerged alongside substantial amounts of credible and fake political information disseminated through social and traditional media. In addition to economic and security issues, corruption was a key issue throughout the presidential, legislative, and municipal election campaigns of 2018. Fake news was a particular concern during the 2018 election campaign, where political parties were accused of disseminating fake news aided by bots to influence voter behavior.[12] The attacks were largely directed against the eventual winner López Obrador and were disseminated via Facebook and WhatsApp.[13] However, many other candidates

[12]For example, due to the way that Facebook's algorithm works, "likes" to a Facebook page or ad increase their visibility. Facebook pages criticizing López Obrador featured posts with thousands of "likes," but no other reactions or comments, suggesting the work of bots. See here for more details.

[13]For example, fake news articles that claimed that López Obrador's wife posted on Twitter that she was disgusted by indigenous people—in a country where official figures indicate that 21.5% of the population is indigenous—were widely shared on Facebook (see here for more details). A fake poll in a major national newspaper suggesting that the PAN candidate was within 5 percentage points of López Obrador's—the latter eventually beat the former by more than 30 percentage points—was also widely circulated by PAN candidates over social media (see here for more details). Fake pictures of rallies with very few attendees were circulated to claim that López Obrador's support was deflating (see here for more details).

across all races were also affected by similar types of attacks.[14] Due to their popularity among Mexican voters, Facebook and WhatsApp were the prime channels for spreading real and fake news in the form of videos, images, and memes. Several Facebook pages that were identified as the most prolific fake news distributors had between one and two million followers around the election.

4 Research Design

Our goal is to identify whether campaign saturation moderates the extent to which incumbent performance information affects voting behavior. We partnered with Borde Político—an NGO primarily based in Mexico City, which uses digital technologies to promote government transparency across the country—to evaluate the impact of their online accountability campaign ahead of the July 1, 2018 elections. Borde Político's information campaign, which focused on the municipal elections, provided voters with information about the FISM program and the share of audited resources that the ASF found to be subject to irregularities. Widespread access to social media enabled the study to randomly vary the share of the municipal population with access to this information via Facebook. This section describes the treatment conditions, sample, experimental design, measurement of outcomes, and estimation, concluding with a discussion of ethical considerations.

4.1 Treatment conditions

Like extant studies which have provided voters with incumbent performance information, the information campaign reported the results of the ASF's audit in a given municipality.[15] Voters that received the information were first informed that the FISM program transfers federal funds to municipalities for social infrastructure projects benefiting the poor. They were then informed of how much money their municipal government received, and the percentage of the audited funds that were subject to irregularities in terms of violating FISM spending regulations.

As part of Borde Político's broader transparency campaigns, this information was disseminated to Facebook users in treated municipalities by a Facebook video ad. Figure 1 shows the slides that make up the 26-second video. The municipality's share of FISM expenditures subject to irregularities is reported in the 17th second. To bolster credibility, the ads were accompanied by a legend indicating that Borde Político is a non-partisan NGO that provides ASF information to inform voters and included links to the Borde Político and ASF websites. Users could also click to access the municipality-specific Facebook page that "boosted" the ad.[16] These pages included

[14] See, for example, here, here, and here.

[15] While the information content provided is similar to prior interventions in Mexico (e.g. Arias, Larreguy, Marshall and Querubín 2019; Chong et al. 2015), this study differs by focusing on the impact of digital dissemination and municipal campaign saturation. By leveraging similar information content, prior studies help benchmark these effects.

[16] A separate page was created for each basic and common knowledge version of the video (see below).

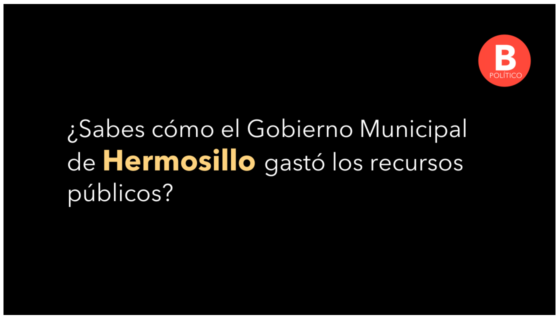

(a) Slide 1 (4 seconds)

(b) Slide 2 (4 seconds)

(c) Slide 3 (4 seconds)

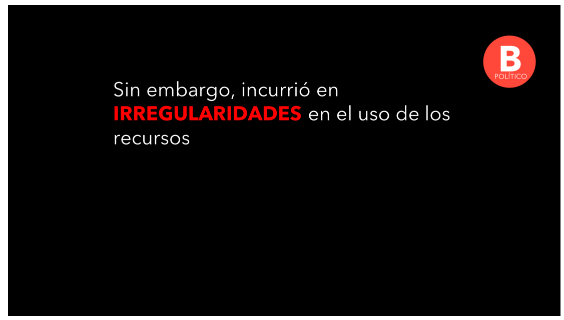

(d) Slide 4 (4 seconds)

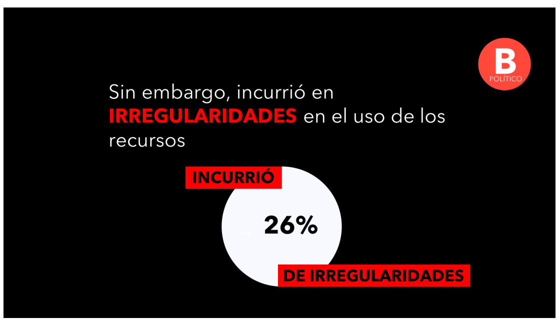

(e) Slide 5 (4 seconds)

(f) Slide 6 (2 seconds)

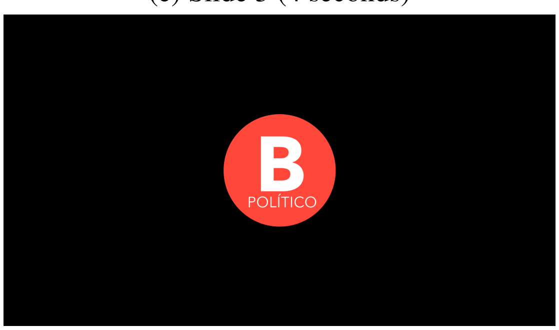

(g) Slide 7 (4 seconds)

Figure 1: Example of the slides included in the ad video (from Hermosillo, Sonora)

Note: In English: slide 1 says "Do you know how the municipal government of Hermosillo spent public monies?;" slide 2 says "In 2016, the municipal government of Hermosillo received funds from the Fund for Social Infrastructure;" slide 3 adds "Received $65 million for infrastructure;" slide 4 says "However, it incurred in irregularities in the spending of the funds;" slide 5 adds "Incurred in 26% of irregularities;" and slide 6 says "Unauthorized spending and targeting people other than the intended beneficiaries are irregularities that cause damage to government finances."

a cover photo highlighting the amount of money received and the fraction of spending subject to irregularities and an infographic reporting this information in greater detail (see the examples in Appendix Figures A1a or A1b). Each municipality ad campaign ran for a week, concluding on June 27, 2018—the last day of official campaigning. Incumbents thus had no time to meaningfully alter their performance in office before the election in response to the ads, and parties had little time to respond during the campaign.

This study's primary innovation is to randomly vary the saturation of Borde Político's Facebook ad campaign across municipalities. Accordingly, Facebook ads were geographically targeted with the capacity to reach 20% of Facebook users of voting age (18+) in low-saturation municipalities, while Facebook ads sought to reach 80% of Facebook users of voting age in high-saturation municipalities.[17] The 20% and 80% saturation levels were chosen to capture a meaningful difference in saturation that could plausibly alter levels of information sharing or coordination, while also maximizing statistical power to estimate direct, indirect, and differential saturation effects (see below).[18] The average municipal ad campaign cost around US$250, representing a tiny fraction of a typical municipal election campaign's budget.

Our lack of control over Facebook's proprietary ad generation algorithm meant that Borde Político's ad campaigns could not be designed to ensure that the ad would reach all voting age adult Facebook users within a targeted location a certain number of times. Rather, for a given investment, Facebook specifies the maximum possible reach of a campaign within a geographic area for a particular demographic. Consequently, the ad campaigns were funded to be able to reach the designated 20% or 80% of voting age Facebook users in low and high saturation municipalities as many times as possible and at equal rates across municipalities. While Facebook does not publicly disclose its constantly-evolving technology used to identify user locations, our conversations with Facebook staff indicated that whether a given Facebook user is targeted by a geographically-constrained ad depends primarily on the location where Facebook believes that users spend most time, based on user-specific GPS data received by Facebook. For most users, this is their home.[19] When a user's GPS data is unavailable, targeting is based on data including the user's IP address, search traffic, and the locations of a user's friends. Since 88% of users accessed the ads via a mobile device, ads are generally likely to be targeted with a high degree of accuracy. We provide evidence consistent with this in our discussion of mechanisms below. We further note that any compari-

[17]In all treated (and some control) municipalities, individual WhatsApp messages were sent to a mean of 50 surveyed voters as part of a concurrent panel survey designed to understand the mechanisms underlying the Facebook campaign. However, since this number represents a negligible fraction of the municipal population, we disregard them when defining municipal treatments. This approach is supported by the lack of a significant difference in electoral outcomes across control municipalities that did and did not contain respondents that received WhatsApp messages.

[18]Following Baird et al. (2018), we minimized the equally-weighted sum of the standard errors for treatment and spillover effects, where municipalities were equally split between control, low saturation, and high saturation.

[19]The geographic areas covered by Borde Político's Facebook ads are large enough that they will often encompass both the home and workplace of Facebook users.

son between treated and control (or spillover) groups would underestimate treatment effects in the presence of mistargeting.

The ad's content was also subtly randomized to explicitly vary common knowledge about the ad campaign's reach. As in George, Gupta and Neggers (2018), Facebook users in some locations were informed that the ad campaign could reach 20% or 80% of voters in their municipality. This entailed adding the slide shown in Appendix Figure A5 at the end of the video.[20] This variant of the basic treatment intended to facilitate voter coordination by increasing the probability of discussion or explicit verbal agreement (e.g. on whether to vote or who to vote for) or altering higher-order beliefs (e.g. about how fellow citizens were likely to vote). While we ultimately observe no discernible differences in viewership of—or reactions to—the ads with and without explicit common knowledge communication (see Appendix Table A10), this variant of the treatment also did not differentially affect voting behavior (see Appendix Table A11). For this reason, we henceforth pool the Facebook ads with and without common knowledge in all analyses.

4.2 Sample of municipalities

Across 2017 and 2018, the ASF released audit reports pertaining to FISM expenditures in 561 municipalities.[21] Of these, 128 municipalities are from the 17 states that held municipal elections in 2018 where the mayor in office before the election was also the mayor that presided over the audited expenditures.[22] These 128 municipalities constituted the sample for Borde Político's Facebook information campaign, and are shaded in Figure 2. The sample collectively contains around 30 million people, roughly a quarter of Mexico's population, and is thus broadly nationally representative.

Figure 3 shows that the majority of ASF audits in these municipalities reported irregularities between 0% and 10%. Exactly zero irregular spending was found in 61 of our 128 municipalities. The mean share of irregular spending across municipalities was 9.2%, with a positive skew driven by several egregious cases. Given the low expectations of politicians recently registered in Mexico (Arias, Larreguy, Marshall and Querubín 2019), the ads likely reported better performance than expected for most voters.

[20]To avoid deception, this information always reflected the true share of people that the campaigns sought to reach.

[21]The two delegaciones in Mexico City were excluded because such delegaciones operated differently from municipalities during the relevant time period.

[22]These states are: Baja California Sur, Campeche, Chiapas, Colima, Estado de México, Guanajuato, Guerrero, Jalisco, Michoacán, Morelos, Nuevo León, Puebla, Querétaro, San Luis Potosí, Sonora, Tabasco, and Yucatán. Municipalities from states like Coahuila, where mayors were elected in 2017 and thus were not responsible for the spending audited by the ASF, were not included in our sample. An additional 7 states held municipal elections, but none of these municipalities were eligible for Borde Político's campaign due to their shorter electoral cycles. The other 7 states did not hold municipal elections.

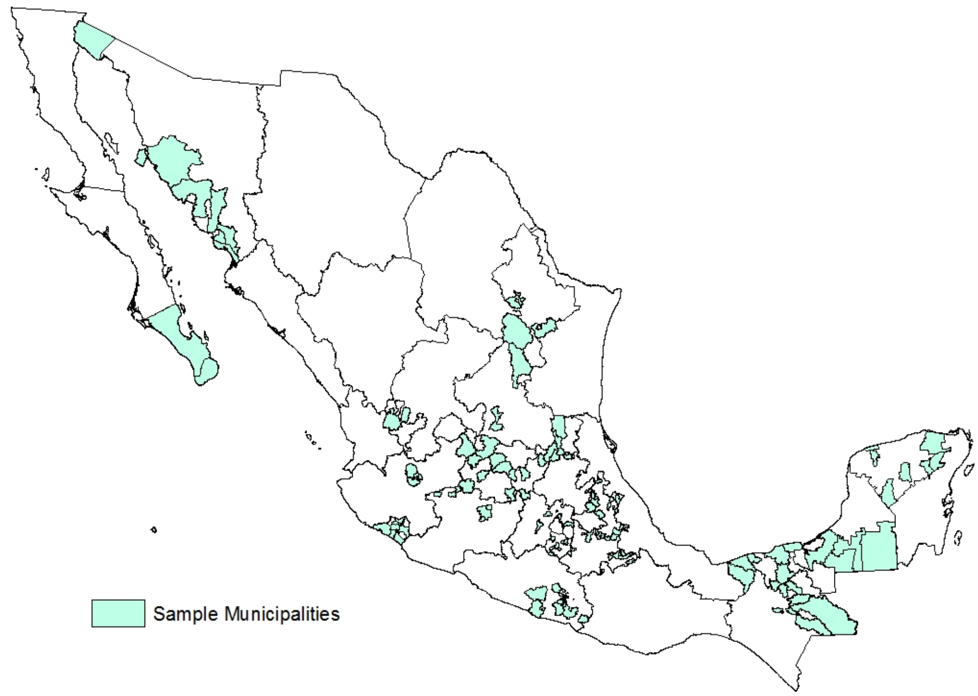

Figure 2: The 128 municipalities included in our sample

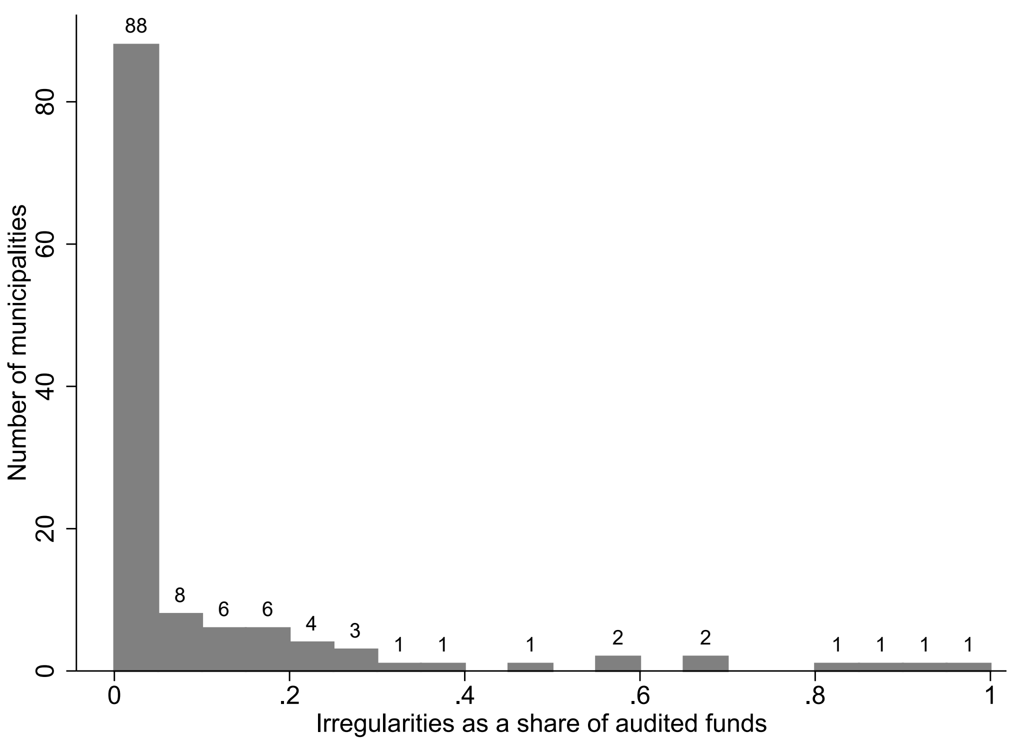

Figure 3: Distribution of irregularities across municipalities in our sample (each band represents intervals of 0.05 shares)

4.3 Experimental design

To identify the electoral effects of Borde Político's information campaign, we designed a two-level randomization strategy. This first assigned campaign saturation at the municipality level and then assigned Facebook ads to segments (defined below) within municipalities selected to receive a non-zero saturation ad campaign. This approach enables us to estimate the effect of information provision on voting behavior, the extent to which behavior spills over to non-treated segments within partially treated municipalities, and whether these effects vary by municipal saturation level.

The saturation of Borde Político's Facebook ad campaign was first randomized at the municipal level as follows. Each municipality was assigned to one of 42 blocks containing 3 municipalities governed by the same incumbent party on the basis of the Mahalanobis distance over 28 covariates, with the exception of 2 rump municipalities that formed an additional block governed by different parties. For simplicity, we henceforth exclude the small block from our analysis.[23] Within each block, one municipality was assigned to each of the following conditions:

1. *Control*: no Facebook ads;

2. *Low saturation information campaign*: Facebook ads were targeted at 20% of adult Facebook users within the municipality; and

3. *High saturation information campaign*: Facebook ads were targeted at 80% of adult Facebook users within a municipality.

Since around 70% of Mexican adults regularly use Facebook, the targeted share of *registered voters* was effectively 14% in low saturation municipalities and 56% in high saturation municipalities. This blocking procedure ensured that each municipality had an equal probability of being treated, without differentially targeting incumbents from any particular political party. Appendix Table A3 shows that campaign saturation is well balanced across predetermined municipal-level covariates.

The targeting of Facebook ads to geographic areas within treated municipalities was then further randomized. To be able to reach up to 20% and 80% of adult Facebook users in low and high saturation cases, we divided each municipality into (multiples of) 5 equally-populated "segments" comprised of electoral precincts—Mexico's smallest electoral geographic unit. In small municipalities, we created five segments. In larger municipalities, where it was feasible to target more segments using Facebook's targeting system, we created multiples of five segments. The resulting 783 segments were defined by contiguous electoral precincts that form compact polygons with

[23] This deviation from our pre-analysis plan was deployed primarily to simplify estimation by maintaining a constant probability of treatment assignment across municipalities (see Appendix section A.6). Appendix Table A2 reports similar results when the two small municipalities in this residual block are included.

similar populations of individuals aged 18 or above (according to the 2010 Census).[24] Complete randomization was then used to assign one fifth of segments within low saturation municipalities, and fourth fifths of segments within high saturation municipalities, to be the targets of Borde Político's Facebook ads.[25] As Appendix Table A4 shows, the segment-level treatment conditions are well balanced across predetermined covariates.

A parallel panel survey was conducted to help illuminate the mechanisms driving any effects of the Facebook ad campaign on voting behavior. This survey yielded complete baseline and endline responses from around 1,800 registered voters that use WhatsApp within the 128 municipalities.[26] The 20-minute baseline survey was conducted over 2-3 weeks in early June 2018, while the 20-minute endline survey was conducted over the month after the election. As an additional randomized treatment, the Facebook ad video and the information accompanying the ad on Facebook, were sent via a WhatsApp message in the week preceding the election to 80% of baseline respondents within both treated municipalities and 23 of the pure control municipalities (see Appendix section A.4). An example of the WhatsApp message is shown in Appendix Figure A2. Treated survey respondents also received a similar followup message that included the infographic shown in Appendix Figure A1.

Unfortunately, we encountered differential attrition of survey respondents across municipalities assigned to different Facebook ad saturation levels (see Appendix Table A6). We are thus unable to leverage variation in saturation to study the effect of municipal Facebook ad saturation on survey outcomes. However, this problem does not arise across treated and control individuals within the municipalities where Facebook ads were delivered (see Appendix Table A5), so our individual-

[24]These segments were generated by the freely-downloadable redistricting program Auto-Redistrict (autoredistrict.org). This software allows users to redistrict blocks of precincts into "districts" to maximize the contiguity, compactness, and equal population of districts. Precinct allocation was then manually adjusted at the margins to smooth edges in order to facilitate ease of targeting with Facebook ads (given the targeting constraint of needing to pick points with 1km radii). The latter adjustment effectively slightly relaxed the population equality constraint. The total number of segments is not a multiple of 5 because one small municipality contains only 3 electoral precincts.

[25]Whether a treated segment would additionally receive common knowledge ads informing voters of how many other voters within their municipality also had access to the ad was also randomized. Within the large majority of low-saturation municipalities with only five segments, only one segment was treated; complete randomization determined whether that segment received the common knowledge treatment. In the few low-saturation municipalities with a multiple of five segments, receiving an equal number of non-common knowledge and common knowledge treated segments was prioritized. Within high-saturation municipalities, half of the treated segments (i.e. 40% of the municipality's segments) received the common knowledge information and the other half did not. As noted above, we focus on the results from analyses that pool ads with and without common knowledge.

[26]The survey was conducted by GeoPoll. They generated a sample based on calling and messaging randomly-generated cell numbers (based on areas codes local to our municipalities), Telmex landline numbers, and the completion of an online Qualtrics survey recruited via a separate Facebook ad campaign solely seeking to recruit survey respondents (which made no reference to Borde Político's information campaign). We aimed to recruit 31 respondents per municipality for a baseline survey, and 20 of these for endline survey. Recruitment rates differed across municipalities, but yielded around 14 endline respondents in the average municipality. To incentivize continued participation, respondents that completed the baseline survey were entered into a lottery to win one of 10 prizes with a value equivalent to a new smart phone; an additional independent lottery was used for the endline as well.

level analyses of survey-level outcomes are restricted to examining within-municipality random variation in being sent the WhatsApp message.

4.4 Measurement of outcomes

Our main outcomes are taken from the precinct-level electoral returns collated by Mexico's state electoral institutes. We focus primarily on the municipal incumbent party's vote share, using the predetermined number of registered voters as the denominator. Since we find a limited impact on turnout, vote share using total votes cast in the denominator reports similar results (see Appendix Tables A13 and A14). Our focus on municipal incumbent parties as the unit of analysis accords with prior studies in Mexico's party-centric electoral context (e.g. Arias, Larreguy, Marshall and Querubín 2019; Chong et al. 2015; Larreguy, Marshall and Snyder 2019). Electoral returns were available for all but two municipalities—Oxchuc, Chiapas and Ayutla de los Libres, Guerrero—which in 2018 adopted a customary system for indigenous peoples to select their mayors that did not involve direct election.

To ascertain engagement with the treatments and possible intermediary mechanisms, we use two additional data sources. First, Facebook analytics data associated with each municipality's ad campaigns allows us to measure, at the municipal level, how many people saw the ads (and for how long) as well as gauge the extent of user interactions. Second, our endline survey also elicited engagement with the WhatsApp messages, respondents' perceptions of the source and credibility of the messages, and various types of beliefs.

4.5 Estimation

Our pre-registered specifications leverage the multiple layers of randomization to identify the effects of Borde Político's information campaign on administrative electoral results.[27] First, we follow Baird et al. (2018) in leveraging the segment-level assignments of Borde Político's Facebook ad campaign to identify the direct and indirect average treatment effects on precinct-level outcomes using the following specification:

$$Y_{psm} = \alpha Y_{psm}^{lag} + \beta Facebook\ ads_{sm} + \gamma Spillover_{sm} + \mu_b + \varepsilon_{psm}, \tag{1}$$

where Y_{psm} is an outcome in precinct p within segment s of municipality m, Y_{psm}^{lag} is a lag of the outcome, $Facebook\ ads_{sm}$ is an indicator for a treated segment s receiving Facebook ads, $Spillover_{sm}$ is an indicator for an untreated segment s located within a treated municipality, and μ_b are fixed

[27]Minor deviations from the pre-analysis plan are explained in Appendix section A.6. Additional specifications that are not presented in the main article due to space constraints are also reported in Appendix section A.6.

effects for the blocks of three similar municipalities within which treatment was assigned. The reference category is the set of precincts from control municipalities assigned to receive no Facebook ads.[28] Observations are weighted by the design's inverse probabilities of treatment assignment and, to weight segments equally, each precinct's share of the segment's 2010 adult population aged 18 or above. Standard errors are clustered by municipality for electoral outcomes throughout because saturation was randomized at the municipal level.

The coefficients β and γ in equation (1) respectively estimate the direct intent to treat effect of the Facebook ad campaign and the indirect spillover effect of being located in a non-treated segment within a treated municipality, both relative to the pure control condition. Rejecting $\gamma = 0$ would suggest that social interactions in response to the Facebook ads influence voter behavior. Since our theoretical expectations are conditional on the content of the information provided by the Facebook ads, we pre-registered no directional hypotheses for the *average* effects within the entire sample.

Second, we identify the differential effects of segment-level Facebook ad treatments across municipal saturation levels using regressions of the following form:

$$\begin{aligned} Y_{psm} &= \alpha Y_{psm}^{lag} + \beta_1 Facebook\ ads\ in\ Low\ Saturation_{sm} + \beta_2 Facebook\ ads\ in\ High\ Saturation_{sm} \\ &\quad + \gamma_1 Spillover\ in\ Low\ Saturation_{sm} + \gamma_2 Spillover\ in\ High\ Saturation_{sm} + \mu_b + \varepsilon_{psm}, \quad (2) \end{aligned}$$

where differences in the effects of segment-level treatments across municipal saturation levels enable us to identify differential treatment effects attributable to high rather than low saturation campaigns. The differences $\beta_2 - \beta_1$ and $\gamma_2 - \gamma_1$ in equation (2) respectively identify the differential effect of the direct and indirect Facebook ad treatments that can be attributed to high saturation. As argued above, we expect $\beta_2 - \beta_1 > (<)0$ and $\gamma_2 - \gamma_1 > (<)0$ when the average effect is positive (negative). Because this expectation does not depend on the direction that the information campaign influences voting behavior, we pre-registered one-sided tests of these hypotheses.

As with all informational interventions, it was not obvious *a priori* how information content would affect support for incumbent parties. Increases or decreases in support are likely to depend on whether reported irregularities exceed voters' prior expectations and whether/how different reported irregularities coordinate voter behavior. Since we could not systematically measure prior beliefs across all segments, we follow Ferraz and Finan (2008) and others in generating bins where reported irregularities are likely to exceed and fall below expectations. In particular, we divide the distribution of reported irregularities into quartiles (see also Larreguy, Marshall and Snyder 2019). The bottom two quartiles are pooled because more than 25% of municipalities registered zero ir-

[28] This includes zero-saturation municipalities where only 0.02% of registered voters received WhatsApp treatments on average.

regularities. Irregularities in the third quartile range from 0.08% to 7.4%, with a mean of 2.2%. Irregularities in the top quartile range from 7.4% to 100%, with a mean of 31.7%. The interactive specifications then extend equations (1) and (2) by including interactions between treatment conditions and a municipality's irregularities quartile. The interactive version of equation (1) entails estimating:

$$\begin{aligned} Y_{psm} \quad = \quad & \alpha Y_{psm}^{lag} + \beta_1 Facebook\ ads_{sm} + \beta_2 \big(Facebook\ ads_{sm} \times Q3\big) \\ & + \beta_3 \big(Facebook\ ads_{sm} \times Q4\big) + \gamma_1 Spillover_{sm} + \gamma_2 \big(Spillover_{sm} \times Q3\big) \\ & + \gamma_3 \big(Spillover_{sm} \times Q4\big) + \delta_1 Q3 + \delta_2 Q4 + \mu_b + \varepsilon_{psm}. \end{aligned} \tag{3}$$

We use two-tailed tests to reflect the theoretical uncertainty over the direction of the effects, although we expect to observe increases in incumbent support in Q1/Q2 and decreases incumbent support in Q4.

4.6 Ethical considerations

Our collaboration with the non-partisan NGO Borde Político meets prevailing ethical standards. First, we took great care to ensure that our evaluation of their campaign complied with all institutional and legal requirements for academic research and NGO activity. The study was approved by each of the three Institutional Review Boards at the universities of the authors, which include a Mexican institution. Second, the intervention also complied with Mexican electoral law. As part of a similar Borde Político information campaign in 2015, which was evaluated by some of this article's authors Arias, Larreguy, Marshall and Querubín 2019),[29] Mexican electoral authorities indicated that electoral law permits NGOs to exercise the freedom of expression they enjoy as collectives of citizens in order to disseminate non-partisan information about municipal government performance. We corroborated this legal interpretation with a local electoral lawyer.

Beyond satisfying institutional and legal requirements, we regard our collaboration with Borde Político to evaluate their information campaign as both ethical and academically valuable for several further reasons. First, the intervention evaluates the impact of information provision by a non-partisan NGO that frequently disseminates politician performance information online, including through its Facebook and Twitter accounts, with the goal of enhancing political accountability. Moreover, all treatments were disseminated on behalf of Borde Político, who suggested delivering the information via Facebook to help understand how the effectiveness of their non-partisan campaigns could be maximized. The collaboration thus facilitated the first rigorous evaluation of the effect of campaign saturation within the context of digital dissemination technologies that Borde

[29] Borde Político disseminated similar information through leaflets.

Político—and many other NGOs—frequently use. Furthermore, the results could also be relevant to policymakers assessing the potential benefits and risks surrounding social media during election campaigns.

Second, in line with electoral law, all possible means were used to ensure that the campaign remained non-partisan. All numerical information, as well as the wording around it, was extracted from the independent and non-partisan ASF's online audit reports. In addition, the sample was chosen in a non-partisan manner: subject to a municipality holding elections in 2018 and the report pertaining to the current incumbent mayor's term in office, all municipalities for which a report was available entered our sample. To ensure that no party was differentially targeted, treatment assignments were blocked on party. We further minimized perceptions of bias by avoiding the use of color schemes associated with any particular party.

Third, learning about the unstudied effects of high saturation campaigns inevitably requires concentrating information dissemination in ways that increase the possibility of affecting voting behavior. However, rather than providing voters with fake or distracting information that could scramble their capacity to vote for their preferred candidate, the transparent and independent information that Borde Político provided was particularly relevant during an election when corruption was a salient issue for voters. Prior studies disseminating similar information show that Mexican voters care about the use of FISM funds, which account for close to a quarter of the municipal budget (Arias, Balán, Larreguy, Marshall and Querubín 2019; Arias, Larreguy, Marshall and Querubín 2019; Chong et al. 2015; Larreguy, Marshall and Snyder 2019). Nevertheless, the ads—which did not mention the upcoming elections—did not explicitly ask voters to respond in any way to the information. We thus expected that the information provided would help voters to reach better-informed voting decisions if they regarded the information as relevant.

Finally, we believe that it is important to inform Borde Político's commendable goal of improving municipal electoral accountability in a country where municipal governance is widely perceived to be tainted by corruption. Such goals often involve NGOs explicitly seeking to influence electoral outcomes. However, the chances of doing so in this particular context were slim because Facebook ads are often internalized by a small fraction of those targeted and because MORENA was widely expected to—and ultimately did—win by a landslide across much of the country.

5 Consumption and comprehension of Facebook ads

Before turning to electoral outcomes, we first examine the information campaign's reach, comprehension, and credibility. We show that Borde Político's Facebook ads achieved significant saturation in proportion with a municipality's intended saturation level, and that such information was internalized and generally regarded as credible by voters.

5.1 The reach of the intervention

Over the course of Borde Político's accountability campaign, the Facebook ads appeared 7.3 million times on the screens of 2.7 million different Facebook users across treated municipalities. Table 1 presents Facebook's ad campaign-level analytics data by municipality, demonstrating that the campaign reached a considerable share of targeted adults in at least a limited way. Appendix Figure A3 plots trends in Facebook ad engagement by day, indicating that the campaign's reach increased over the course of the campaign's week.

Table 1 first shows that the ad campaign ultimately reached a substantial fraction of targeted adults, broadly in proportion with the campaign's intended saturation level. Our estimates come from the following regression:

$$Y_{psm} = \beta_1 Low\ saturation_m + \beta_2 High\ saturation_m + \mu_b + \varepsilon_{psm}, \tag{4}$$

where $Low\ saturation_m$ and $High\ saturation_m$ are, respectively, indicators for the municipal-level low and high saturation treatment conditions. Column (1) reports that the Facebook ad appeared (on Facebook's "News Feed") as paid content 0.32 times for every adult member of the population in low saturation municipalities, and 0.95 times per adult in high saturation municipalities. Column (2) respectively reports a further 0.03 and 0.04 impressions per adult coming from organic views, which arose when friends on Facebook encountered the ad because Facebook ad viewers shared, commented, or reacted to (e.g. liked) an ad. To adjust for the number of voters targeted, these numbers can be divided by the saturation level. Turning to unique Facebook users in columns (3), the campaign reached more than one third of its intended population in the average municipality, i.e. around half of the population of Facebook users. Column (4) again indicates that comparatively few additional views were generated organically. Although we cannot establish the intersection of respondents reached through paid-for ads and organic views or the location of organic views, these numbers imply that the mean Facebook user encountered the ad around three times. Appendix Figure A4 shows that, while there is variation in Facebook user engagement within assigned saturation levels, most 80% saturation municipalities experienced notably greater engagement than even the most-engaged 20% saturation municipalities.

A non-trivial number of Facebook users also engaged with the ad. Column (5) shows that 2-3% of targeted voting age adults clicked on the ad, e.g. by sharing, liking, or commenting on the ad or clicked through to the Facebook page. Furthermore, columns (7) and (9) respectively report that there were 0.04 views per adult of at least 3 seconds and 0.02 views per adult of at least 10 seconds in low saturation municipalities, while there were 0.11 and 0.06 such views per adult in high saturation municipalities. Column 11 shows that the corresponding numbers are 0.015 and 0.04 per adult, respectively, for the share of adults that watched the entire ad. The share of

Table 1: Effect of municipal treatments on municipal Facebook engagement

	Municipal counts per capita (normalized by 2015 adult population)										
	Paid-for impressions	Organic impressions	Paid-for unique viewers	Organic unique viewers	Unique user page engagements	Total views (of 3 seconds)	Unique views (of 3 seconds)	Total views (of 10 seconds)	Unique views (of 10 seconds)	Total views (of entire video)	Unique views (of entire video)
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)
Panel A: average treatment effects											
High saturation	0.949***	0.042***	0.279***	0.024***	0.013***	0.181***	0.113***	0.075***	0.060***	0.043***	0.039***
	(0.103)	(0.007)	(0.026)	(0.004)	(0.002)	(0.018)	(0.011)	(0.007)	(0.006)	(0.004)	(0.004)
Low saturation	0.323***	0.030***	0.084***	0.016***	0.007***	0.064***	0.037***	0.028***	0.021***	0.017***	0.015***
	(0.080)	(0.009)	(0.019)	(0.005)	(0.002)	(0.014)	(0.008)	(0.006)	(0.005)	(0.004)	(0.003)
Observations	124	124	124	124	124	124	124	124	124	124	124
R^2	0.64	0.54	0.71	0.54	0.54	0.68	0.69	0.67	0.68	0.65	0.66
Control outcome mean	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
Control outcome std. dev.	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
Test: Low saturation = High saturation (p value, 2-sided)	0.000	0.157	0.000	0.114	0.002	0.000	0.000	0.000	0.000	0.000	0.000
Test: 4 * Low saturation = High saturation (p value, 2-sided)	0.258	0.019	0.420	0.022	0.050	0.159	0.238	0.124	0.178	0.067	0.084
Panel B: heterogeneity by irregularities quartile											
High saturation	0.703***	0.037**	0.216***	0.019**	0.013***	0.132***	0.085***	0.060***	0.048***	0.036***	0.032***
	(0.119)	(0.013)	(0.032)	(0.007)	(0.003)	(0.022)	(0.014)	(0.011)	(0.008)	(0.007)	(0.006)
High saturation × Q3	0.180	0.024	0.036	0.016	0.003	0.046	0.025	0.017	0.013	0.011	0.011
	(0.279)	(0.025)	(0.068)	(0.014)	(0.006)	(0.050)	(0.030)	(0.023)	(0.018)	(0.015)	(0.013)
High saturation × Q4	0.937**	0.012	0.240**	0.011	-0.000	0.166**	0.097**	0.050*	0.041*	0.021	0.020*
	(0.312)	(0.021)	(0.074)	(0.012)	(0.005)	(0.055)	(0.032)	(0.020)	(0.016)	(0.011)	(0.010)
Low saturation	0.383**	0.026*	0.096**	0.014*	0.007*	0.065**	0.039**	0.031**	0.024**	0.019**	0.017**
	(0.133)	(0.012)	(0.030)	(0.006)	(0.003)	(0.021)	(0.013)	(0.010)	(0.008)	(0.007)	(0.006)
Low saturation × Q3	-0.131	0.005	-0.032	0.002	-0.003	-0.008	-0.007	-0.011	-0.009	-0.006	-0.005
	(0.222)	(0.022)	(0.058)	(0.012)	(0.005)	(0.041)	(0.024)	(0.019)	(0.014)	(0.011)	(0.010)
Low saturation × Q4	-0.103	0.016	-0.017	0.010	0.000	-0.001	-0.000	-0.003	-0.003	0.000	0.000
	(0.225)	(0.021)	(0.052)	(0.012)	(0.005)	(0.039)	(0.023)	(0.016)	(0.012)	(0.011)	(0.009)
Observations	124	124	124	124	124	124	124	124	124	124	124
R^2	0.73	0.59	0.77	0.59	0.56	0.74	0.75	0.71	0.72	0.68	0.68
Control outcome mean	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
Control outcome std. dev.	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00

Notes: Each specification is estimated using OLS, and includes randomization block fixed effects. The two municipalities for which electoral data is unavailable are excluded. * denotes $p < 0.1$, ** denotes $p < 0.05$, *** denotes $p < 0.01$ from two-sided t tests.

adults that watched for at least 10 seconds best approximates the share of citizens exposed to the information on the fraction of spending subject to irregularities, which was reported 17 seconds through the ad. Nevertheless, Facebook users that did not get that far could still have responded to the ad by thinking more on the issue or by discussing it with others. An average of around 15% of targeted voting age adults (or around 20% of targeted Facebook users) thus substantively engaged with the ads, by watching at least 3 seconds of an ad, over the duration of Borde Político's information campaign. Accounting for potential information diffusion and coordination that does not occur on Facebook, the ad campaign could have thus influenced a significant fraction of the electorate in treated segments and municipalities.

Although Facebook's algorithm ultimately dictates when ads appear, the intended 1:4 ratio of exposure to Borde Político's ads across low and high saturation treatments was generally maintained. The p values associated with the first test at the foot of panel A demonstrate that levels of engagement in high saturation municipalities were systematically greater by all non-organic metrics of engagement. While the high saturation treatment's reach was more limited relative to its target, the second test shows that we cannot statistically reject a 1:4 ratio for most measures of engagement.

The extent of Facebook user engagement with the ads does not generally vary with the irregularities reported in the ad. Interacting equation (4) with irregularities quartiles, panel B shows that—with the exception of high saturation municipalities with the highest amount of irregularities—voter engagement did not vary with the level of irregularities reported. For most comparisons, this suggests that differential treatment effects across quartiles of the irregularities distribution are unlikely to be driven by differential access to the Facebook ads. The exception in Q4 may reflect the relative lack of negative reactions from Facebook users, which in turn encourages Facebook's ad assignment algorithm to increase publication of the ad. The results in Q4 should, therefore, be interpreted with greater caution.

5.2 Voter comprehension of the treatment

We next use the small-scale WhatsApp experiment conducted within our panel survey to examine whether respondents understood and internalized the ad's content weeks after receiving it. We leverage within-municipality variation in the receipt of a WhatsApp message containing the Facebook ad by estimating regressions of the form:

$$Y_{im} = \beta WhatsApp\ ad_{im} + \mu_m + \varepsilon_{im}, \tag{5}$$

where Y_{im} is an outcome for individual i in municipality m, and the municipality fixed effects, μ_m, mean that we exploit variation within the municipalities where the Facebook ads were delivered by

Table 2: Comprehension and internalization of the information campaign ads

	Remembers WhatsApp message	**Knows content of message**	**Knows % of irregularities**
	(1)	(2)	(3)
Panel A: average treatment effects			
WhatsApp ad	0.089***	0.025**	0.024***
	(0.016)	(0.010)	(0.007)
R^2	0.083	0.051	0.061
Panel B: heterogeneity by irregularities quartile			
WhatsApp ad	0.044**	0.017	0.015*
	(0.020)	(0.013)	(0.009)
WhatsApp ad × Q3	0.109**	0.017	-0.002
	(0.052)	(0.032)	(0.022)
WhatsApp ad × Q4	0.119***	0.023	0.033**
	(0.036)	(0.023)	(0.015)
R^2	0.091	0.052	0.064
Observations	1,490	1,476	1,434
Outcome range	{0,1}	{0,1}	{0,1}
Control outcome mean	0.056	0.021	0.004
Control outcome std. dev.	0.229	0.144	0.059

Notes: Each specification is estimated using OLS, and includes municipality fixed effects. All observations are weighted by the inverse probability of treatment assignment. Robust standard errors are in parentheses. * denotes $p < 0.1$, ** denotes $p < 0.05$, *** denotes $p < 0.01$ from two-sided t tests.

WhatsApp message to survey respondents.[30] Robust standard errors are used in all analyses, and all observations are weighted by their inverse probability of treatment assignment.

Column (1) of Table 2 first shows that treated respondents were almost 9 percentage points more likely to report having received the WhatsApp message several weeks after the election ($p < 0.01$). Facebook analytics data further indicate that 32% of baseline respondents who received the WhatsApp message treatment clicked on the link to the Facebook page contained within the message.

Respondents also internalized the ad's content. Columns (2) and (3) report that treated voters were 2.5 percentage points more likely to correctly identify (from a list of 4 options) that the message contained information about total FISM resources or the share subject to irregularities

[30]As noted above, we do not examine the effects of municipal saturation due to differential attrition across municipality-level treatment conditions.

($p < 0.05$) and were also 2 percentage points more likely to correctly identify (within 10 percentage point bands) the percentage of funds audited by the ASF that were subject to irregularities ($p < 0.01$). These estimates suggest that a quarter of the respondents who recalled the ad internalized the information. Since the post-election survey was conducted in the weeks after the election, this figure likely understates the share of respondents that recalled the information on election day. Moreover, panel B shows that the information was somewhat more memorable when a large fraction of irregularities was reported, although this translated less systematically into accurate recall of the ad's content.

Together, these results suggest that the ad was likely to have been comprehensible to the many users who watched it on Facebook, and also suggest that was likely to have been recalled by voters at the time of the election.

5.3 Perceived credibility of the ads

While the Facebook ads had significant reach and were likely to have been internalized by interested voters, they may have only influenced voter beliefs or behaviors if they were regarded as credible. We assess this by asking respondents that remember receiving the WhatsApp message about its provenance and credibility. As Figure 4a illustrates, around a quarter of respondents correctly identified the ad's source as being from an NGO. Somewhat surprisingly given the generally favorable indicators of incumbent performance, one third regarded it as coming from opposition parties. Nevertheless, Figure 4b suggests that the majority of voters regarded the information contained within the message as somewhat or very credible. Such levels of credibility are notable in an electoral context where many voters were aware of and concerned by fake news. Moreover, the majority of respondents believed that such a Facebook ad campaign would be seen by many others in their municipality.

6 Effects of Facebook ads on voting behavior

We now present our precinct-level results. We first show that access to Borde Político's large-scale Facebook ads campaign had limited impact on incumbent party vote share and turnout *on average*. However, once we account for the information's content, we find that the information campaign significantly increased the vote share of incumbent parties whose mayors were shown to have presided over zero or negligible irregularities. Most importantly, we demonstrate that these increases in support for the best-performing incumbents are greater among the treated segments of high saturation than low saturation municipalities. We discuss spillover effects in the mechanisms section (see section 7).

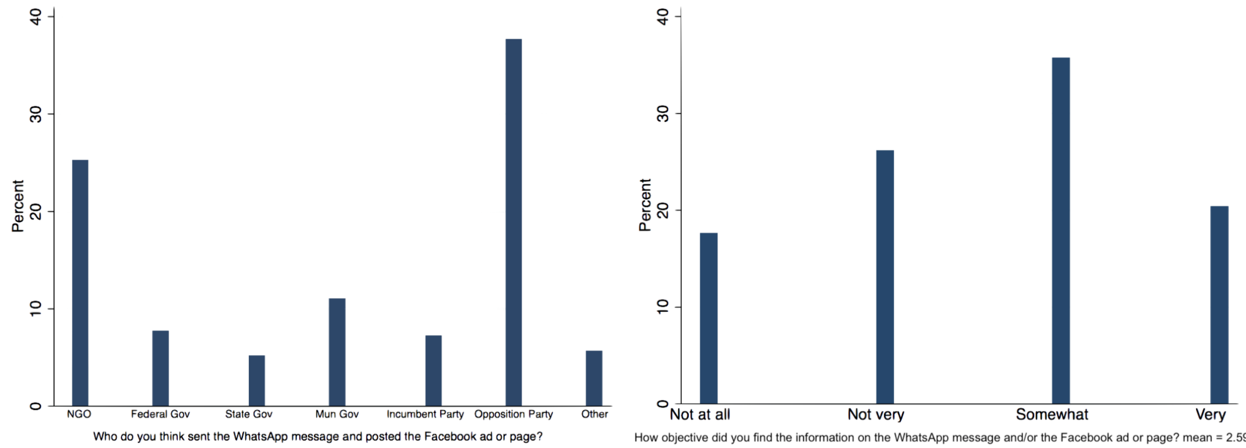

(a) Perception of source

(b) Perception of credibility

Figure 4: Source and credibility of WhatsApp message/Facebook ad

Note: All data is from the subsample of endline respondents that remembered receiving the WhatsApp message.

6.1 Average treatment effects

Table 3 examines the effects of the Facebook ad campaigns on municipal incumbent party vote share and turnout.[31] For each outcome, we examine the intent to treat effects of the segment-level treatment assignments. The block fixed effects and lagged dependent variables account for around 60% of the variation in our outcomes; this entails relatively precise standard errors of around a single percentage point.

We first find tentative evidence that the Facebook ads increased incumbent party support in the *average* municipality. Integrating across saturation and irregularities levels, column (1) shows that the segment-level Facebook ads treatment increased the incumbent party's vote share by 2.1 percentage points ($p < 0.1$). This represents around a 12% increase in incumbent support, relative to the 18% of registered voters that turned out for the incumbent party in the control group. Column (2) reports a 2.4 percentage points effect in high saturation municipalities, although the difference relative to the 0.7 percentage point effect in low saturation municipalities is not quite statistically significant at the 10% level ($p = 0.11$, one-sided). These estimates are consistent with voters in the many municipalities observing 0% malfeasance either updating favorably about, or seeking to coordinate around, the incumbent party. The next subsection shows that, as expected, the average effect pools across differential effects by the level of reported irregularities.

The Facebook ad campaign also slightly increased aggregate turnout. Column (3) reports a 1.3 percentage point increase in turnout due to the ad campaign ($p < 0.1$). Column (4) again indicates

[31] To save space, we restrict attention to incumbent party vote share, as a share of registered voters. Table A13 shows similar results for incumbent vote share, as a share of turnout.

Table 3: Effect of Facebook ads on precinct-level municipal incumbent party vote share (share of registered voters) and turnout

	Incumbent party vote (share of registered voters)		**Turnout**	
	(1)	(2)	(3)	(4)
Facebook ads	0.021*		0.013*	
	(0.011)		(0.008)	
Spillover	0.009		0.005	
	(0.011)		(0.008)	
Facebook ads in high saturation		0.024*		0.013
		(0.013)		(0.009)
Facebook ads in low saturation		0.007		0.001
		(0.013)		(0.010)
Spillover in high saturation		0.020		0.008
		(0.013)		(0.009)
Spillover in low saturation		0.005		0.002
		(0.013)		(0.009)
Observations	13,251	13,251	13,251	13,251
R^2	0.55	0.57	0.63	0.64
Control outcome mean	0.18	0.18	0.64	0.64
Control outcome std. dev.	0.11	0.11	0.12	0.12
Test: spillover $\geq$ direct (p value, 1-sided)	0.070		0.065	
Test within ads treatment: low $\geq$ high (p value, 1-sided)		0.110		0.081
Test within spillovers: low $\geq$ high (p value, 1-sided)		0.142		0.241

Notes: Each specification is estimated using OLS, and includes a lagged dependent variable and randomization block fixed effects. All segments are weighted equally and by the inverse probability of treatment assignment. Standard errors clustered by municipality are in parentheses. * denotes $p < 0.1$, ** denotes $p < 0.05$, *** denotes $p < 0.01$ from two-sided t tests.

that this is predominantly driven by the Facebook ads in high saturation municipalities, where the differential is statistically significant at the $p < 0.1$ level. The estimated effects on turnout are broadly in line with the positive effects of Facebook's own mobilization campaigns in the U.S. (Bond et al. 2012). However, they are substantively small in magnitude, representing only a 2% increase in turnout relative to the baseline 64% in the control group.

6.2 Heterogeneity by reported irregularities

The impact of providing Facebook ads reporting *any* level of incumbent expenditure irregularities resulted in a borderline statistically significant increase in support for the incumbent party. However, we next show that—as with many informational interventions—the small average effects mask substantial heterogeneity with respect to the content of the information provided. We focus on the incumbent party vote, since the lack of differential effects on aggregate turnout (see Appendix Table A15) indicate that the results are driven primarily by shifts in municipal incumbent party support.

The results in Table 4 show that Borde Político's Facebook information campaign significantly increased the vote share of the least malfeasant municipal incumbent parties. In the municipalities where directly treated segments were informed of zero or negligible irregularities (i.e. quartiles Q1 and Q2, which are pooled because more than 25% of municipalities registered zero irregularities), column (1) shows that the Facebook ads treatment increased the incumbent party's vote share by 4.1 percentage points ($p < 0.05$)—an increase of more than 20%, relative to the control group mean. In contrast, the negative interaction coefficients for the third and fourth quartiles (i.e. Q3 and Q4) indicate that the conditional average treatment effect is—as expected—smaller for higher levels of reported irregularities. The difference in the effect of Facebook ads between Q1/Q2 and Q4 is statistically significant ($p < 0.1$). The tests at the foot of column (1) indicate that the overall effect in each quartile is only significantly different from zero for the municipal incumbent parties with the cleanest spending records. The stronger electoral impact of good performance is in line with Arias, Larreguy, Marshall and Querubín's (2019) prior findings in a smaller set of Mexican municipalities in 2015 (Arias, Larreguy, Marshall and Querubín 2019), where they argue that limited sanctioning of poor performance reflects voters' pessimistic prior beliefs and the benefits of reduced uncertainty about the incumbent's type. The limited sanctioning in Q4 could also reflect a lack of scope to reduce the incumbent party's vote share, given MORENA's electoral success in 2018 and the fact that MORENA was the incumbent in only one treated municipality.

While access to ads is randomly assigned, the share of irregularities reported is not. To address the possibility that heterogeneity in response to Facebook ads across municipalities where the ASF found different levels of irregularities instead reflects other differences across these municipalities, we further adjust for the interaction between treatment conditions and 11 predetermined covariates at the municipal level. First, we adjust for the other quantitative information conveyed by the ad—the financial year to which the audit pertained and the amount of FISM funds received by the municipality in that year—to address the possibility that other municipality-specific elements of the ad drive voting behavior. Second, we adjust for the prior municipal incumbent party vote share to address the possibility that audits report revelations differentially impact the parties of more or less popular local governments. Third, we further adjust for eight demographic and socioeconomic variables that could both be correlated with reported irregularities and facilitate coordinated responses to ads or proxy for voters' prior beliefs, attentiveness or access to Facebook ads, or capacity to comprehend such ads.[32] Summary statistics for all interactive covariates are shown in Appendix

[32]These municipal-level covariates are: the 2010 adult population; average years of schooling in 2010; the share illiterate in 2010; the average number of occupants per room in 2010; the average number of children per woman in 2010; the share of households with electricity, water, and drainage in 2010; the share of the municipal population that is working age in 2010; and the share of households with internet at home in 2010.

Table 4: Effect of Facebook ads on precinct-level municipal incumbent party vote share (share of registered voters), by quartile of the sample irregularities distribution

	Incumbent party vote (share of registered voters)	
	(1)	(2)
Facebook ads	0.041**	0.057***
	(0.016)	(0.020)
Facebook ads × Q3	-0.046	-0.026
	(0.032)	(0.034)
Facebook ads × Q4	-0.051*	-0.079*
	(0.028)	(0.040)
Spillover	0.023	0.034*
	(0.015)	(0.018)
Spillover × Q3	-0.001	-0.008
	(0.026)	(0.031)
Spillover × Q4	-0.046	-0.071*
	(0.032)	(0.039)
Observations	13,251	13,251
R^2	0.59	0.65
Control outcome mean	0.18	0.18
Control outcome std. dev.	0.11	0.11
Test: null effect of Facebook ads in Q3 (p value, 2-sided)	0.849	0.244
Test: null effect of Facebook ads in Q4 (p value, 2-sided)	0.661	0.389
Test: null effect of spillover in Q3 (p value, 2-sided)	0.259	0.246
Test: null effect of spillover in Q4 (p value, 2-sided)	0.418	0.179
Interactive covariates		✓

Notes: Each specification is estimated using OLS, and includes a lagged dependent variable and randomization block fixed effects. Specifications including interactive covariates further include interactions between treatment conditions and the following municipal-level covariates: year of audit; amount of FISM funds received; population aged above 18; lagged incumbent party vote share; average years of schooling; share of the population that is illiterate; average occupants per room; average number of children per woman; the share of the population with electricity, water, and drainage in their home; the working age share of the population; and the share of households with internet at home. The omitted irregularities category is Q1/Q2. All segments are weighted equally and by the inverse probability of treatment assignment. Standard errors clustered by municipality are in parentheses. * denotes $p < 0.1$, ** denotes $p < 0.05$, *** denotes $p < 0.01$ from two-sided t tests.

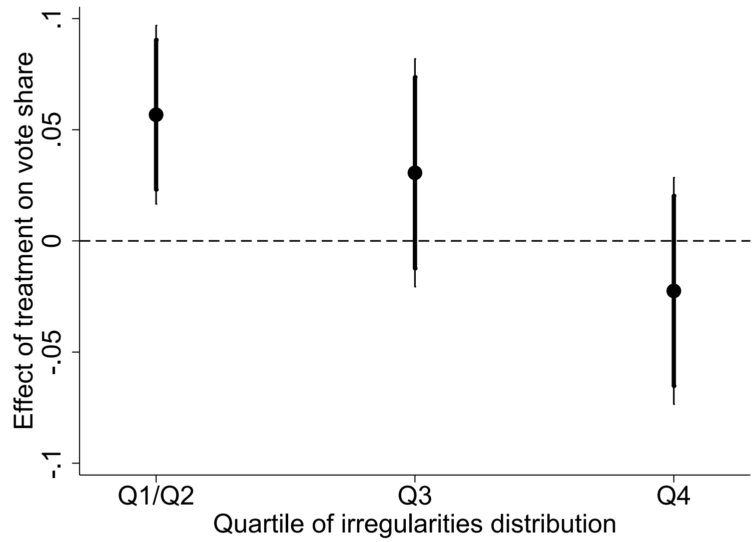

Figure 5: Conditional average treatment effect of Facebook ads, by irregularities quartile (with 90% and 95% confidence intervals)

Note: All estimates are from specification in column (2) of Table 4 with interactive covariates.

Table A3.

The estimates that adjust for interactions with these potential confounds in column (2) are similar to the results in column (1) without such interactive covariates. Again, we find a clear positive effect of Facebook ads in Q1/Q2 ($p < 0.01$). In Q4, we observe a decline in the incumbent party's vote share of 2.2 percentage points, although the test at the foot of the table indicates that this effect is not statistically significant at conventional levels. Appendix Table A16 reports similar results when quartiles Q3 and Q4 are pooled or irregularities are instead dichotomized as above 0% or 5% irregularities.

The overall effects, which are illustrated graphically in Figure 5 for the specification including interactive covariates, are notable for their magnitude in comparison with previous findings. The 5.7 percentage point increase in incumbent party vote share in Q1/Q2 exceeds estimates from most prior experimental and quasi-experimental studies. Prior studies—which generally involve far lower levels of saturation—typically report effects of up to 3 percentage points, even when murder charges against candidates were reported via SMS in India (George, Gupta and Neggers 2018).[33] Only in the case of relatively widespread media reporting of similar types of malfeasance in Brazil (Ferraz and Finan 2008) do we observe comparable impacts on vote shares. As the next subsection shows, the large effects are driven by the high saturation variant of the campaign.

[33]Most prior studies measure incumbent party vote share as a proportion of voters that turned out. The direct comparison in these cases is to Appendix Table A14, where normalizing by those that turned out increases the magnitude of the coefficients.

6.3 Differential effects of information campaign saturation

Having established that Borde Político's accountability campaign rewarded "clean" municipal incumbent parties, we next examine the extent to which campaign saturation accentuates voters' electoral response to information provision. As argued above, large-scale information campaigns may stimulate belief updating and voter coordination that would not have occurred if incumbent performance information had only been provided to a small fraction of the electorate.

We test this hypothesis by leveraging the random assignment of treated municipalities to receive 20% or 80% information campaign saturation levels. Since Table 1 showed that the low saturation treatment was slightly more effective than the high saturation treatment at reaching targeted voters, our estimates of equation (2) pertain to a case where effective saturation in high saturation municipalities was roughly 3 times greater than saturation in low saturation municipalities.

The results in columns (1) and (2) of Table 5 demonstrate that greater levels of ad saturation indeed amplify the effects of Facebook ads reporting zero or negligible irregularities. Regardless of whether interactions between treatment conditions and municipal-level covariates are adjusted for, reporting performance in Q1 or Q2 of the irregularities distribution increased the share of voters that voted for the incumbent party in treated segments by 5-6 percentage points ($p < 0.05$)—or around half a standard deviation—in high saturation municipalities. In contrast, the fourth row of the table shows that the effect in Q1/Q2 is slightly less than half this size in the directly treated segments of low saturation municipalities, and is statistically indistinguishable from zero. Furthermore, the tests at the foot of the table show that the 3 percentage point difference between the small effects in low saturation municipalities and the large effects in high saturation municipalities is statistically significantly ($p < 0.05$, one-tailed, when interactive covariates are adjusted for). Consistent with the lack of an effect in Q3 and Q4, we do not observe clear differential effects across saturation levels within Q3 and Q4.[34]

These comparisons between directly treated segments and control municipalities indicate that a campaign's saturation can substantially increase its effect within areas directly targeted by Facebook ads. This finding aligns with extant evidence suggesting that information campaigns are more likely to influence the voting behavior of treated voters when they are disseminated at a large scale, and is thus consistent with saturation driving the larger effects attributed to broadcast media outlets. While the framing and credibility functions of the media may also be important, our findings thus suggest that saturation is central to the impact of information distributed via modern communication technologies. We next explore some of the mechanisms underpinning these effects.

[34]The null effect in Q4 is more negative in low saturation municipalities in column (2), although this effect is not robust in column (1).

Table 5: Effect of Facebook ads on precinct-level municipal incumbent party vote share (share of registered voters), by information campaign saturation

	Incumbent party vote (share of registered voters)	
	(1)	(2)
Facebook ads in high saturation	0.049***	0.060**
	(0.018)	(0.026)
Facebook ads in high saturation × Q3	-0.069*	-0.044
	(0.038)	(0.041)
Facebook ads in high saturation × Q4	-0.062**	-0.049
	(0.031)	(0.054)
Facebook ads in low saturation	0.023	0.027
	(0.018)	(0.022)
Facebook ads in low saturation × Q3	-0.003	-0.032
	(0.029)	(0.036)
Facebook ads in low saturation × Q4	-0.048	-0.077*
	(0.035)	(0.046)
Spillover in high saturation	0.047**	0.056**
	(0.018)	(0.026)
Spillover in high saturation × Q3	-0.076*	-0.050
	(0.039)	(0.042)
Spillover in high saturation × Q4	-0.060**	-0.045
	(0.030)	(0.055)
Spillover in low saturation	0.014	0.020
	(0.017)	(0.021)
Spillover in low saturation × Q3	0.026	-0.002
	(0.029)	(0.035)
Spillover in low saturation × Q4	-0.039	-0.065
	(0.036)	(0.045)
Observations	13,251	13,251
R^2	0.63	0.70
Control outcome mean	0.18	0.18
Control outcome std. dev.	0.11	0.11
Test: larger effect of Facebook ads in high (vs. low) in Q1/Q2 (*p* value, 1-sided)	0.065	0.025
Test: larger effect of Facebook ads in high (vs. low) in Q3 (*p* value, 1-sided)	0.092	0.262
Test: larger effect of Facebook ads in high (vs. low) in Q4 (*p* value, 1-sided)	0.704	0.041
Test: larger effect of spillovers in high (vs. low) in Q1/Q2 (*p* value, 1-sided)	0.020	0.020
Test: larger effect of spillovers in high (vs. low) in Q3 (*p* value, 1-sided)	0.052	0.580
Test: larger effect of spillovers in high (vs. low) in Q4 (*p* value, 1-sided)	0.803	0.026
Interactive covariates		✓

Notes: Each specification is estimated using OLS, and includes a lagged dependent variable and randomization block fixed effects. Specifications including interactive covariates further include interactions between treatment conditions and the following municipal-level covariates: year of audit; amount of FISM funds received; population aged above 18; lagged incumbent party vote share; average years of schooling; share of the population that is illiterate; average occupants per room; average number of children per woman; the share of the population with electricity, water, and drainage in their home; the working age share of the population; and the share of households with internet at home. The omitted irregularities category is Q1/Q2. All segments are weighted equally and by the inverse probability of treatment assignment. Standard errors clustered by municipality are in parentheses.* denotes $p < 0.1$, ** denotes $p < 0.05$, *** denotes $p < 0.01$ from two-sided t tests. The one-sided coefficient tests at the foot of the table test whether the effect of high saturation is larger in a given direction in high saturation than in low saturation municipalities.

7 Mechanisms

Our main results show that Facebook ads reporting municipal government expenditure irregularities increased votes for the incumbent parties that oversaw negligible irregularities, and that this effect was amplified in high saturation municipalities. If the ad campaign only influenced those that viewed the ad directly on Facebook, our results suggest high rates of persuasion. Indeed, the interpolated persuasion rate implied by assuming that only the voting behavior of the approximately 7% of targeted voters who reached the 17th second of the (paid-for or organically-generated) ad—when the share of irregularities appeared in the video—is 84%.[35] Even assuming that only 3 seconds of the ad is enough to increase support for the incumbent party by raising the issue's salience, the implied persuasion rate is still 39%.

Given that such rates far exceed those documented in prior studies of voting (DellaVigna and Gentzkow 2010), the relatively large effects of Borde Político's accountability campaign on voting behavior suggest either unusually high persuasion rates among atomized viewers of the ad, substantial interaction—such as through information diffusion or voter coordination—between voters that did and did not consume the ad, or that other actors capable of influencing large numbers of voters responded to Borde Político's ad campaign. We next document sizable spillovers within treated municipalities and limited responses to the information campaign from political campaigns and media outlets. Together, these findings indicate that social interactions between voters can underpin the effectiveness of mass online information campaigns.

7.1 Information campaigns and social interactions

Our survey data indicate that the information contained in Borde Político's campaign, and voters' reactions to it, could plausibly have influenced the vote choices of voters that did not directly view the ad through social interactions. This would most likely operate through general political discussion between individuals. For instance, 76% of survey respondents reported discussing politics at least once a week with family and acquaintances in person or over the phone and 64% reported doing so through social media (i.e. Facebook, Twitter, and WhatsApp). Also, 55% of survey respondents who recalled receiving the WhatsApp message reported discussing its content with their family and acquaintances. Furthermore, 66% of respondents suggested that discussions with others

[35] Linear interpolation based on columns (9) and (11) of Table 1 suggests that $100 \times \frac{0.015+(0.021-0.015)\times\frac{26-17}{26-10}}{0.2} \approx 9.2\%$ of targeted voters in low saturation municipalities and $100 \times \frac{0.038+(0.060-0.038)\times\frac{26-17}{26-10}}{0.8} \approx 6.3\%$ of targeted voters in high saturation municipalities reached the 17th second of the 26-second ad. Given that treated segments in high saturation municipalities are 4 times more prevalent than treated segments in low saturation municipalities in our sample, this yields a sample average exposure rate of around 6.9%. Following DellaVigna and Gentzkow (2010), the persuasion rate implied by column (2) of Table 4 is then: $100 \times \frac{0.057}{0.069}\frac{1}{1-0.18} \approx 84\%$.

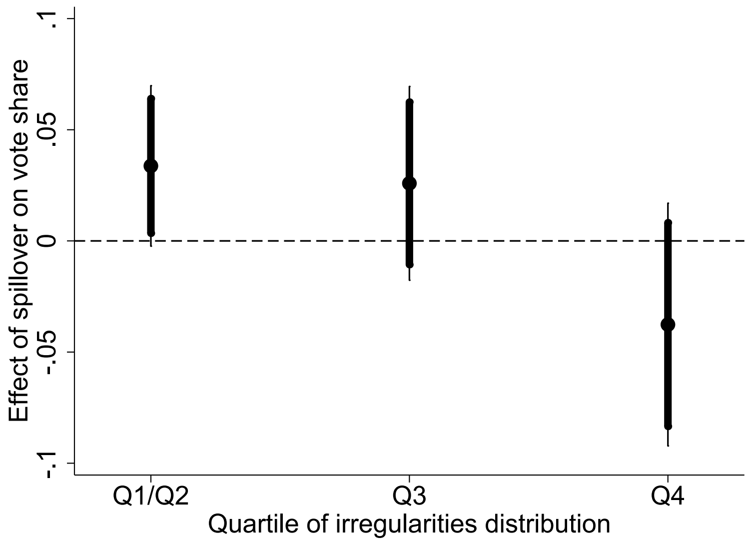

Figure 6: Conditional average treatment effect of spillovers, by irregularities quartile (with 90% and 95% confidence intervals)

Note: All estimates are from specification in column (2) of Table 4 with interactive covariates.

before the election helped them choose which candidate to vote for, while 84% of respondents reported that their expectations of others' vote choices affected their own vote choice. In contrast, the ad campaigns appear to have initiated relatively little activity online through Facebook's platform. As Table 1 shows, organic views of the ad only increased viewers beyond the paid-for Facebook ads by around 10%.

To establish whether the social interactions suggested by survey respondents drive voting behavior, we compare segments within treated municipalities that were randomly assigned not to receive direct access to Facebook ads with segments within control municipalities. If Facebook's ad generation algorithm only reaches voters within targeted locations, voters in spillover segments could only have been influenced by Borde Político's ad campaign via some form of interaction with voters in directly treated segments.

As shown in Tables 4 and 5, the Facebook ads produced spillover effects in line with the larger effects observed in directly-treated electoral precincts. Figure 6 first shows that the presence of Facebook ads within the average treated municipality increased the incumbent party's vote share in untreated segments of municipalities in Q1/Q2 of the irregularities distribution by 3.4 percentage points. This effect is around half the size of the effect in directly treated segments and is not quite statistically significant. Column (2) of Table 5 shows that this effect rises to 5.6 percentage points in spillover segments within high saturation municipalities ($p < 0.05$), and falls to a statistically insignificant 2 percentage points in low saturation municipalities. As with the effects of Facebook ads in segments that were treated directly, the tests at the foot of Table 5 demonstrate that the greater spillover effect within high saturation municipalities, relative to low saturation municipalities, is also statistically significant ($p < 0.05$, one-sided test). This again suggests that explicit or tacit cross-segment interactions between voters amplify an information campaign's effects. Saturation thus amplifies both the direct effect of information *and* the associated spillovers. The latter suggests that social interactions amplify effects throughout targeted municipalities, and thus also imply that social interactions could account for much of the effect within treated segments as well. If the difference between the effects in directly and indirectly treated segments is interpreted as the persuasion effect of watching the ad directly, our results suggest that most of the impact of Facebook ads is due to interactions occurring *after* the ad is viewed on Facebook.

While most Facebook users saw Borde Político's informational ads on mobile devices providing GPS data that enable Facebook to target ads at users based on their home location, spillover effects could still potentially reflect direct treatment arising from inaccurate geographic targeting of Facebook ads. However, two tests suggest that mistargeting does not drive voting behavior. First, we should expect to observe larger spillover effects in electoral precincts with access to 3G+ (3G, 4G, or LTE) mobile telecommunication signals if Facebook mistargeted ads to users that live—and thus vote—in untreated segments. A lack of 3G technology, which renders videos on mobile devices slow and pixelated, is likely to impede other forms of social interaction far less. Restricting our sample to spillover and control segments, we assess this by examining the interaction between spillovers and an indicator for electoral precincts where 3G+ coverage reaches at least 75% of the population's home. The statistically insignificant differential effect coefficients for the spillover effect in Q1/Q2 in columns (1) and (2) of Appendix Table A17 provide no evidence supporting the possibility of mistargeting.

Second, if treatments were mistargeted and voters considered politicians from the same party to be similarly malfeasant, we should expect to observe changes in the vote share of the incumbent party of a nearby treated municipality in precincts that are adjacent to that nearby treated municipality. To test for this possibility, while still leveraging experimental variation, we restrict our sample to precincts from non-experimental municipalities that are within 5 kilometers of only one experimental segment (see Appendix section A.13 for further details). Appendix Table A18 shows that proximity to neither (directly or indirectly) treated segments nor high or low municipal saturation levels in the nearby experimental municipality systematically affected the relevant vote share, within any irregularities quartile. These null estimates thus also provide no evidence to suggest that Facebook ads were significantly mistargeted.

7.2 Limited political and media responses to the Facebook ad campaigns

The main alternative mechanism that could account for both the campaign's large direct and spillover effects and saturation's amplifying effects is that the intervention also changed the municipal-level strategies of other actors with the capacity to influence voters *en masse*. Most plausibly, the increase in incumbent party vote share in Q1/Q2 in directly and indirectly treated segments could result from incumbent parties or media outlets incorporating this information into large-scale campaign activities or news reports. Scope to do so was limited because the ads were distributed in the last week of the electoral campaign, while the results of the ASF reports were already available to parties and media outlets before Borde Político's intervention. Nevertheless, recent studies document sophisticated campaign responses to pre-election information dissemination campaigns (Arias, Larreguy, Marshall and Querubín 2019; Banerjee et al. 2011; Cruz, Keefer and Labonne 2019) and political debates (Bidwell, Casey and Glennerster 2019; Bowles and Larreguy 2018).

While we cannot determine exactly what occurred on the ground, we assess these alternative explanations by examining politician engagement with Borde Político's campaign online and reporting on the ASF reports (and corruption more generally) by local newspapers. We find little support for either alternative potential mechanism.

7.2.1 Online campaign responses to the Facebook ads

Online campaigning is now common in Mexico, where candidates use social media accounts to announce their campaign promises, publicize their slogans, and criticize other candidates. However, we were able to detect very few responses by candidates for the municipal presidency to the Facebook ads and associated pages.[36] Across all the Facebook and Twitter accounts that we identified as belonging to candidates in our sample of municipalities, we were only able to detect two responses to Borde Político's information campaign on Facebook and none on Twitter. The challenger candidate of the Alternative Sonora Movement in Huatabampo, Sonora, shared on his Facebook page that 30% of the FISM funds spent by the PRI municipal government were subject to irregularities. Similarly, the Citizen's Movement challenger candidate eventually elected as municipal president in Venustiano Carranza, Michoacán, shared that the PRD municipal government incurred 14% in irregularities.

We additionally scraped thousands of comments, reactions, and shares relating to all Borde Político's Facebook ads and pages. Again, we observe negligible activity among Facebook users identified as running for other offices. A single PAN candidate for federal deputy liked the ad reporting 30% of irregularities by the PRI municipal government in Ciudad Valles, and one PRD candidate for federal deputy challenged the 61% of irregularities of the PRD municipal government reported in Cuautla, Morelos, arguing that there were no irregularities.[37] These scattered responses

[36]Note that we cannot identify those Facebook users that shared and reacted to the Facebook ads and pages that do not have a public profile. However, this is unlikely to be a problem since candidates and political operatives are likely to have a public profile.

[37]Among the many other reactions by Facebook users, we were only able to identify 8 reactions that appeared to be from possible party operatives (people who regularly posted in favor of their party/candidate or explicitly mentioned working for the party or candidate's campaign).

cannot plausibly account for the changes in voting behavior in treated segments at the scale that we observe.

7.2.2 Local media reporting of the ASF audits after the ad campaign

To examine whether media reporting related to Borde Político's Facebook ad campaign amplified the campaign's direct effects and induced within-municipality spillovers, we collected online data from 263 local newspapers that serve 92 of the 124 (74%) municipalities in our final sample.[38] The majority of local newspapers in Mexico provide significant amounts of content—including from their print editions—on their websites, sometimes including full versions of the print editions as well. While we could not obtain radio and television content, local newspapers are an important source of news content for local broadcast media outlets (Larreguy, Lucas and Marshall 2016).

Over the 10 days between the start of the Facebook ad campaigns and the election, we searched for a variety of terms related to both the specific content of the Facebook ads and the more general references to corruption. General references to corruption were included because the ad campaign could have increased demand for related information, which could also have influenced voting behavior if such demand was met by media outlets before the election. We estimate the effect of the campaign on such media reports using municipal-level regressions analogous to equation (4).

We again find little evidence to suggest that Borde Político's Facebook ad campaign induced a response from local media outlets. First, we were unable to detect any newspaper articles referencing the Facebook campaign, Borde Político's dissemination of information, or the ASF's reports in the pre-election period. Given that media outlets do sometimes report on the outcomes of ASF audits (Larreguy, Marshall and Snyder 2019), this suggests that the newspapers in the municipalities in our sample were unaware of the Facebook ads, had already reported on ASF (and did not see the need to discuss them further), or lacked incentives or resources to report on the issue during a nationwide election campaign. Second, we find that the ad campaigns did not significantly increase reporting more generally on issues related to corruption before the election. Indeed, Table 6 shows that neither total mentions of corruption nor circulation-weighted mentions of corruption by local newspapers were systematically affected by the presence of a municipal Facebook ad campaign, the level of irregularities reported, or the campaign's saturation. These results suggest that the changes in vote choice induced by the Facebook ads were not driven by media coverage of the campaign.

[38] Since our randomization ensures that newspaper circulation within a given municipality is orthogonal to municipal treatment assignment and some small municipalities may not be served by local newspapers, we retain all municipalities for this analysis.

Table 6: Effect of Facebook ad campaigns on municipal-level newspaper reporting on corruption in general

	Total articles on corruption						Total articles on corruption (circulation-weighted)					
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)
Any Saturation	-6.209	-5.483	-14.962				-6.209	-5.533	-15.042			
	(4.393)	(8.429)	(12.038)				(4.382)	(8.429)	(12.028)			
Any Saturation × Q3		-10.003	2.179					-9.803	2.352			
		(13.033)	(15.896)					(12.964)	(15.859)			
Any Saturation × Q4		4.608	25.325					4.593	25.391			
		(9.819)	(22.613)					(9.808)	(22.606)			
High saturation				-4.125	-4.548	-24.509				-4.173	-4.609	-24.494
				(4.746)	(9.965)	(17.266)				(4.735)	(9.965)	(17.256)
High saturation × Q3					-3.511	17.325					-3.390	17.381
					(13.155)	(23.292)					(13.121)	(23.260)
High saturation × Q4					0.998	54.470					0.875	54.246
					(12.537)	(33.992)					(12.504)	(33.982)
Low saturation				-8.370*	-6.636	-18.081				-8.321*	-6.676	-18.112
				(4.687)	(7.282)	(14.416)				(4.675)	(7.283)	(14.409)
Low saturation × Q3					-15.577	3.257					-15.305	3.449
					(15.274)	(19.237)					(15.174)	(19.190)
Low saturation × Q4					7.718	35.836					7.790	35.886
					(8.791)	(27.853)					(8.800)	(27.856)
Observations	124	124	124	124	124	124	124	124	124	124	124	124
R^2	0.37	0.40	0.60	0.37	0.42	0.68	0.37	0.40	0.60	0.37	0.42	0.68
Control outcome mean	8.51	8.51	8.51	8.51	8.51	8.51	8.46	8.46	8.46	8.46	8.46	8.46
Control outcome std. dev.	31.53	31.53	31.53	31.53	31.53	31.53	31.45	31.45	31.45	31.45	31.45	31.45
Test: low = high (p value, 2-sided)				0.208						0.218		
Test: effect of treatment in Q3 (p value, 2-sided)		0.126	0.165					0.127	0.166			
Test: effect of treatment in Q4 (p value, 2-sided)		0.794	0.396					0.778	0.397			
Test: effect of low saturation in Q3 (p value, 2-sided)					0.123	0.222					0.125	0.225
Test: effect of low saturation in Q4 (p value, 2-sided)					0.710	0.251					0.704	0.251
Test: effect of high saturation in Q3 (p value, 2-sided)					0.216	0.585					0.215	0.587
Test: effect of high saturation in Q4 (p value, 2-sided)					0.608	0.109					0.587	0.111
Interactive covariates			✓			✓			✓			✓

Notes: Each specification is estimated using OLS, and includes randomization block fixed effects. The omitted irregularities category is Q1/Q2. Robust standard errors are in parentheses. * denotes $p < 0.1$, ** denotes $p < 0.05$, *** denotes $p < 0.01$ from two-sided t tests.

7.3 Discussion of mechanisms

We have shown that information campaign saturation induced and amplified voter responses to Borde Político's Facebook ads primarily through off-Facebook social interactions. Establishing whether social interactions changed voting behavior by inducing voters to update their beliefs about the incumbent's type and/or effort or coordinate their vote choices is challenging for several reasons that we now discuss. We also consider research designs that might illuminate these factors.

First, the comparative statics implied by the belief updating and coordination mechanisms are often identical (Arias, Balán, Larreguy, Marshall and Querubín 2019; Zhuravskaya, Petrova and Enikolopov forthcoming). Any effects produced by either channel will generally increase in the extent of social interactions between voters, e.g. through information diffusion or greater scope for coordination. While shutting down the coordination channel is essentially impossible in the field, ensuring that information about the incumbent is *fully internalized* by treated voters—such that further interaction with the information, due to informational spillovers, could not further alter their beliefs about incumbent irregularities—is almost as difficult. Indeed, this cannot realistically be achieved at scale with mass media interventions that provide voters with access to information.

However, as Chen and Yang (2019) show, it is possible to induce voters to consume information. One approach to illuminating mechanisms other than belief updating about the incumbent's type and/or effort, induced by providing ASF audit results, would then be to incentivize a small number of voters to acquire and internalize the information provided by the mass campaign (e.g. Chen and Yang 2019), and then estimate treatment effects for compliers across municipalities with different levels of saturation or voters receiving different information about the campaign's scale. Alternatively, a detailed in-person explanation of the information may be even more effective at ensuring that some citizens could not plausibly update further after receiving the information from others (e.g. Bhandari, Larreguy and Marshall 2019). Such designs could be implemented through parallel surveys, and—under the (partially testable) assumption that voters do not update from the information on other margins after discussion with others—facilitate estimation of differential saturation effects that are not driven by information diffusion. Unfortunately, this study lacked the resources to implement such high-powered information engagement treatments.

Second, even survey questions asking directly about particular mechanisms may struggle to distinguish between mechanisms. The challenge is exacerbated in our context by the difficulty of ensuring that respondents engage with the mass-level version of the treatment that may be necessary for the activation of social mechanisms that occur outside the survey. Interpreting the survey evidence in our case was further limited by relatively low levels of engagement with the Facebook ads in our respondent sample and differential attrition across municipal treatment conditions. Nevertheless, we do observe some evidence consistent with both the belief updating and voter coordination mechanisms: noisy estimates for belief updating within our panel survey broadly align with voting behavior (see Appendix Table A19), while the descriptive data above suggests that information sharing and vote coordination are common in this setting.

8 Conclusions

We show that non-partisan information campaigns on social media, especially where their level of electorate saturation is high, can substantially promote electoral accountability. In particular, our evaluation of Borde Político's anti-malfeasance campaign experimentally demonstrates that a large-scale Facebook ad campaign in Mexico substantially increased support for the municipal incumbent parties reported to have performed well. This effect is particularly large among treated voters in high saturation municipalities, and appears to be driven by (off-Facebook) interactions between voters induced by the campaign, rather than responses by political campaigns or rebroadcasting by other media outlets.

These results suggest that the high saturation of many mass broadcast or social media messages may explain the greater impact of information dissemination conducted by mass media outlets, which we systemically document in Table A1. In this regard, our findings advance existing understandings of the effects of information dissemination on electoral accountability, on which extant evidence is decidedly mixed (Dunning et al. forthcoming). Of course, media outlets also do more than just provide access to information for large audiences. They can distort, filter, and frame content in different ways (see e.g. Prat and Strömberg 2013; Strömberg 2015). Future research might more directly establish the extent to which saturation and these other potential mechanisms drive the role of mass media in promoting or hindering electoral accountability.

While we show that information campaign saturation and social effects are key mechanisms underpinning information's effects on voting behavior, finer-grained distinctions such as that between beliefs updating and coordination remain important topics for future research. Such distinctions may have important implications for policy makers seeking to establish in which contexts large-scale digital information provision could do most to increase accountability. Moreover, creative refinements to information campaigns could potentially be designed to accentuate particular mechanisms.

Finally, the relatively large effects of online information campaigns on vote shares pose interesting questions for democracy. While we focused on the potential benefits for electoral accountability of a non-partisan NGO campaign providing information gleaned from publicly-accessible independent audit reports, our findings also suggest that digital technologies could be used for more nefarious means. The use of social media around elections has recently come under intense scrutiny due to Russia's involvement in the 2016 U.S. presidential election, but also the frequent spreading of politically-relevant content via WhatsApp in developing contexts. Our results suggest that such online information campaigns could have substantial electoral impacts, and thus highlight the importance of election regulators in ensuring that elections are not hijacked by fake news.

References

Acemoglu, Daron, Tarek A. Hassan and Ahmed Tahoun. 2017. “The power of the street: Evidence from egypt’s arab spring.” *Review of Financial Studies* 31(1):1–42.

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. forthcoming. “When Does Information Influence Voters? The Joint Importance of Salience and Coordination.” *Comparative Political Studies* .

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.

Allcott, Hunt, Luca Braghieri, Sarah Eichmeyer and Matthew Gentzkow. 2019. “The welfare effects of social media.” NBER working paper 25514.

Allcott, Hunt and Matthew Gentzkow. 2017. “Social media and fake news in the 2016 election.” *Journal of Economic Perspectives* 31(2):211–36.

Alt, James E., Amalie Jensen, Horacio Larreguy, David D. Lassen and John Marshall. 2019. “Contagious political concerns: How unemployment information passed between weak ties influences Danish voters.” Working paper.

Anderson, Simon P. and John McLaren. 2012. “Media mergers and media bias with rational consumers.” *Journal of the European Economic Association* 10(4):831–859.

Arias, Eric, Horacio Larreguy, John Marshall and Pablo Querubín. 2019. “Priors rule: When do malfeasance revelations help or hurt incumbent parties?” Working paper.

Arias, Eric, Pablo Balán, Horacio Larreguy, John Marshall and Pablo Querubín. 2019. “Information provision, voter coordination, and electoral accountability: Evidence from Mexican social networks.” *American Political Science Review* 113(2):475–498.

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

Avis, Eric, Claudio Ferraz and Frederico Finan. 2018. "Do government audits reduce corruption? Estimating the impacts of exposing corrupt politicians." *Journal of Political Economy* 126(5):1912–1964.

Baird, Sarah, J. Aislinn Bohren, Craig McIntosh and Berk Özler. 2018. "Optimal design of experiments in the presence of interference." *Review of Economics and Statistics* 100(5):844–860.

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

Baron, David P. 2006. "Persistent media bias." *Journal of Public Economics* 90(1):1–36.

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

Berry, Christopher R. and William G. Howell. 2007. "Accountability and local elections: Rethinking retrospective voting." *Journal of Politics* 69(3):844–858.

Besley, Timothy and Andrea Prat. 2006. "Handcuffs for the grabbing hand? Media capture and government accountability." *American Economic Review* 96(3):720–736.

Besley, Timothy and Robin Burgess. 2002. "The political economy of government responsiveness: Theory and evidence from India." *Quarterly Journal of Economics* 117(4):1415–1451.

Bhandari, Abhit, Horacio Larreguy and John Marshall. 2019. "Able and mostly willing: An empirical anatomy of information's effect on voter-driven accountability in Senegal." Working paper.

Bidwell, Kelly, Katherine Casey and Rachel Glennerster. 2019. "Debates: Voting and Expenditure Responses to Political Communication." Working Paper.

Boas, Taylor C., F. Daniel Hidalgo and Marcus A. Melo. 2019. "Norms versus Action: Why Voters Fail to Sanction Malfeasance in Brazil." *American Journal of Political Science* 63(2):385–400.

Bond, Robert M., Christopher J. Fariss, Jason J. Jones, Adam D.I. Kramer, Cameron Marlow, Jaime E. Settle and James H. Fowler. 2012. "A 61-million-person experiment in social influence and political mobilization." *Nature* 489:295–298.

Bowles, Jeremy and Horacio Larreguy. 2018. "Who Debates,Wins? Evidence on Debate Participation and Radio Broadcasting in a Liberian Election." Working paper.

Buntaine, Mark T., Ryan Jablonski, Daniel L. Nielson and Paula M. Pickering. 2018. "SMS texts on corruption help Ugandan voters hold elected councillors accountable at the polls." *Proceedings of the National Academy of Sciences* 115(26):6668–6673.

Bursztyn, Leonardo, Davide Cantoni, David Y. Yang, Noam Yuchtman and Y. Jane Zhang. 2019. "Persistent Political Engagement: Social Interactions and the Dynamics of Protest Movements." Working Paper.

Campante, Filipe, Ruben Durante and Francesco Sobbrio. 2017. "Politics 2.0: The multifaceted effect of broadband internet on political participation." *Journal of the European Economic Association* 16(4):1094–1136.

Chen, Yuyu and David Y. Yang. 2019. "The Impact of Media Censorship: 1984 or Brave New World?" *American Economic Review* 109(6):2294–2332.

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.

Costas, Elena, Albert Solé-Ollé and Pilar Sorribas-Navarro. 2011. "Do voters really tolerate corruption? Evidence from Spanish Mayors." Working paper.

Cruz, Cesi, Philip Keefer and Julien Labonne. 2019. "Buying Informed Voters: New Effects of Information on Voters and Candidates." Working paper.

Cruz, Cesi, Philip Keefer, Julien Labonne and Francesco Trebbi. 2019. "Making Policies Matter: Voter Responses to Campaign Promises." Working paper.

Da Silveira, Bernardo S. and Joao M.P. De Mello. 2011. "Campaign advertising and election outcomes: Quasi-natural experiment evidence from gubernatorial elections in Brazil." *Review of Economic Studies* 78(2):590–612.

de Figueiredo, Miguel F.P., F. Daniel Hidalgo and Yuri Kasahara. 2014. "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.

DellaVigna, Stefano, John A. List, Ulrike Malmendier and Gautam Rao. 2014. "Voting to tell others." NBER working paper No. 19832.

DellaVigna, Stefano and Matthew Gentzkow. 2010. "Persuasion: Empirical Evidence." *Annual Review of Economics* 2(1):643–669.

DellaVigna, Stefano, Ruben Enikolopov, Vera Mironova, Maria Petrova and Ekaterina Zhuravskaya. 2014. “Cross-border media and nationalism: Evidence from Serbian radio in Croatia.” *American Economic Journal: Applied Economics* 6(3):103–32.

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

Enikolopov, Ruben, Maria Petrova and Aleksey Makarin. 2018. “Social Media and Protest Participation: Evidence from Russia.” Working Paper.

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

Fafchamps, Marcel, Ana Vaz and Pedro C. Vicente. Forthcoming. “Voting and Peer Effects: Experimental Evidence from Mozambique.” *Economic Development and Cultural Change* .

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.

Fergusson, Leopoldo. 2014. “Media markets, special interests, and voters.” *Journal of Public Economics* 109(C):13–26.

Fergusson, Leopoldo and Carlos Molina. 2019. “Facebook Causes Protests.” Working Paper.

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.

García-Jimeno, Camilo, Angel Iglesias and Pinar Yildirim. 2018. “Women, Rails and Telegraphs: An Empirical Study of Information Diffusion and Collective Action.” NBER Working Paper 24495.

Gentzkow, Matthew, Jesse M. Shapiro and Michael Sinkinson. 2011. “The Effect of Newspaper Entry and Exit on Electoral Politics.” *American Economic Review* 101(7):2980–3018.

George, Siddharth, Sarika Gupta and Yusuf Neggers. 2018. “Coordinating Voters against Criminal Politicians: Evidence from a Mobile Experiment in India.” Working paper.

Gerber, Alan S., James G. Gimpel, Donald P. Green and Daron R. Shaw. 2011. “How large and long-lasting are the persuasive effects of televised campaign ads? Results from a randomized field experiment.” *American Political Science Review* 105(1):135–150.

Hootsuite and We are Social. 2018*a*. “2018 Digital Yearbook.”.
URL: *https://hootsuite.com/pages/digital-in-2018*

Hootsuite and We are Social. 2018*b*. “Digital in 2018 in Mexico.”.
URL: *https://hootsuite.com/pages/digital-in-2018*

Howard, Philip N. 2010. *The digital origins of dictatorship and democracy: Information technology and political Islam*. Oxford University Press.

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

Jones, Jason J., Robert M. Bond, Eytan Bakshy, Dean Eckles and James H. Fowler. 2017. “Social influence and political mobilization: Further evidence from a randomized experiment in the 2012 US presidential election.” *PloS one* 12(4):e0173851.

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.

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, Christopher Lucas and John Marshall. 2016. “When do media stations support political accountability? A field experiment in Mexico.” AEA RCT Registry AEARCTR-0001186.

Larreguy, Horacio and Joana C. M. Monteiro. 2019. “Media Networks and Political Accountability: Evidence from Radio Networks in Brazil.” Working paper.

Larreguy, Horacio, John Marshall and Jr. Snyder, James M. 2018. “Leveling the playing field: How equalizing access to campaign advertising helps locally non-dominant parties in consolidating democracies.” *Journal of the European Economic Association* 16(6):1812–1849.

Larreguy, Horacio, John Marshall and Jr. Snyder, James M. 2019. “Publicizing Malfeasance: When the local media structure faciliates electoral accountability in Mexico.” Working paper.

Larson, Jennifer M. 2017. “Networks and Interethnic Cooperation.” *Journal of Politics* 79(2):546–559.

Liberini, Federica, Michela Redoanoz, Antonio Russo, Angel Cuevas and Ruben Cuevas. 2018. “Politics in the Facebook Era: Evidence from the 2016 US Presidential Elections.” Working paper.

Lierl, Malte and Marcus Holmlund. forthcoming. Performance Information and Voting Behavior in Burkina Faso’s Municipal Elections: Separating the Effects of Information Content and Information Delivery. In *Information and Accountability: A New Method for Cumulative Learning*, ed. Thad Dunning, Guy Grossman, Macartan Humphreys, Susan Hyde and Craig McIntosh. Cambridge University Press.

Lynch, Marc. 2011. “After Egypt: The limits and promise of online challenges to the authoritarian Arab state.” *Perspectives on Politics* 9(2):301–310.

Manacorda, Marco and Andrea Tesei. 2018. “Liberation technology: mobile phones and political mobilization in Africa.” Working paper.

Marshall, John. 2019*a*. “Signaling sophistication: How social expectations can increase political information acquisition.” *Journal of Politics* 81(1):167–186.

Marshall, John. 2019*b*. “Tuning in, voting out: News consumption cycles, homicides, and electoral accountability in Mexico.” Working paper.

Martin, Gregory J. and Ali Yurukoglu. 2017. “Bias in cable news: Persuasion and polarization.” *American Economic Review* 107(9):2565–2599.

McCann, James A. and Jorge I. Domínguez. 1998. “Mexicans react to electoral fraud and political corruption: an assessment of public opinion and voting behavior.” *Electoral Studies* 17(4):483–503.

Miner, Luke. 2015. “The unintended consequences of Internet diffusion: Evidence from Malaysia.” *Journal of Public Economics* 132:66–78.

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

Myerson, Roger B. 1999. “Theoretical comparisons of electoral systems.” *European Economic Review* 43(4-6):671–697.

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

Panagopoulos, Costas and Donald P. Green. 2008. “Field experiments testing the impact of radio advertisements on electoral competition.” *American Journal of Political Science* 52(1):156–168.

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

Pierskalla, Jan H. and Florian M. Hollenbach. 2013. “Technology and collective action: The effect of cell phone coverage on political violence in Africa.” *American Political Science Review* 107(2):207–224.

Platas Izama, Melina and Pia Raffler. 2018. “The Limits of Partisanship: How Information Can Encourage Crossing Party Lines.” Working paper.

Poushter, Jacob, Caldwell Bishop and Hanyu Chwe. 2018. “Social media use continues to rise in developing countries but plateaus across developed ones.” *Pew Research Center* 22.

Prat, Andrea and David Strömberg. 2013. The Political Economy of Mass Media. In *Advances in Economics and Econometrics: Volume 2, Applied Economics: Tenth World Congress*. Vol. 50 Cambridge University Press p. 135.

Reinikka, Ritva and Jakob Svensson. 2011. “The power of information in public services: Evidence from education in Uganda.” *Journal of Public Economics* 95(7-8):956–966.

Rodríguez-Castelán, Carlos, Kiyomi Cadena and Laura Moreno. 2018. “Evaluación de largo plazo de los impactos distributivos del Fondo de Infraestructura Social Municipal (FISM).” Working paper.

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

Simpser, Alberto. 2012. “Does electoral manipulation discourage voter turnout? Evidence from Mexico.” *Journal of Politics* 74(3):782–795.

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

Spenkuch, Jorg L. and David Toniatti. 2018. "Political Advertising and Election Outcomes." *Quarterly Journal of Economics* 133(4):1981–2036.

Steinert-Threlkeld, Zachary C. 2017. "Spontaneous collective action: Peripheral mobilization during the Arab Spring." *American Political Science Review* 111(2):379–403.

Strömberg, David. 2015. "Media and Politics." *Annual Review of Economics* 7:173–205.

Vasudevan, Srinivasan. 2019. "Diminishing the Effectiveness of Vote Buying: Experimental Evidence from a Persuasive Radio Campaign in India." Working paper.

Zhuravskaya, Ekaterina, Maria Petrova and Ruben Enikolopov. forthcoming. "Political Effects of the Internet and Social Media." *Annual Review of Economics* .

A Appendix

Contents

A.1 Mixed evidence that information has improved electoral accountability

A substantial body of evidence evaluating the impact of information on electoral accountability in developing contexts—where political accountability remains limited (e.g. Khemani et al. 2016; Pande 2011)—has accumulated over the past decade.[39] Table A1 summarizes the results of studies that leverage either field or natural experiments to estimate the effects of providing information documenting at least one aspect of incumbent performance to voters before elections on administrative or self-reported measures of votes for the incumbent or the incumbent party.[40] We distinguish studies in terms of: (i) whether the findings are broadly consistent with standard theories of electoral accountability; and (ii) whether treatment saturation—the fraction of the electorate that had direct access to information corresponding to a given incumbent politician or party—was low (targeting less than 10% of the electorate represented by the incumbent about which information was provided), medium (10-40% of the electorate), or high (greater than 40% of the electorate). Appendix A.2 discusses the classification of particular studies in greater detail.

As Table A1 shows, the information campaigns most likely to support electoral accountability involve higher levels of saturation. Most of the studies examining high saturation campaigns leverage spatial variation in access to media outlets that are likely to report on local incumbents' performance. Ferraz and Finan's (2008) finding that the sanctioning of incumbents based on the outcomes of independent audit reports in Brazil is driven by municipalities with access to a local radio station, while Larreguy, Marshall and Snyder (2019) further exploit plausibly exogenous variation in access to local media to demonstrate that voter rewards for and punishment of malfeasance revelations from a similar audit program in Mexico is caused by the media. Marshall (2019*b*) also finds that access to local media in Mexico drives sanctioning of municipal incumbents overseeing spikes in homicide rates before elections. Although we focus on voter responses to information provided before elections in this article, other studies show that politician behavior in office responds to the presence of the media (e.g. Avis, Ferraz and Finan 2018; Besley and Burgess 2002; Larreguy and Monteiro 2019; Reinikka and Svensson 2011).

Although they also do not vary saturation levels, a number of medium and high saturation field experiments also provide compelling evidence consistent with electoral accountability. Randomizing the mass provision of incumbent performance information shortly before elections, these studies show—across a variety of continents—that voters reward incumbents at the ballot box for

[39] For studies from developed contexts, see, for example, Berry and Howell (2007), Costas, Solé-Ollé and Sorribas-Navarro (2011), Fergusson (2014), Kendall, Nannicini and Trebbi (2015), and Snyder and Strömberg (2010).

[40] We exclude studies exclusively measuring vote intentions, on the basis that they are normally elicited shortly after surveys, may be vulnerable to social desirability biases, and are hypothetical by construction. We also exclude recent articles examining debates between candidates (e.g. Bowles and Larreguy 2018; Bidwell, Casey and Glennerster 2019; Platas Izama and Raffler 2018), on the basis that a wide range of information—which may or may not include incumbent performance indicators—is provided.

Table A1: The electoral effects of providing incumbent performance information, by treatment saturation

		Consistency of findings with standard theories of electoral accountability		
		Mostly null findings	*Mixed findings*	*Mostly consistent findings*
Saturation of information dissemination	*Low (0-10%)*	Boas, Hidalgo and Melo (2019), Humphreys and Weinstein (2012), Lierl and Holmlund (forthcoming).	Adida et al. (forthcoming).	Bhandari, Larreguy and Marshall (2019).
	Medium (10-40%)		Buntaine et al. (2018), Chong et al. (2015), de Figueiredo, Hidalgo and Kasahara (2014).	Arias, Larreguy, Marshall and Querubín (2019), Banerjee et al. (2011), George, Gupta and Neggers (2018).
	High (40-100%)	Cruz, Keefer and Labonne (2019).		Cruz et al. (2019), Ferraz and Finan (2008), Larreguy, Marshall and Snyder (2019), Marshall (2019*b*).

Notes: Underlined articles leverage experimental variation in access to information. See section A.2 for more details on the classification of articles and findings.

fulfilling campaign promises (Cruz et al. 2019), exerting greater legislative effort (Banerjee et al. 2011), and engaging in less malfeasance than expected (Arias, Larreguy, Marshall and Querubín 2019), and also sanction politicians accused of severe crimes like murder (George, Gupta and Neggers 2018). Several medium saturation studies also find some evidence that voters support (reject) better(worse)-performing incumbents, but observe different degrees of electoral accountability across different layers of government (Buntaine et al. 2018) or find that voters primarily sanction challengers (Chong et al. 2015; de Figueiredo, Hidalgo and Kasahara 2014). While one high saturation campaign informing Filipino voters of the municipal development funds available to often-corrupt incumbents had limited impact on voting behavior, Cruz, Keefer and Labonne (2019) also report an increase in vote buying in treated areas commensurate with treated voters' initial intentions to abandon incumbents immediately after receiving the information.

In contrast, many low saturation interventions fail to detect effects of incumbent performance information on voting behavior. Indeed, randomized information provision pertaining to policy decisions and legislative effort did little to influence voters in northern Brazil (Boas, Hidalgo and Melo 2019), Burkina Faso (Lierl and Holmlund forthcoming), or Uganda (Humphreys and Weinstein 2012). All but the last of these studies come from the recent Metaketa initiative, which coordinated similar accountability experiments across six developing countries and found negligible effects on average (Dunning et al. forthcoming).[41] Adida et al.'s (forthcoming) study is an exception with mixed findings showing that legislator performance that exceeded expectations was punished in a very low saturation version of the intervention and, when combined with civics train-

[41]The Metaketa studies providing incumbent performance information are: Adida et al. (forthcoming), Arias, Larreguy, Marshall and Querubín (2019), Boas, Hidalgo and Melo (2019), Buntaine et al. (2018), and Lierl and Holmlund (forthcoming); one intervention in India was withdrawn, while another focused on debates rather than performance.

ing, was rewarded in a somewhat higher saturation version of the same intervention. The single low saturation field experiment to report that incumbent performance consistently impacted vote choices comes from rural Senegal (Bhandari, Larreguy and Marshall 2019), where there is evidence of substantial information diffusion within treated villages.

In sum, our review of extant studies indicates that saturation may moderate the impact of information dissemination on electoral accountability. However, the notable correlation in Table A1 is, at best, suggestive of a causal relationship. First, saturation is likely correlated with many potential confounds, including the type of content provided, how prominently and persuasively information is communicated, or the predisposition of voters in a given context to respond. Second, studies examining high saturation information provision have generally leveraged observational data from natural experiments that may only be published when significant results are found. Since such designs may be under-powered to detect small effects or more vulnerable than pre-registered experiments to remain in the "file drawer," the observed cross-study correlation could instead reflect publication biases. Third, the one study that has experimentally varied electorate-level saturation introduced only limited variation in saturation and no experimental condition involved high levels of saturation.[42]

A.2 Classification of extant informational interventions

As noted above, we reviewed field and natural experimental studies estimating the effect of providing voters with incumbent performance information in developing contexts. We thus excluded studies in developed contexts and studies in developing contexts that provided non-performance information (e.g. candidate debates). We also excluded that studies that only contained self-reported outcomes measured immediately after treatment. As noted above, saturation is defined by the share of voters with *direct access* to incumbent performance information. We, therefore, do not count untreated or indirectly treated voters within treated clusters in our computations of saturation.

Below we summarize the studies included in Table A1, and discuss our coding decisions:

1. In Adida et al. (forthcoming), we code the borderline statistically significant negative result for low dosage in column (2) of Table 4 that reports the impact of positive information on incumbent vote share as "mostly null findings" because the direction goes against the expectations of standard electoral accountability models. The borderline statistically significantly positive effect of providing better-than-expected incumbent performance information for high dosage in column (4) is coded as "mixed findings," given that this effect only holds when incumbent performance information is accompanied by civic training. A typical com-

[42] Adida et al. (forthcoming) compared 15 low dosage communes with 15 higher dosage communes, which reached 12-15% voters within around 4% and 30% of villages within each commune respectively.

mune contains around 50 villages. We code their low dosage case, where 2 villages per commune (or around 4%) were treated, as "low" saturation. We also code their high dosage case, where 15 villages per commune (or 30%) were treated, as "low" saturation because only 12-15% of households within each treated village had direct access to treatment. Overall, we thus code this study as providing "mixed findings."

2. In Arias, Larreguy, Marshall and Querubín (2019), we code the significant positive effect of the information treatment for 0% malfeasance and the significantly negative interaction between treatment and incumbent malfeasant spending on incumbent party vote share, both from column (4) of Table 4, as "mostly consistent findings." The campaign's saturation is classified as "medium" because around 20% of precincts were treated in a typical municipality, although only up to 200 leaflets were distributed to households (more than half of the households) within a given electoral precinct.

3. In Banerjee et al. (2011), we code the significant positive effects of the interaction between treatment and overall incumbent quality in column (4) and incumbent performance in column (5) of Table 4 on incumbent party vote share "mostly consistent findings." We code saturation as "medium" due to the fact that 200 polling stations—each with roughly 1,000 voters—are treated out of 775 selected polling stations in ten constituencies with high slum density—each with approximately 100,000 citizens.

4. In Bhandari, Larreguy and Marshall (2019), we code the significant positive interaction effect between the information treatment and the incumbent local performance index on validated reported vote for the incumbent in Table 8 and polling station-level incumbent party vote share in Table 9 as "mostly consistent findings." Saturation is coded as "low:" although 375 villages were treated across 5 constituencies in Senegal, each containing approximately 300 villages, only 9 voters per village were directly treated.

5. In Boas, Hidalgo and Melo (2019), we code the zero treatment effects of reporting either approved or rejected account through a field experiment reported in Figure 3 as "mostly null findings." Saturation is coded as "low" because the information was randomly distributed to 1,600 registered voters across 47 municipalities in Brazil.

6. In Buntaine et al. (2018), we code the zero effects of either good or bad news about incumbent performance on chairman vote and the statistically significantly positive (negative) effects of good (bad) news about incumbent performance on councilor vote as "mixed findings." Saturation is coded as "medium" because messages were sent to 16,083 citizens in 762 villages and we expect these to cover a medium share of the villages under the councilors in the experimental sample. This study also varied saturation across villages, as opposed to across

electorates, but found little evidence to suggest an impact of localized saturation (see Figure 2).

7. In Chong et al. (2015), we code the effects of the treatments as "mixed." Table 4 shows that the provision of incumbent corruption information reduced turnout for the incumbent and challenger where high levels of incumbent corruption (top tercile) were reported. However, such effects are greater for the challenger than the incumbent, which suggest that overall corruption revelations favored corrupt incumbents. Saturation is coded as "medium" because, although all electoral precincts in each of the 12 sample municipalities were assigned to one of four treatment conditions (including a pure control), the authors ended up pooling three of these conditions as a control group, leaving a quarter of precincts per municipality treated by their definition.

8. In Cruz, Keefer and Labonne (2019), the results are clearly "most null findings," as they find no effect of the information treatment on reported vote for the incumbent in panel A of Table A.22 and vote share for the incumbent in Panel B. Saturation is saturation coded as "high" because 142 villages were treated out of 284 in the sample in 12 municipalities, each with 20-25 village approximately; this implies that 47-60% of villages were treated in each municipality. In each village, flyers containing treatment information were distributed to all households.

9. In Cruz et al. (2019), we code the results as "mostly consistent findings" on the basis that the authors report significantly positive effects of an interaction between the information treatment and the incumbent keeping their policy promises ahead of entering office on reported vote for the incumbent in Table 4. Saturation is coded as "high" because 104 villages were treated (although not all treated villages received prior incumbent promises) out of 158 in the sample in 7 municipalities, each with 20-25 village approximately; this implies that 59-74% of villages were treated in each municipality. In each village, flyers containing treatment information were distributed to all households.

10. In de Figueiredo, Hidalgo and Kasahara (2014), we code the findings as "mixed" based on the significant negative effect of information showing the challenger to be corrupt on vote share for the challenger in Table 2 and the insignificant positive effect of information also showing the incumbent to be corrupt on vote share for the incumbent in Table 3. Saturation is coded as "medium" on the basis that 187,177 fliers with candidate information were delivered to 200 out of 1,759 precincts in Sao Paulo.

11. In George, Gupta and Neggers (2018), we code the significant negative effects of revelations of candidate murder-related charges on candidate vote share in Table 6 as "mostly consistent

findings." We code saturation as "medium" because the authors treated 80% of an experimental sample of 4,131 villages, out of a total of 9,627 villages from 39 assembly constituencies.

12. In Humphreys and Weinstein (2012), we code the null treatment effect on reported vote of reporting relatively good news about incumbent performance relative to the prior beliefs in Table 3 as "mostly null findings." Saturation is coded as "low" since, despite the many information dissemination efforts, treatment information reached a very small share of the electorate. Specifically, the authors undertook two main information dissemination efforts prior to the election: community-wide workshops and blanket treatment of polling stations with scorecard results. However, while workshops averaged about 120 people in attendance and in each workshop and 1500 copies of the local language scorecard were handed out to be shared more broadly, only one workshop was conducted per constituency, each averaging roughly 50,000 voters. Similarly, only voters in two polling station areas per constituency were assigned to receive scorecard results.

13. In Lierl and Holmlund (forthcoming), we code the null treatment effects of both good and bad information about incumbent performance reported in Table 3 as "mostly null effects." Saturation is coded as "low" because the information was only randomly distributed to 752 study participants across 38 municipalities in Burkina Faso.

14. Finally, Ferraz and Finan (2008), Larreguy, Marshall and Snyder (2019), and Marshall (2019*b*) are all articles that analyze the effect of media revelation of incumbent malfeasance or municipal violence on incumbent vote share, finding large sanctioning or reward effects that reflect performance (of some form) in each case; we thus code each as showing "mostly consistent findings." Due to the large coverage or circulation of the media outlets that they focus on, the saturation in each study is coded as "high."

A.3 Infographics available on the Facebook pages

Figures A1a or A1b show examples of the infographic available on the Facebook page associated with the ads, with the former reporting 0% irregularities and the latter type reporting greater than 0% irregularities.

A.4 Additional information about WhatsApp treatments

Figure A2 shows the message received before the election by treated WhatsApp survey respondents. Treatments were sent via WhatsApp in the days before the election (after the baseline survey). All respondents in 20 of the 43 control municipalities received no WhatsApp treatments,

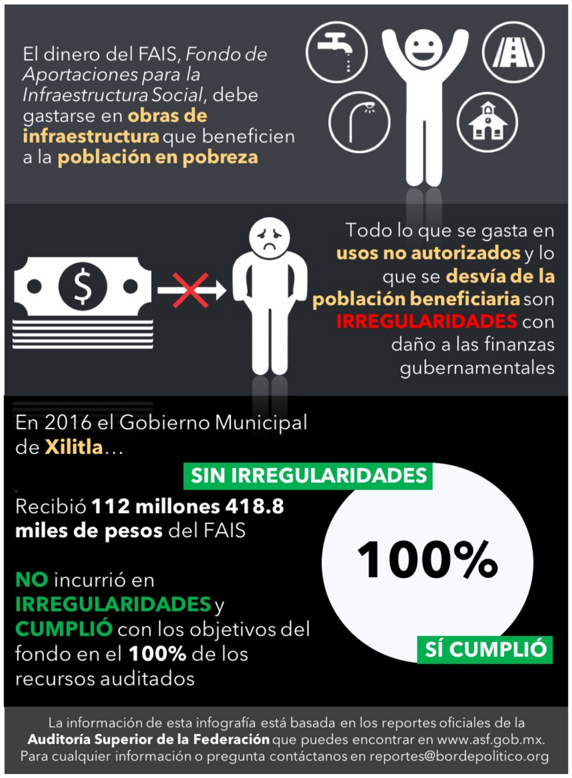

(a) Example of an infographic from a municipality (Xilitla, San Luis Potosí) where irregularities were 0%

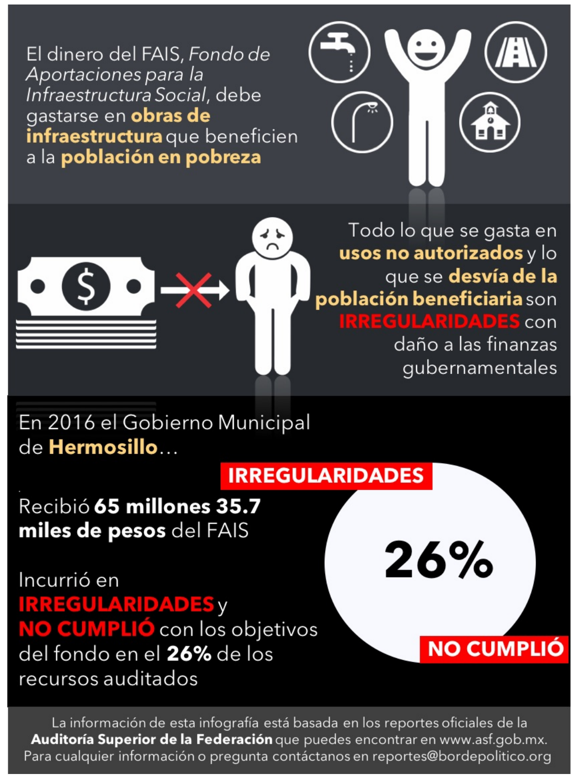

(b) Example of an infographic from a municipality (Hermosillo, Sonora) where irregularities were above 0%

Figure A1: Examples of treatment infographics

Note: In English, the first panel of the left infographic says "The funds from the FAIS, the Fund for Transfers for Municipal Social Infrastructure, must be spent on infrastructure projects benefiting the poor;" the second panel says "All funds spent on unauthorized projects or not benefiting the poor constitute irregularities harming public finances;" the third panel says "In 2016, the municipal government of Xilitla received 112 million 419.8 thousand pesos" and "It did not incur in any regularity fulfilling the intended purposes of the fund in 100% of the audited funds;" and the fourth panel says "The information from the infographic is from the ASF's official audit reports that can be accessed at asf.gov.mx" and "To request more information or make an inquiry, you can contact us at reportes@bordepolitico.org." The first, second, and last panel of the right infographic do not change, and the third panel says "In 2016, the municipal government of Hermosillo received 65 million 35.7 thousand pesos" and "It incurred in regularities not fulfilling the intended purposes of the fund in 26% of the audited funds."

while 80% of respondents in low- and high-saturation treated municipalities and the remaining 23 control municipalities received the WhatsApp treatment. To open the WhatsApp video, respondents needed to click to download the video. Within low- and high-saturation municipalities, these treatments were equally split between common knowledge and no common knowledge information, while all treated respondents in the 23 zero-saturation municipalities that received WhatsApp treatments did not receive the common knowledge treatment (since treatment had not been distributed at scale). These treatments were assigned within blocks of five similar respondents (based on baseline survey data) within municipalities.

A.5 Engagement with Facebook ads

Figure A3 plots trends in Facebook user engagement with Borde Político's ad campaign over the pre-election period. We were not able to obtain day-by-day data on watching at least 10 seconds of the ad.

Figure A4 further reports the distribution of aggregate engagement with the Facebook ads within low and high saturation municipalities.

A.6 Deviation from the pre-analysis plan

All analyses follow our pre-analysis plan (available at socialscienceregistry.org/trials/3135), with the following exception:

1. The one block of municipalities containing only two municipalities was excluded from the main analyses. We excluded this "rump" block (which contained two small municipalities comprising 23 precincts) on the basis that the treatment assignment probabilities vary from the other 42 blocks and that no electoral data was available for one segment within one of these municipalities. In the presence of block fixed effects, the former issue could be addressed to yield unbiased estimates of ATEs and CATEs by interacting all treatments (and covariate interactions) with a fixed effect for this block. For simplicity, we choose not to do this for our main estimates; accordingly, no municipal-level weights are required. However, Table A2 reports similar results for our main estimates when these two municipalities are included in regressions that further use inverse weights to adjust for the different probability of treatment assignment in the 2-municipality block.

A.7 Balance tables

Table A3 shows balance tests based on the municipal-level treatment assignments in equation (4), where for each predetermined variable we report the p value associated with an F test of the re-

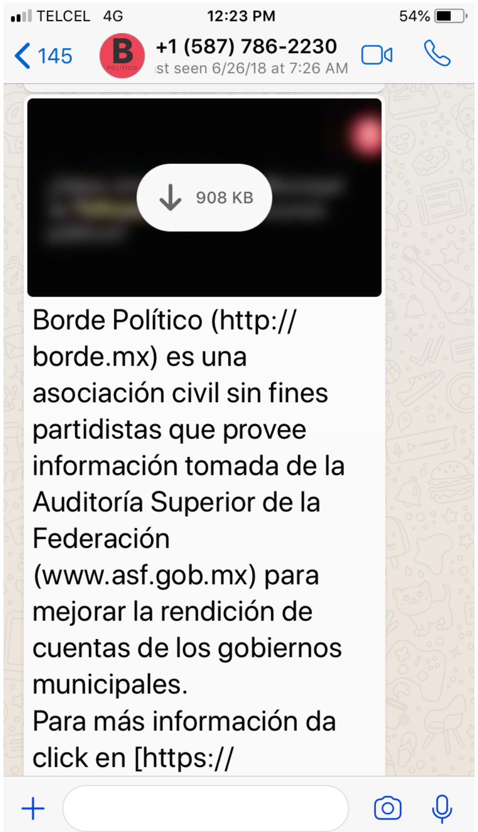

(a) WhatsApp message containing ad (top part of the message)

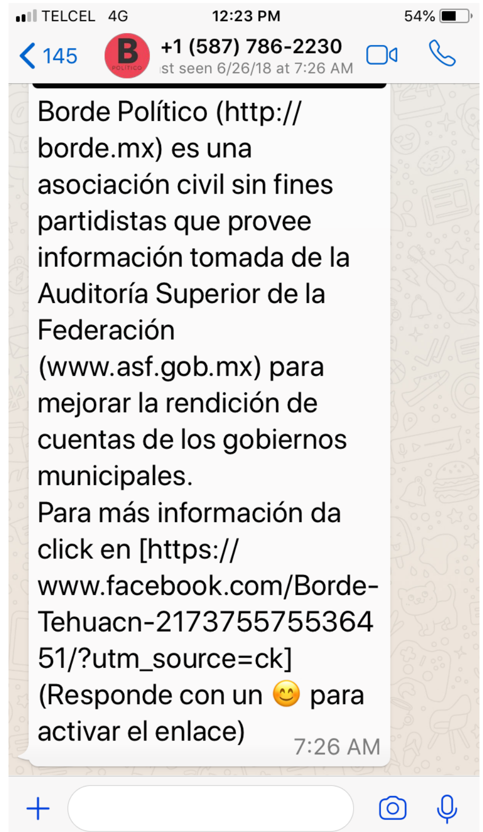

(b) WhatsApp message containing ad (bottom part of the message)

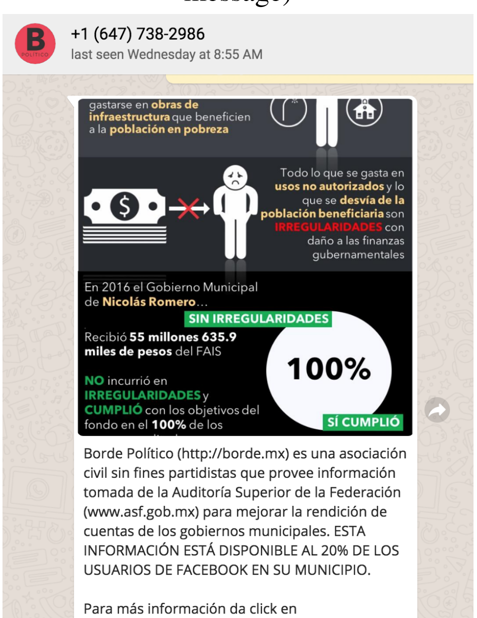

(c) Reminder WhatsApp message containing infographic (top)

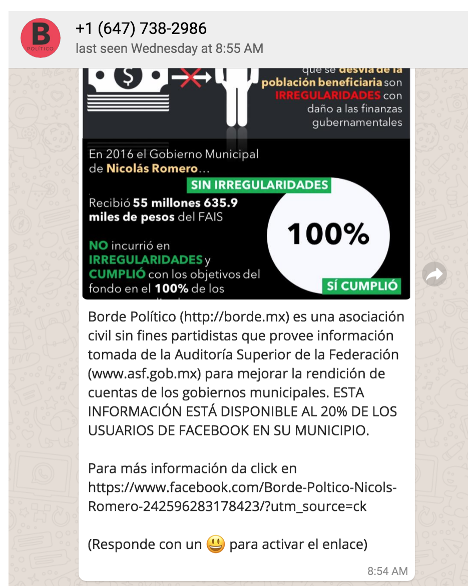

(d) Reminder WhatsApp message containing infographic (bottom)

Figure A2: Example of the slides included in the ad video (from Nicolás Romero, Estado de México)

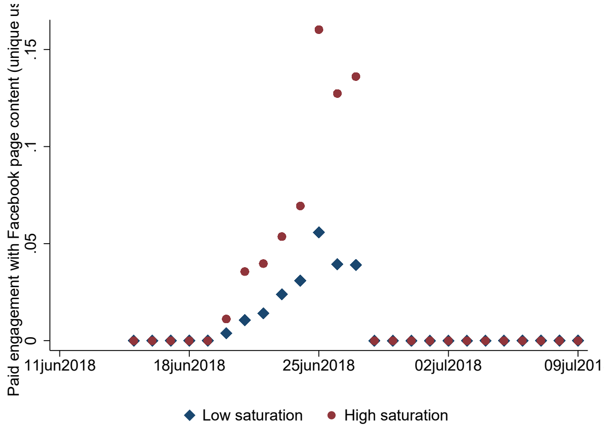

(a) Paid-for impressions

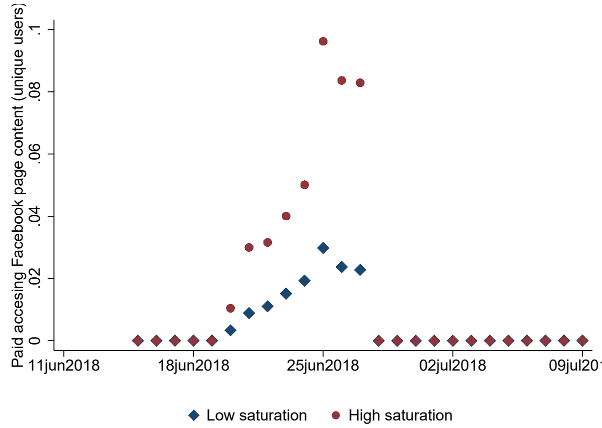

(b) Paid-for unique viewers

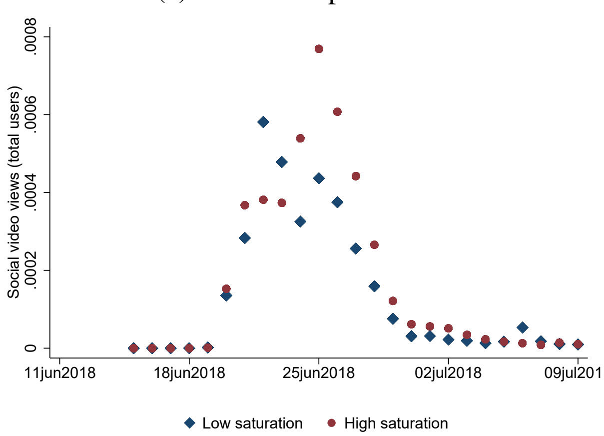

(c) Organic impressions

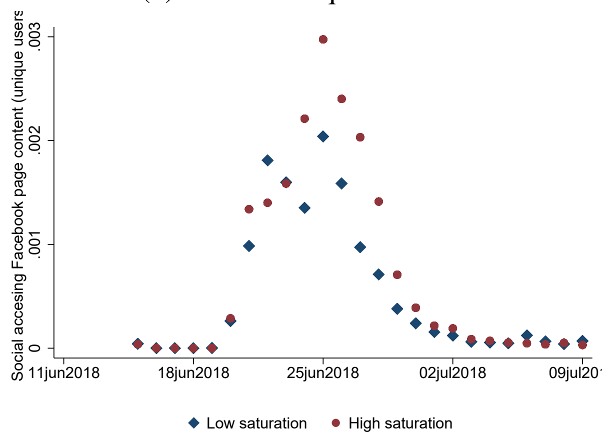

(d) Organic unique viewers

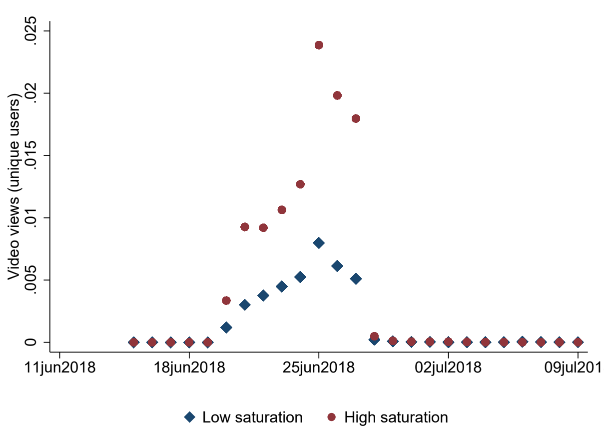

(e) Unique views (of at least 3 seconds)

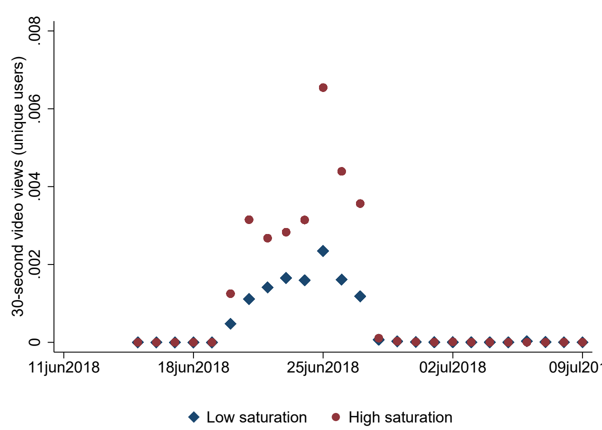

(f) Unique views (of entire video)

Figure A3: Trends in engagement with the Facebook ads during the information campaign

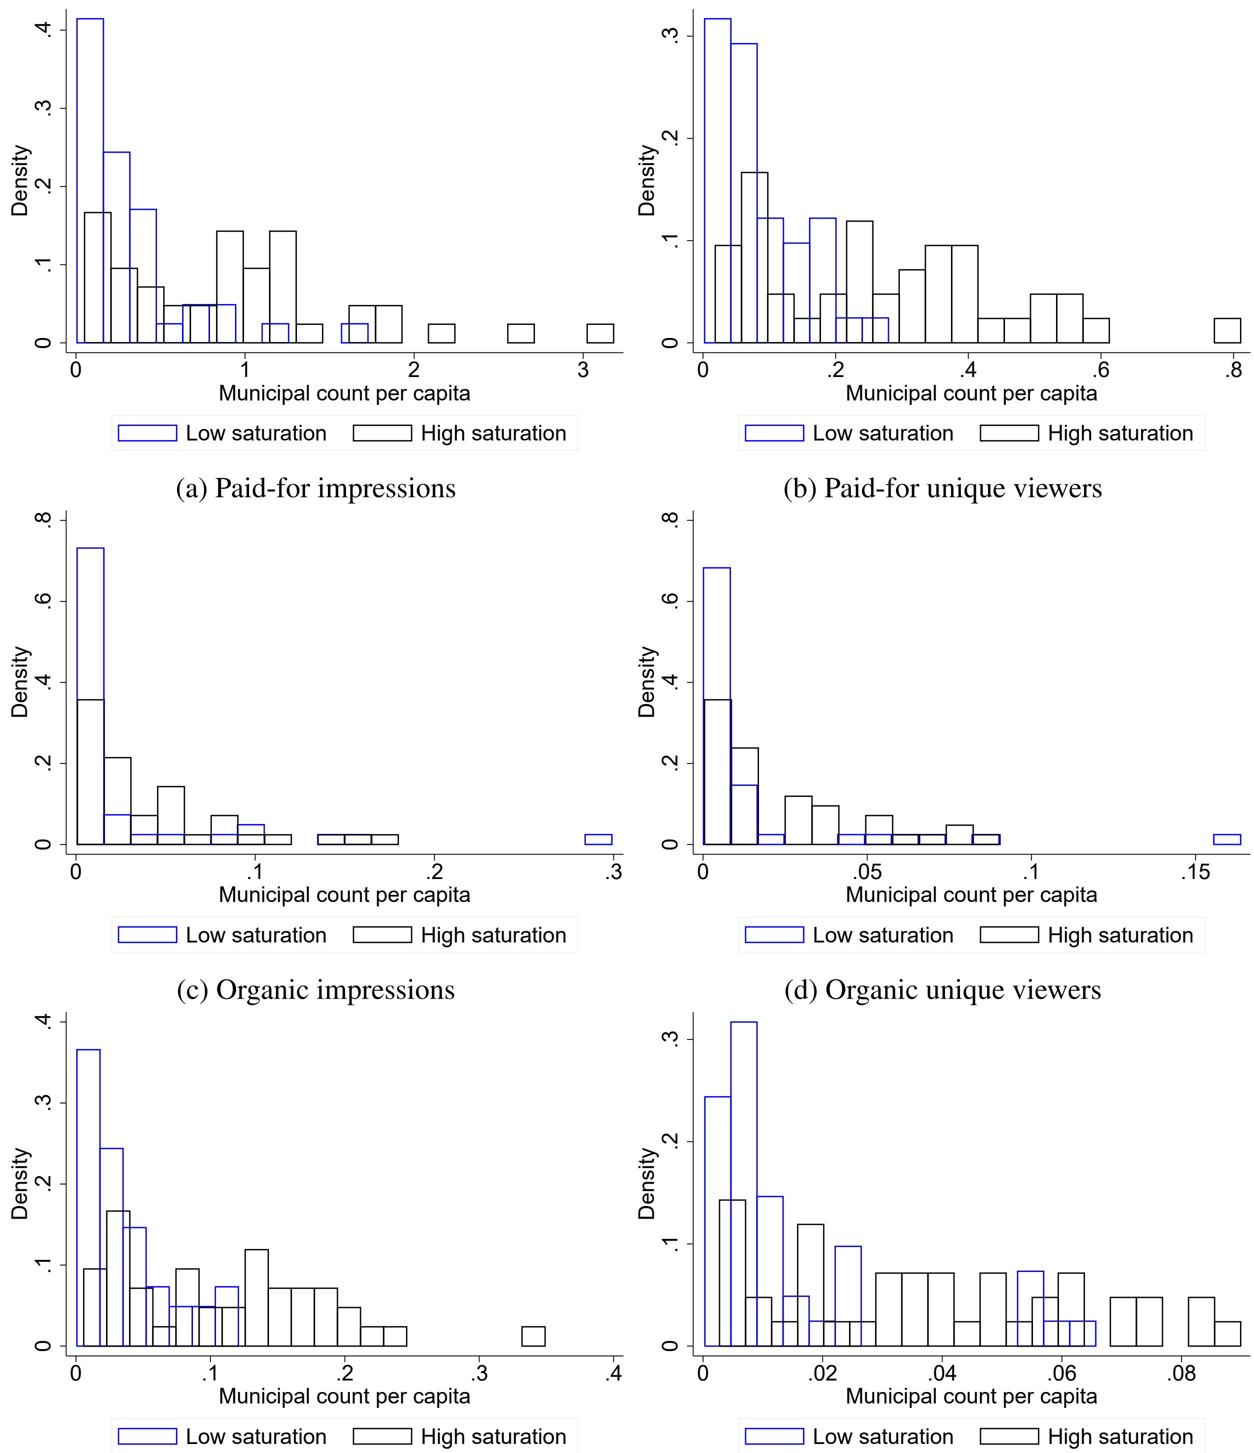

(a) Paid-for impressions

(b) Paid-for unique viewers

(c) Organic impressions

(d) Organic unique viewers

(e) Unique views (of at least 3 seconds)

(f) Unique views (of entire video)

Figure A4: Distribution of engagement with the Facebook ads during the information campaign, by municipal saturation

Table A2: Effect of Facebook ads on precinct-level municipal incumbent party vote share (share of registered voters), by quartile of the sample irregularities distribution and including the block containing only two municipalities

	Incumbent party vote (share of registered voters)			
	(1)	(2)	(3)	(4)
Facebook ads	0.040**	0.058***		
	(0.016)	(0.021)		
Facebook ads × Q3	-0.049	-0.031		
	(0.032)	(0.035)		
Facebook ads × Q4	-0.049*	-0.087**		
	(0.028)	(0.040)		
Spillover	0.023	0.035*		
	(0.015)	(0.018)		
Spillover × Q3	-0.005	-0.013		
	(0.026)	(0.031)		
Spillover × Q4	-0.045	-0.078**		
	(0.032)	(0.039)		
Facebook ads in high saturation			0.049***	0.065**
			(0.018)	(0.027)
Facebook ads in high saturation × Q3			-0.076**	-0.055
			(0.038)	(0.042)
Facebook ads in high saturation × Q4			-0.059*	-0.070
			(0.031)	(0.054)
Facebook ads in low saturation			0.024	0.031
			(0.018)	(0.022)
Facebook ads in low saturation × Q3			-0.010	-0.043
			(0.030)	(0.036)
Facebook ads in low saturation × Q4			-0.049	-0.090**
			(0.035)	(0.045)
Spillover in high saturation			0.044**	0.060**
			(0.018)	(0.027)
Spillover in high saturation × Q3			-0.081**	-0.060
			(0.039)	(0.043)
Spillover in high saturation × Q4			-0.055*	-0.065
			(0.030)	(0.055)
Spillover in low saturation			0.015	0.025
			(0.017)	(0.021)
Spillover in low saturation × Q3			0.019	-0.013
			(0.029)	(0.036)
Spillover in low saturation × Q4			-0.039	-0.078*
			(0.036)	(0.044)
Observations	13,274	13,274	13,274	13,274
R^2	0.59	0.64	0.62	0.70
Control outcome mean	0.18	0.18	0.18	0.18
Control outcome std. dev.	0.11	0.11	0.11	0.11
Test: null effect of Facebook ads in Q3 (p value, 2-sided)	0.710	0.297		
Test: null effect of Facebook ads in Q4 (p value, 2-sided)	0.693	0.265		
Test: null effect of spillover in Q3 (p value, 2-sided)	0.362	0.322		
Test: null effect of spillover in Q4 (p value, 2-sided)	0.426	0.122		
Test: larger effect of Facebook ads in high (vs. low) in Q1/Q2 (p value, 1-sided)			0.074	0.027
Test: larger effect of Facebook ads in high (vs. low) in Q3 (p value, 1-sided)			0.086	0.253
Test: larger effect of Facebook ads in high (vs. low) in Q4 (p value, 1-sided)			0.727	0.937
Test: larger effect of spillovers in high (vs. low) in Q1/Q2 (p value, 1-sided)			0.032	0.024
Test: larger effect of spillovers in high (vs. low) in Q3 (p value, 1-sided)			0.051	0.377
Test: larger effect of spillovers in high (vs. low) in Q4 (p value, 1-sided)			0.820	0.971
Interactive covariates		✓		✓

Notes: Each specification is estimated using OLS, and includes a lagged dependent variable and randomization block fixed effects and an interaction between treatment. Specifications including interactive covariates further include interactions between treatment conditions and the following municipal-level covariates: year of audit; amount of FISM funds received; population aged above 18; lagged incumbent party vote share; average years of schooling; share of the population that is illiterate; average occupants per room; average number of children per woman; the share of the population with electricity, water, and drainage in their home; the working age share of the population; and the share of households with internet at home. The omitted irregularities category is Q1/Q2. All segments are weighted equally and by the inverse probability of treatment assignment. Standard errors clustered by municipality are in parentheses.* denotes $p < 0.1$, ** denotes $p < 0.05$, *** denotes $p < 0.01$ from two-sided t tests. The one-sided coefficient tests at the foot of the table test whether the effect of high saturation is larger in a given direction in high saturation than in low saturation municipalities.

Figure A5: Example of an infographic from a municipality (Hermosillo, Sonora) where irregularities were above 0%

striction that $\beta_1 = \beta_2 = 0$. Only 3 of the 60 tests show a statistically significant difference at the $p < 0.1$ level.

Table A4 shows balance tests based on the segment-level treatment assignments in equation (2), where for each predetermined variable we report the p value associated with an F test of the restriction that $\beta_1 = \beta_2 = \gamma_1 = \gamma_2 = 0$. Only 1 of 62 tests show a statistically significant difference at the $p < 0.1$ level. Table A5 reports balance tests for the survey-level samples.

A.8 Alternative preregistered specifications

While we prefer the estimation strategies used in the main article, the pre-analysis plan also specified several additional approaches to estimation in some cases. Focusing on the main results in Table 4, these decisions do not influence our findings. Indeed, the principal results in Table 4 are robust to excluding the lagged outcome (see Table A7), further weighting precincts by the number of registered voters (see Table A8), and weighting municipalities with different numbers of segments equally (see Table A9). The estimates without a lagged dependent variable are, unsurprisingly, less precisely estimated.

A.9 The common knowledge variant of the Facebook ads

As noted in the main text, some segments received Facebook ads with additional information designed to facilitate common knowledge about the campaign's scale. Figure A5 provides an example of this in the case of the 20% saturation campaign in the municipality of Hermosillo, Sonora. This slide was the penultimate slide in the video, and thus appeared right before the concluding slide (which contained no text).

As noted in the main text, the common knowledge treatment variant was pooled with the basic Facebook ads because we observed indistinguishable levels of engagement with these ads and

Table A3: Municipal-level Facebook saturation treatment condition balance tests

Covariate	Control mean	Control std. dev.	F test (p value)	Covariate	Control mean	Control std. dev.	F test (p value)
Year of audit	2,015.85	0.35	0.228	Share illiterate in 2010	0.13	0.11	0.440
FISM resources received	101,887.58	130,758.39	0.212	Share with no schooling in 2010	0.12	0.09	0.837
FISM resources used	90,442.80	114,967.96	0.158	Share with at least incomplete primary schooling in 2010	0.88	0.09	0.837
FISM resources audited	76,844.26	89,970.69	0.262	Share with at least complete primary schooling in 2010	0.71	0.16	0.649
Share of resources subject to irregularities	0.09	0.18	0.159	Share with at least incomplete secondary schooling in 2010	0.53	0.20	0.615
Municipal adult population in 2010	135,401.00	199,014.69	0.823	Share with at least complete secondary schooling in 2010	0.47	0.21	0.606
Registered voters in 2015	2,247.71	2,443.80	0.715	Share with higher education in 2010	0.25	0.19	0.502
Municipal incumbent vote share (reg. voters) in 2015	0.26	0.10	0.825	Share disabled in 2010	0.04	0.02	0.409
Municipal turnout in 2015	0.60	0.14	0.890	Share economically active in 2010	0.36	0.07	0.420
Federal MORENA-PT-PES vote share (reg. voters) in 2015	0.05	0.05	0.913	Share without health care in 2010	0.34	0.18	0.634
Federal PAN-PRD-MC vote share (reg. voters) in 2015	0.24	0.12	0.976	Share with state workers healthcare in 2010	0.04	0.05	0.117
Federal PRI-PVEM-NA vote share (reg. voters) in 2015	0.24	0.13	0.378	Average number of occupants per dwelling in 2010	4.27	0.63	0.474
Federal turnout in 2015	0.56	0.14	0.853	Average number of occupants per room in 2010	1.39	0.45	0.100*
Number of households in 2010	812.05	987.63	0.882	Share of households with more than 2 rooms in 2010	0.56	0.16	0.102
Private dwellings in 2010	1,030.64	1,416.93	0.761	Share of households with more than 3 rooms in 2010	0.65	0.19	0.101
Population in 2010	3,412.95	3,785.47	0.929	Share of households with non-dirt floor in 2010	0.89	0.12	0.464
Share women in 2010	0.51	0.02	0.293	Share of households with a toilet in 2010	0.89	0.17	0.032**
Share working age in 2010	0.61	0.06	0.491	Share of households with water in 2010	0.84	0.22	0.327
Share over 65 in 2010	0.06	0.03	0.791	Share of households with drainage in 2010	0.79	0.26	0.904
Share married in 2010	0.56	0.05	0.133	Share of households with electricity in 2010	0.95	0.08	0.418
Average children per woman in 2010	2.65	0.56	0.554	Share of households with water, drainage, and electricity in 2010	0.70	0.30	0.850
Share of households with a male head in 2010	0.78	0.08	0.166	Share of households with a washing machine in 2010	0.47	0.29	0.162
Share born out of state in 2010	0.12	0.15	0.553	Share of households with a landline telephone in 2010	0.29	0.24	0.319
Share Catholic in 2010	0.84	0.16	0.450	Share of households with a radio telephone in 2010	0.70	0.16	0.154
Share non-Catholic in 2010	0.11	0.13	0.147	Share of households with a fridge telephone in 2010	0.65	0.29	0.160
Share without religion in 2010	0.03	0.04	0.636	Share of households with a cell phone in 2010	0.47	0.30	0.585
Share indigenous speakers in 2010	0.21	0.35	0.059*	Share of households with a television in 2010	0.82	0.21	0.132
Average years of schooling in 2010	7.34	2.26	0.727	Share of households with a car in 2010	0.31	0.22	0.138
Average years of schooling for women in 2010	7.14	2.27	0.510	Share of households with a computer in 2010	0.18	0.19	0.558
Average years of schooling for men in 2010	7.55	2.28	0.782	Share of households with internet in 2010	0.12	0.16	0.560

Notes: The F tests report the p value associated with the test of the restriction that $\beta_1 = \beta_2 = 0$ in equation 4. Each specification is estimated using OLS, and includes randomization block fixed effects. All observations are weighted equally by municipality. * denotes $p < 0.1$, ** denotes $p < 0.05$, *** denotes $p < 0.01$ from two-sided t tests.

Table A4: Segment-level Facebook ad treatment condition balance tests

Covariate	Control mean	Control std. dev.	F test (p value)	Covariate	Control mean	Control std. dev.	F test (p value)
Year of audit	2,015.87	0.33	0.396	Average years of schooling for men in 2010	7.90	2.42	0.886
FISM resources received	109,254.87	134,275.11	0.396	Share illiterate in 2010	0.12	0.11	0.185
FISM resources used	95,119.56	120,793.90	0.134	Share with no schooling in 2010	0.11	0.09	0.286
FISM resources audited	81,344.42	89,436.74	0.299	Share with at least incomplete primary schooling in 2010	0.89	0.09	0.286
Share of resources subject to irregularities	0.08	0.16	0.351	Share with at least complete primary schooling in 2010	0.73	0.16	0.448
Municipal adult population in 2010	209,177.31	274,260.97	0.875	Share with at least incomplete secondary schooling in 2010	0.56	0.21	0.767
Segment adult population in 2010	23,131.00	20,766.99	0.572	Share with at least complete secondary schooling in 2010	0.50	0.21	0.749
Precinct adult population in 2010	2,059.45	2,358.88	0.697	Share with higher education in 2010	0.28	0.21	0.849
Registered voters in 2015	2,314.20	2,594.90	0.796	Share disabled in 2010	0.04	0.02	0.639
Municipal incumbent vote share (reg. voters) in 2015	0.43	0.11	0.622	Share economically active in 2010	0.37	0.07	0.172
Municipal turnout in 2015	0.59	0.14	0.601	Share without health care in 2010	0.34	0.17	0.775
Federal MORENA-PT-PES vote share (reg. voters) in 2015	0.05	0.04	0.593	Share with state workers healthcare in 2010	0.04	0.05	0.274
Federal PAN-PRD-MC vote share (reg. voters) in 2015	0.23	0.11	0.935	Average number of occupants per dwelling in 2010	4.24	0.63	0.332
Federal PRI-PVEM-NA vote share (reg. voters) in 2015	0.23	0.12	0.364	Average number of occupants per room in 2010	1.33	0.45	0.101
Federal turnout in 2015	0.55	0.14	0.386	Share of households with more than 2 rooms in 2010	0.59	0.17	0.199
Number of households in 2010	827.99	1,036.54	0.800	Share of households with more than 3 rooms in 2010	0.68	0.20	0.132
Private dwellings in 2010	1,058.14	1,520.19	0.916	Share of households with non-dirt floor in 2010	0.90	0.12	0.426
Population in 2010	3,453.91	3,955.84	0.752	Share of households with a toilet in 2010	0.90	0.16	0.110
Share women in 2010	0.51	0.02	0.502	Share of households with water in 2010	0.85	0.22	0.135
Share working age in 2010	0.61	0.06	0.417	Share of households with drainage in 2010	0.82	0.25	0.193
Share over 65 in 2010	0.06	0.03	0.963	Share of households with electricity in 2010	0.96	0.08	0.234
Share married in 2010	0.55	0.05	0.761	Share of households with water, drainage, and electricity in 2010	0.73	0.30	0.372
Average children per woman in 2010	2.59	0.57	0.664	Share of households with a washing machine in 2010	0.51	0.30	0.120
Share of households with a male head in 2010	0.77	0.08	0.283	Share of households with a landline telephone in 2010	0.33	0.25	0.790
Share born out of state in 2010	0.12	0.15	0.323	Share of households with a radio telephone in 2010	0.72	0.16	0.238
Share Catholic in 2010	0.84	0.16	0.370	Share of households with a fridge telephone in 2010	0.68	0.29	0.133
Share non-Catholic in 2010	0.10	0.12	0.257	Share of households with a cell phone in 2010	0.51	0.30	0.187
Share without religion in 2010	0.03	0.04	0.994	Share of households with a television in 2010	0.84	0.20	0.126
Share indigenous speakers in 2010	0.18	0.33	0.050**	Share of households with a car in 2010	0.34	0.23	0.213
Average years of schooling in 2010	7.68	2.38	0.904	Share of households with a computer in 2010	0.21	0.21	0.898
Average years of schooling for women in 2010	7.47	2.38	0.786	Share of households with internet in 2010	0.15	0.18	0.919

Notes: The F tests report the p value associated with the test of the restriction that $\beta_1 = \beta_2 = \gamma_1 = \gamma_2 = 0$ in equation 1. Each specification is estimated using OLS, and includes randomization block fixed effects. All segments are weighted equally and by the inverse probability of treatment assignment. * denotes $p < 0.1$, ** denotes $p < 0.05$, *** denotes $p < 0.01$ from two-sided t tests.

Table A5: WhatsApp treatment balance tests

	(1) Age	(2) Education	(3) Female	(4) Information acquisition index
WhatsApp ad	-0.264	-0.051	0.034	0.023
	(0.644)	(0.047)	(0.025)	(0.031)
Observations	1,490	1,490	1,490	1,490
R^2	0.071	0.092	0.053	0.082
Control outcome mean	29.65	5.036	0.329	0.153
Control outcome std. dev.	12.31	0.891	0.469	0.581

	(5) Voted in last election	(6) Voted for incumbent	(7) Use of WhatsApp	(8) Use of Facebook
WhatsApp ad	0.021	0.004	-0.003	-0.013**
	(0.024)	(0.018)	(0.002)	(0.005)
Observations	1,490	1,490	1,490	1,490
R^2	0.108	0.163	0.023	0.036
Control outcome mean	0.526	0.163	1	0.997
Control outcome std. dev.	0.461	0.370	0	0.059

	(9) Intend to vote	(10) Intend to vote for incumbent	(11) Others' intention to vote	(12) Others' intention to vote for incumbent
WhatsApp ad	4.365**	3.880**	0.281	2.341*
	(2.132)	(1.649)	(0.869)	(1.344)
Observations	1,490	1,490	1,490	1,490
R^2	0.066	0.130	0.105	0.110
Control outcome mean	77.08	12.43	74.37	45.39
Control outcome std. dev.	42.10	30.55	17.28	26.07

	(13) Priors about irregularities	(14) Prior precision about irregularities	(15) Incumbent seeking reelection	(16) Amount of FISM funds audited
WhatsApp ad	-0.141	0.001	-0.009	849.9
	(1.357)	(0.012)	(0.009)	(872.4)
Observations	1,490	1,490	1,490	1,490
R^2	0.085	0.133	0.605	0.793
Control outcome mean	61.99	0.059	0.406	83,751.2
Control outcome std. dev.	24.94	0.236	0.492	63,621.2

	(17) Adult Population	(18) Average years of schooling	(19) Share illiterate	(20) Average number of occupants per room
WhatsApp ad	9,958.9*	0.059**	-0.003*	-0.007
	(5,546.9)	(0.027)	(0.002)	(0.005)
Observations	1,490	1,490	1,490	1,490
R^2	0.875	0.820	0.705	0.741
Control outcome mean	420,540.2	8.992	0.053	1.095
Control outcome std. dev.	349,921.3	1.161	0.045	0.192

	(21) Average children per women	(22) Share of households with electricity, water and drainage	(23) Share of population that is working age	(24) Share of households with internet
WhatsApp ad	-0.016**	0.007	0.001*	0.005**
	(0.007)	(0.005)	(0.001)	(0.002)
Observations	1,490	1,490	1,490	1,490
R^2	0.713	0.707	0.717	0.780
Control outcome mean	2.217	0.862	0.648	0.249
Control outcome std. dev.	0.257	0.142	0.026	0.102

Notes: Each specification is estimated using OLS. Columns (1) to (14) include municipality fixed effects. Columns (15) to (24) include randomization block fixed effects. All observations are weighted by the inverse probability of treatment assignment. Observations from pure control municipalities are excluded. Robust standard errors are reported in parentheses in columns (1)-(14), while standard errors clustered by municipality are in parentheses in columns (15) to (24). * denotes $p < 0.1$, ** denotes $p < 0.05$, *** denotes $p < 0.01$ from two-sided t tests.

Table A6: Facebook treatment balance tests

	(1) Age	(2) Education	(3) Female	(4) Information acquisition index
High Saturation Municipal treatment	0.985	-0.105	-0.008	0.018
	(0.649)	(0.067)	(0.023)	(0.046)
Low Saturation Municipal treatment	0.595	-0.014	-0.028	0.032
	(0.713)	(0.060)	(0.026)	(0.039)
Observations	2,020	2,020	2,020	2,020
R^2	0.021	0.048	0.030	0.042
Control outcome mean	29.57	5.077	0.343	0.187
Control: Std.Dev.	12.17	0.932	0.475	0.572

	(5) Voted in last election	(6) Voted for incumbent	(7) Use of WhatsApp	(8) Use of Facebook
High Saturation Municipal treatment	0.041	0.038	-0.004*	-0.010*
	(0.033)	(0.042)	(0.002)	(0.005)
Low Saturation Municipal treatment	0.001	-0.054	-0.004*	-0.010*
	(0.033)	(0.042)	(0.002)	(0.006)
Observations	2,020	2,020	2,020	2,020
R^2	0.053	0.062	0.017	0.018
Control outcome mean	0.558	0.184	1	0.990
Control outcome std. dev.	0.468	0.387	0	0.098

	(9) Intend to vote	(10) Intend to vote for incumbent	(11) Others' intention to vote	(12) Others' intention to vote for incumbent
High Saturation Municipal treatment	3.470	1.791	-0.979	-2.339
	(2.476)	(4.126)	(1.178)	(2.007)
Low Saturation Municipal treatment	1.659	-1.772	-2.436*	-3.146*
	(2.088)	(3.586)	(1.396)	(1.664)
Observations	2,020	2,020	2,020	2,020
R^2	0.015	0.072	0.047	0.050
Control outcome mean	80.19	15.23	74.93	47.69
Control outcome std. dev.	39.89	33.40	17.36	26.12

	(13) Priors about irregularities	(14) Prior precision about irregularities	(15) Incumbent seeking reelection	(16) Amount of FISM funds audited
High Saturation Municipal treatment	-2.146	0.007	-0.049	-53,939.3***
	(1.564)	(0.020)	(0.165)	(20,380.4)
Low Saturation Municipal treatment	-0.961	0.019	0.109	-28,495.7
	(1.411)	(0.015)	(0.159)	(21,893.5)
Observations	2,020	2,020	2,020	2,020
R^2	0.025	0.054	0.441	0.645
Control outcome mean	62.94	0.044	0.385	81,749.3
Control outcome std. dev.	24.29	0.204	0.487	75,832.3

	(17) Adult Population	(18) Average years of schooling	(19) Share illiterate	(20) Average number of occupants per room
High Saturation Municipal treatment	-51,046.1	0.465**	-0.016**	-0.101**
	(52,867.0)	(0.232)	(0.008)	(0.048)
Low Saturation Municipal treatment	4,487.7	0.264	-0.015*	-0.111**
	(61,647.5)	(0.252)	(0.008)	(0.049)
Observations	2,020	2,020	2,020	2,020
R^2	0.813	0.692	0.622	0.588
Control outcome mean	399,423.1	8.983	0.058	1.114
Control outcome std. dev.	343,963.2	1.284	0.058	0.245

	(21) Average children per women	(22) Share of households with electricity, water and drainage	(23) Share of population that is working age	(24) Share of households with internet
High Saturation Municipal treatment	-0.117**	0.063**	0.010**	0.051**
	(0.055)	(0.029)	(0.004)	(0.023)
Low Saturation Municipal treatment	-0.089	0.036	0.011**	0.039
	(0.055)	(0.030)	(0.005)	(0.024)
Observations	2,020	2,020	2,020	2,020
R^2	0.600	0.609	0.688	0.708
Control outcome mean	2.275	0.852	0.644	0.238
Control outcome std. dev.	0.284	0.156	0.031	0.111

Notes: Each specification is estimated using OLS, and includes randomization block fixed effects. Standard errors clustered at the municipal level are in parentheses. * denotes $p < 0.1$, ** denotes $p < 0.05$, *** denotes $p < 0.01$ from two-sided t tests.

Table A7: Effect of Facebook ads on precinct-level municipal incumbent party vote share (share of registered voters), by quartile of the sample irregularities distribution and excluding the lagged dependent variable

	Incumbent party vote (share of registered voters)			
	(1)	(2)	(3)	(4)
Facebook ads	0.036*	0.056***		
	(0.019)	(0.021)		
Facebook ads × Q3	-0.039	-0.027		
	(0.039)	(0.035)		
Facebook ads × Q4	-0.043	-0.061		
	(0.032)	(0.041)		
Spillover	0.014	0.033*		
	(0.018)	(0.019)		
Spillover × Q3	0.016	-0.014		
	(0.035)	(0.033)		
Spillover × Q4	-0.029	-0.052		
	(0.035)	(0.039)		
Facebook ads in high saturation			0.047**	0.055**
			(0.021)	(0.026)
Facebook ads in high saturation × Q3			-0.073*	-0.043
			(0.044)	(0.042)
Facebook ads in high saturation × Q4			-0.059	-0.020
			(0.037)	(0.056)
Facebook ads in low saturation			0.020	0.036
			(0.021)	(0.022)
Facebook ads in low saturation × Q3			0.022	-0.039
			(0.038)	(0.040)
Facebook ads in low saturation × Q4			-0.036	-0.059
			(0.039)	(0.047)
Spillover in high saturation			0.040*	0.048*
			(0.021)	(0.027)
Spillover in high saturation × Q3			-0.072	-0.045
			(0.048)	(0.043)
Spillover in high saturation × Q4			-0.053	-0.018
			(0.036)	(0.057)
Spillover in low saturation			0.005	0.020
			(0.021)	(0.021)
Spillover in low saturation × Q3			0.045	-0.013
			(0.037)	(0.038)
Spillover in low saturation × Q4			-0.018	-0.035
			(0.040)	(0.046)
Observations	13,278	13,278	13,278	13,278
R^2	0.45	0.57	0.51	0.62
Control outcome mean	0.18	0.18	0.18	0.18
Control outcome std. dev.	0.11	0.11	0.11	0.11
Test: null effect of Facebook ads in Q3 (p value, 2-sided)	0.919	0.295		
Test: null effect of Facebook ads in Q4 (p value, 2-sided)	0.800	0.834		
Test: null effect of spillover in Q3 (p value, 2-sided)	0.269	0.419		
Test: null effect of spillover in Q4 (p value, 2-sided)	0.616	0.499		
Test: larger effect of Facebook ads in high (vs. low) in Q1/Q2 (p value, 1-sided)			0.061	0.152
Test: larger effect of Facebook ads in high (vs. low) in Q3 (p value, 1-sided)			0.017	0.345
Test: larger effect of Facebook ads in high (vs. low) in Q4 (p value, 1-sided)			0.556	0.058
Test: larger effect of spillovers in high (vs. low) in Q1/Q2 (p value, 1-sided)			0.014	0.062
Test: larger effect of spillovers in high (vs. low) in Q3 (p value, 1-sided)			0.020	0.646
Test: larger effect of spillovers in high (vs. low) in Q4 (p value, 1-sided)			0.554	0.057
Interactive covariates		✓		✓

Notes: Each specification is estimated using OLS, and includes randomization block fixed effects. Specifications including interactive covariates further include interactions between treatment conditions and the following municipal-level covariates: year of audit; amount of FISM funds received; population aged above 18; lagged incumbent party vote share; average years of schooling; share of the population that is illiterate; average occupants per room; average number of children per woman; the share of the population with electricity, water, and drainage in their home; the working age share of the population; and the share of households with internet at home. The omitted irregularities category is Q1/Q2. All segments are weighted equally and by the inverse probability of treatment assignment. Standard errors clustered by municipality are in parentheses.* denotes $p < 0.1$, ** denotes $p < 0.05$, *** denotes $p < 0.01$ from two-sided t tests. The one-sided coefficient tests at the foot of the table test whether the effect of high saturation is larger in a given direction in high saturation than in low saturation municipalities.

Table A8: Effect of Facebook ads on precinct-level municipal incumbent party vote share (share of registered voters), by quartile of the sample irregularities distribution and weighting by registered voters

	Incumbent party vote (share of registered voters)			
	(1)	(2)	(3)	(4)
Facebook ads	0.043**	0.060***		
	(0.018)	(0.021)		
Facebook ads × Q3	-0.058*	-0.044		
	(0.032)	(0.034)		
Facebook ads × Q4	-0.043*	-0.113**		
	(0.025)	(0.044)		
Spillover	0.025	0.041**		
	(0.017)	(0.020)		
Spillover × Q3	-0.026	-0.043		
	(0.028)	(0.033)		
Spillover × Q4	-0.033	-0.108**		
	(0.028)	(0.043)		
Facebook ads in high saturation			0.049**	0.069**
			(0.021)	(0.027)
Facebook ads in high saturation × Q3			-0.070*	-0.050
			(0.036)	(0.041)
Facebook ads in high saturation × Q4			-0.046*	-0.087
			(0.027)	(0.054)
Facebook ads in low saturation			0.019	0.031
			(0.020)	(0.023)
Facebook ads in low saturation × Q3			-0.035	-0.081**
			(0.033)	(0.036)
Facebook ads in low saturation × Q4			-0.034	-0.104**
			(0.030)	(0.047)
Spillover in high saturation			0.054**	0.073**
			(0.022)	(0.028)
Spillover in high saturation × Q3			-0.100**	-0.072*
			(0.040)	(0.042)
Spillover in high saturation × Q4			-0.051*	-0.079
			(0.028)	(0.058)
Spillover in low saturation			0.012	0.029
			(0.020)	(0.023)
Spillover in low saturation × Q3			0.003	-0.054
			(0.031)	(0.036)
Spillover in low saturation × Q4			-0.019	-0.095**
			(0.032)	(0.047)
Observations	13,251	13,251	13,251	13,251
R^2	0.60	0.66	0.64	0.74
Control outcome mean	0.17	0.17	0.17	0.17
Control outcome std. dev.	0.10	0.10	0.10	0.10
Test: null effect of Facebook ads in Q3 (*p* value, 2-sided)	0.561	0.523		
Test: null effect of Facebook ads in Q4 (*p* value, 2-sided)	0.998	0.085		
Test: null effect of spillover in Q3 (*p* value, 2-sided)	0.969	0.948		
Test: null effect of spillover in Q4 (*p* value, 2-sided)	0.688	0.036		
Test: larger effect of Facebook ads in high (vs. low) in Q1/Q2 (*p* value, 1-sided)			0.043	0.010
Test: larger effect of Facebook ads in high (vs. low) in Q3 (*p* value, 1-sided)			0.430	0.017
Test: larger effect of Facebook ads in high (vs. low) in Q4 (*p* value, 1-sided)			0.172	0.946
Test: larger effect of spillovers in high (vs. low) in Q1/Q2 (*p* value, 1-sided)			0.010	0.005
Test: larger effect of spillovers in high (vs. low) in Q3 (*p* value, 1-sided)			0.113	0.108
Test: larger effect of spillovers in high (vs. low) in Q4 (*p* value, 1-sided)			0.233	0.979
Interactive covariates		✓		✓

Notes: Each specification is estimated using OLS, and includes a lagged dependent variable and randomization block fixed effects and an interaction between treatment. Specifications including interactive covariates further include interactions between treatment conditions and the following municipal-level covariates: year of audit; amount of FISM funds received; population aged above 18; lagged incumbent party vote share; average years of schooling; share of the population that is illiterate; average occupants per room; average number of children per woman; the share of the population with electricity, water, and drainage in their home; the working age share of the population; and the share of households with internet at home. The omitted irregularities category is Q1/Q2. All precincts are weighted voting age population and by the inverse probability of treatment assignment. Standard errors clustered by municipality are in parentheses.* denotes $p < 0.1$, ** denotes $p < 0.05$, *** denotes $p < 0.01$ from two-sided t tests. The one-sided coefficient tests at the foot of the table test whether the effect of high saturation is larger in a given direction in high saturation than in low saturation municipalities.

Table A9: Effect of Facebook ads on precinct-level municipal incumbent party vote share (share of registered voters), by quartile of the sample irregularities distribution and weighting municipalities with different numbers of segments equally

	Incumbent party vote (share of registered voters)			
	(1)	(2)	(3)	(4)
Facebook ads	0.036*	0.056***		
	(0.019)	(0.021)		
Facebook ads × Q3	-0.039	-0.027		
	(0.039)	(0.035)		
Facebook ads × Q4	-0.043	-0.061		
	(0.032)	(0.041)		
Spillover	0.014	0.033*		
	(0.018)	(0.019)		
Spillover × Q3	0.016	-0.014		
	(0.035)	(0.033)		
Spillover × Q4	-0.029	-0.052		
	(0.035)	(0.039)		
Facebook ads in high saturation			0.047**	0.055**
			(0.021)	(0.026)
Facebook ads in high saturation × Q3			-0.073*	-0.043
			(0.044)	(0.042)
Facebook ads in high saturation × Q4			-0.059	-0.020
			(0.037)	(0.056)
Facebook ads in low saturation			0.020	0.036
			(0.021)	(0.022)
Facebook ads in low saturation × Q3			0.022	-0.039
			(0.038)	(0.040)
Facebook ads in low saturation × Q4			-0.036	-0.059
			(0.039)	(0.047)
Spillover in high saturation			0.040*	0.048*
			(0.021)	(0.027)
Spillover in high saturation × Q3			-0.072	-0.045
			(0.048)	(0.043)
Spillover in high saturation × Q4			-0.053	-0.018
			(0.036)	(0.057)
Spillover in low saturation			0.005	0.020
			(0.021)	(0.021)
Spillover in low saturation × Q3			0.045	-0.013
			(0.037)	(0.038)
Spillover in low saturation × Q4			-0.018	-0.035
			(0.040)	(0.046)
Observations	13,278	13,278	13,278	13,278
R^2	0.45	0.57	0.51	0.62
Control outcome mean	0.18	0.18	0.18	0.18
Control outcome std. dev.	0.11	0.11	0.11	0.11
Test: null effect of Facebook ads in Q3 (p value, 2-sided)	0.919	0.295		
Test: null effect of Facebook ads in Q4 (p value, 2-sided)	0.800	0.834		
Test: null effect of spillover in Q3 (p value, 2-sided)	0.269	0.419		
Test: null effect of spillover in Q4 (p value, 2-sided)	0.616	0.499		
Test: larger effect of Facebook ads in high (vs. low) in Q1/Q2 (p value, 1-sided)			0.061	0.152
Test: larger effect of Facebook ads in high (vs. low) in Q3 (p value, 1-sided)			0.017	0.345
Test: larger effect of Facebook ads in high (vs. low) in Q4 (p value, 1-sided)			0.556	0.058
Test: larger effect of spillovers in high (vs. low) in Q1/Q2 (p value, 1-sided)			0.014	0.062
Test: larger effect of spillovers in high (vs. low) in Q3 (p value, 1-sided)			0.020	0.646
Test: larger effect of spillovers in high (vs. low) in Q4 (p value, 1-sided)			0.554	0.057
Interactive covariates		✓		✓

Notes: Each specification is estimated using OLS, and includes a lagged dependent variable and randomization block fixed effects and an interaction between treatment. Specifications including interactive covariates further include interactions between treatment conditions and the following municipal-level covariates: year of audit; amount of FISM funds received; population aged above 18; lagged incumbent party vote share; average years of schooling; share of the population that is illiterate; average occupants per room; average number of children per woman; the share of the population with electricity, water, and drainage in their home; the working age share of the population; and the share of households with internet at home. The omitted irregularities category is Q1/Q2. All municipalities are weighted equally and by the inverse probability of treatment assignment. Standard errors clustered by municipality are in parentheses.* denotes $p < 0.1$, ** denotes $p < 0.05$, *** denotes $p < 0.01$ from two-sided t tests. The one-sided coefficient tests at the foot of the table test whether the effect of high saturation is larger in a given direction in high saturation than in low saturation municipalities.

no differential effects on voting behavior. First, Table A10 shows that the common knowledge treatment was no more likely to be engaged with than the non-common knowledge variant of the treatment. We observe no statistically significant difference in user interactions for the common knowledge campaigns. Second, Table A11 ultimately shows that no notable difference in the effect of the common knowledge and non-common knowledge variants of the Facebook ads. These findings may not be especially surprising, given that the common knowledge information only appeared toward the end of the ad or otherwise required that users read the comments or page associated with the ad.

A.10 Null effects of municipal level treatments on incumbent election victory

Table A12 reports no discernible effect of the municipal-level treatments on whether the municipal incumbent party was re-elected.

A.11 Results for other voting outcome measures

Table A13 reports the average treatment on incumbent vote share as a share of turnout.

Tables A14 and A15 report the treatment effects by irregularities quartile and saturation level for incumbent vote share, as a share of turnout, and turnout. The vote share results are similar to those reported in the main text, while the effects on turnout suggest that aggregate changes in turnout are not driving the aggregate changes in incumbent party vote share.

A.12 Robustness to alternative operationalization of irregularities

Table A16 shows that the results are robust to using two alternative operationalizations of the content report. First, columns (1)-(4) compare municipalities with exactly 0% irregularities to municipalities with some irregularities. Second, columns (5)-(8) compare municipalities with at most 5% irregularities to municipalities with greater than 5% irregularities.

A.13 Tests of Facebook mistargeting

Tables A17 and A18 report the results of the geographic mistargeting tests described in the main text. Table A17 shows that precincts in spillover segments with high levels of 3G+ coverage are no more likely to change their voting behavior than precincts without good 3G+ coverage. As noted in the main article, this suggests that geographic mistargeting is unlikely to explain our findings because geographic mistargeting should have larger effects in nearby locations with easy access to Facebook.

Table A10: Effect of municipal treatments on municipal Facebook engagement, by common knowledge treatment

	Municipal counts per capita (normalized by 2015 adult population)								
	Paid-for impressions (1)	Organic impressions (2)	Paid-for unique viewers (3)	Organic unique viewers (4)	Unique user page engagements (5)	Total views (of 3 seconds) (6)	Unique views (of 3 seconds) (7)	Total views (of entire video) (8)	Unique views (of entire video) (9)
Facebook ads (no common knowledge)	0.748***	0.038***	0.215***	0.021***	0.011***	0.143***	0.088***	0.035***	0.031***
	(0.092)	(0.006)	(0.025)	(0.004)	(0.001)	(0.017)	(0.010)	(0.004)	(0.003)
Facebook ads (common knowledge)	0.682***	0.036***	0.200***	0.020***	0.010***	0.132***	0.082***	0.032***	0.028***
	(0.094)	(0.007)	(0.025)	(0.004)	(0.002)	(0.017)	(0.010)	(0.004)	(0.004)
Observations	173	173	173	173	173	173	173	173	173
R^2	0.50	0.56	0.51	0.56	0.50	0.52	0.52	0.53	0.52
Control outcome mean	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
Control outcome std. dev.	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
Test: No common knowledge = Common knowledge (p value, 2-sided)	0.511	0.747	0.567	0.669	0.426	0.562	0.558	0.416	0.429

Notes: Each specification is estimated using OLS, and includes randomization block fixed effects. Standard errors are clustered by municipality. * denotes $p < 0.1$, ** denotes $p < 0.05$, *** denotes $p < 0.01$ from two-sided t tests.

Table A11: Differential effects of Facebooks ads in the 2018 municipal elections, by common knowledge treatment

	Incumbent party vote (share of turnout)			
	(1)	(2)	(3)	(4)
Facebook	0.020*	0.039**		
	(0.011)	(0.016)		
Facebook × Q3		-0.045		
		(0.033)		
Facebook × Q4		-0.042		
		(0.028)		
Facebook + CK	0.020*	0.042**		
	(0.011)	(0.016)		
Facebook + CK × Q3		-0.043		
		(0.031)		
Facebook + CK × Q4		-0.063**		
		(0.029)		
Spillover	0.008			
	(0.012)			
Spillover		0.022		
		(0.016)		
Spillover × Q3		0.004		
		(0.029)		
Spillover × Q4		-0.044		
		(0.032)		
Facebook ads in high saturation			0.024*	0.050**
			(0.013)	(0.019)
Facebook ads in high saturation × Q3				-0.066*
				(0.038)
Facebook ads in high saturation × Q4				-0.065**
				(0.032)
Facebook ads + CK in high saturation			0.023*	0.051***
			(0.013)	(0.019)
Facebook ads in low saturation			0.004	0.022
			(0.015)	(0.022)
Facebook ads in low saturation × Q3				-0.034
				(0.031)
Facebook ads in low saturation × Q4				-0.029
				(0.033)
Facebook ads + CK in low saturation			0.011	0.025
			(0.015)	(0.020)
Facebook ads + CK in low saturation × Q3				0.006
				(0.033)
Facebook ads + CK in low saturation × Q4				-0.072
				(0.043)
Spillover in high saturation			0.020	0.048**
			(0.013)	(0.019)
Spillover in low saturation			0.006	0.014
			(0.013)	(0.018)
Spillover in low saturation × Q3				0.028
				(0.031)
Spillover in low saturation × Q4				-0.039
				(0.036)
Spillover in high saturation × Q3				-0.077**
				(0.038)
Spillover in high saturation × Q4				-0.065**
				(0.031)
Facebook ads + CK in high saturation × Q3				-0.074**
				(0.036)
Facebook ads + CK in high saturation × Q4				-0.068**
				(0.032)
Observations	13,251	13,251	13,251	13,251
R^2	0.56	0.60	0.58	0.64
Control outcome mean	0.18	0.18	0.18	0.18
Control outcome std. dev.	0.11	0.11	0.11	0.11

Notes: Each specification is estimated using OLS, and includes a lagged dependent variable and randomization block fixed effects. All segments are weighted equally and by the inverse probability of treatment assignment. Standard errors clustered by municipality are in parentheses. * denotes $p < 0.1$, ** denotes $p < 0.05$, *** denotes $p < 0.01$ from two-sided t tests.

Table A12: Effects of municipal saturation treatments on municipal election outcomes

	Municipal incumbent party re-elected					
	(1)	(2)	(3)	(4)	(5)	(6)
Any Saturation	-0.078	-0.020				
	(0.076)	(0.133)				
Any Saturation × Q3		-0.065				
		(0.227)				
Any Saturation × Q4		-0.151				
		(0.232)				
High saturation			-0.047	0.062		
			(0.090)	(0.147)		
High saturation × Q3				-0.253		
				(0.285)		
High saturation × Q4				-0.253		
				(0.238)		
Low saturation			-0.109	-0.119		
			(0.092)	(0.157)		
Low saturation × Q3				0.152		
				(0.233)		
Low saturation × Q4				-0.042		
				(0.285)		
Saturation					-0.020	0.150
					(0.111)	(0.169)
Saturation × Q3						-0.412
						(0.333)
Saturation × Q4						-0.347
						(0.267)
Observations	124	124	124	124	124	124
R^2	0.48	0.53	0.48	0.55	0.47	0.54
Control outcome mean	0.39	0.39	0.39	0.39	0.39	0.39
Control outcome std. dev.	0.49	0.49	0.49	0.49	0.49	0.49
Test: low = high (p value, 2-sided)			0.534			
Test: effect of treatment in Q3 (p value, 2-sided)		0.587				0.331
Test: effect of treatment in Q4 (p value, 2-sided)		0.338				0.313
Test: effect of low saturation in Q3 (p value, 2-sided)				0.822		
Test: effect of low saturation in Q4 (p value, 2-sided)				0.474		
Test: effect of high saturation in Q3 (p value, 2-sided)				0.391		
Test: effect of high saturation in Q4 (p value, 2-sided)				0.276		

Notes: Each specification is estimated using OLS, and includes randomization block fixed effects. Robust standard errors are in parentheses. * denotes $p < 0.1$, ** denotes $p < 0.05$, *** denotes $p < 0.01$ from two-sided t tests.

Table A13: Effect of Facebook ads on precinct-level municipal incumbent party vote share (share of turnout)

	Incumbent party vote (share of turnout)		
	(1)	(2)	(3)
High saturation	0.019		
	(0.018)		
Low saturation	0.015		
	(0.018)		
Facebook ads		0.026*	
		(0.016)	
Spillover		0.011	
		(0.016)	
Facebook ads in high saturation			0.031*
			(0.018)
Facebook ads in low saturation			0.010
			(0.018)
Spillover in high saturation			0.028
			(0.019)
Spillover in low saturation			0.008
			(0.019)
Observations	13,251	13,251	13,251
R^2	0.50	0.51	0.53
Control outcome mean	0.28	0.28	0.28
Control outcome std. dev.	0.15	0.14	0.14
Test: low $\geq$ high (p value, 1-sided)	0.832		
Test: spillover $\geq$ direct (p value, 1-sided)		0.107	
Test within ads treatment: low $\geq$ high (p value, 1-sided)			0.143
Test within spillovers: low $\geq$ high (p value, 1-sided)			0.158

Notes: Each specification is estimated using OLS, and includes a lagged dependent variable and randomization block fixed effects. All segments are weighted equally and by the inverse probability of (municipal or segment, as appropriate) treatment assignment. Standard errors clustered by municipality are in parentheses. * denotes $p < 0.1$, ** denotes $p < 0.05$, *** denotes $p < 0.01$ from two-sided t tests.

Table A14: Effect of Facebook ads on precinct-level municipal incumbent party vote share (share of turnout), by quartile of the sample irregularities distribution

	Incumbent party vote (share of turnout)			
	(1)	(2)	(3)	(4)
Facebook ads	0.058	0.088		
	(0.022)	(0.028)		
Facebook ads × Q3	-0.087	-0.071		
	(0.044)	(0.047)		
Facebook ads × Q4	-0.061	-0.118		
	(0.038)	(0.055)		
Spillover	0.029	0.055		
	(0.021)	(0.026)		
Spillover × Q3	0.000	-0.027		
	(0.039)	(0.045)		
Spillover × Q4	-0.047	-0.096		
	(0.043)	(0.054)		
Facebook ads in high saturation			0.074	0.092
			(0.026)	(0.035)
Facebook ads in high saturation × Q3			-0.136	-0.097
			(0.053)	(0.058)
Facebook ads in high saturation × Q4			-0.079	-0.078
			(0.044)	(0.073)
Facebook ads in low saturation			0.025	0.046
			(0.026)	(0.029)
Facebook ads in low saturation × Q3			0.003	-0.072
			(0.044)	(0.052)
Facebook ads in low saturation × Q4			-0.038	-0.086
			(0.050)	(0.065)
Spillover in high saturation			0.072	0.086
			(0.026)	(0.036)
Spillover in high saturation × Q3			-0.142	-0.096
			(0.054)	(0.059)
Spillover in high saturation × Q4			-0.079	-0.069
			(0.043)	(0.075)
Spillover in low saturation			0.012	0.034
			(0.025)	(0.028)
Spillover in low saturation × Q3			0.048	-0.017
			(0.044)	(0.052)
Spillover in low saturation × Q4			-0.029	-0.074
			(0.050)	(0.063)
Observations	13,251	13,251	13,251	13,251
R^2	0.56	0.63	0.59	0.68
Control outcome mean	0.28	0.28	0.28	0.28
Control outcome std. dev.	0.14	0.14	0.14	0.14
Test: null effect of Facebook ads in Q3 (p value, 2-sided)	0.412	0.625		
Test: null effect of Facebook ads in Q4 (p value, 2-sided)	0.916	0.405		
Test: null effect of spillover in Q3 (p value, 2-sided)	0.352	0.404		
Test: null effect of spillover in Q4 (p value, 2-sided)	0.626	0.294		
Test: larger effect of Facebook ads in high (vs. low) in Q1/Q2 (p value, 1-sided)			0.030	0.026
Test: larger effect of Facebook ads in high (vs. low) in Q3 (p value, 1-sided)			0.016	0.677
Test: larger effect of Facebook ads in high (vs. low) in Q4 (p value, 1-sided)			0.586	0.143
Test: larger effect of spillovers in high (vs. low) in Q1/Q2 (p value, 1-sided)			0.010	0.017
Test: larger effect of spillovers in high (vs. low) in Q3 (p value, 1-sided)			0.010	0.448
Test: larger effect of spillovers in high (vs. low) in Q4 (p value, 1-sided)			0.688	0.055
Interactive covariates		✓		✓

Notes: Each specification is estimated using OLS, and includes a lagged dependent variable and randomization block fixed effects and an interaction between treatment. Specifications including interactive covariates further include interactions between treatment conditions and the following municipal-level covariates: year of audit; amount of FISM funds received; population aged above 18; lagged incumbent party vote share; average years of schooling; share of the population that is illiterate; average occupants per room; average number of children per woman; the share of the population with electricity, water, and drainage in their home; the working age share of the population; and the share of households with internet at home. The omitted irregularities category is Q1/Q2. All segments are weighted equally and by the inverse probability of treatment assignment. Standard errors clustered by municipality are in parentheses.* denotes $p < 0.1$, ** denotes $p < 0.05$, *** denotes $p < 0.01$ from two-sided t tests. The one-sided coefficient tests at the foot of the table test whether the effect of high saturation is larger in a given direction in high saturation than in low saturation municipalities.

Table A15: Effect of Facebook ads on precinct-level municipal turnout, by quartile of the sample irregularities distribution

	Incumbent party vote (share of turnout)			
	(1)	(2)	(3)	(4)
Facebook ads	0.008	-0.007		
	(0.012)	(0.011)		
Facebook ads × Q3	0.025	0.046		
	(0.020)	(0.017)		
Facebook ads × Q4	-0.012	0.023		
	(0.022)	(0.023)		
Spillover	0.006	-0.008		
	(0.013)	(0.010)		
Spillover × Q3	0.005	0.020		
	(0.020)	(0.018)		
Spillover × Q4	-0.019	0.001		
	(0.020)	(0.020)		
Facebook ads in high saturation			0.004	-0.002
			(0.013)	(0.011)
Facebook ads in high saturation × Q3			0.038	0.043
			(0.023)	(0.019)
Facebook ads in high saturation × Q4			-0.006	0.030
			(0.025)	(0.025)
Facebook ads in low saturation			0.006	-0.010
			(0.014)	(0.011)
Facebook ads in low saturation × Q3			0.004	0.036
			(0.023)	(0.018)
Facebook ads in low saturation × Q4			-0.027	-0.009
			(0.024)	(0.021)
Spillover in high saturation			-0.001	-0.004
			(0.013)	(0.012)
Spillover in high saturation × Q3			0.032	0.031
			(0.023)	(0.020)
Spillover in high saturation × Q4			-0.001	0.026
			(0.023)	(0.025)
Spillover in low saturation			0.007	-0.005
			(0.014)	(0.011)
Spillover in low saturation × Q3			-0.003	0.014
			(0.023)	(0.019)
Spillover in low saturation × Q4			-0.022	-0.011
			(0.023)	(0.022)
Observations	13,251	13,251	13,251	13,251
R^2	0.63	0.69	0.65	0.71
Control outcome mean	0.64	0.64	0.64	0.64
Control outcome std. dev.	0.12	0.12	0.12	0.12
Test: null effect of Facebook ads in Q3 (p value, 2-sided)	0.026	0.001		
Test: null effect of Facebook ads in Q4 (p value, 2-sided)	0.792	0.364		
Test: null effect of spillover in Q3 (p value, 2-sided)	0.430	0.343		
Test: null effect of spillover in Q4 (p value, 2-sided)	0.383	0.628		
Test: larger effect of Facebook ads in high (vs. low) in Q1/Q2 (p value, 1-sided)			0.582	0.863
Test: larger effect of Facebook ads in high (vs. low) in Q3 (p value, 1-sided)			0.024	0.120
Test: larger effect of Facebook ads in high (vs. low) in Q4 (p value, 1-sided)			0.782	0.001
Test: larger effect of spillovers in high (vs. low) in Q1/Q2 (p value, 1-sided)			0.223	0.549
Test: larger effect of spillovers in high (vs. low) in Q3 (p value, 1-sided)			0.053	0.241
Test: larger effect of spillovers in high (vs. low) in Q4 (p value, 1-sided)			0.782	0.045
Interactive covariates		✓		✓

Notes: Each specification is estimated using OLS, and includes a lagged dependent variable and randomization block fixed effects and an interaction between treatment. Specifications including interactive covariates further include interactions between treatment conditions and the following municipal-level covariates: year of audit; amount of FISM funds received; population aged above 18; lagged incumbent party vote share; average years of schooling; share of the population that is illiterate; average occupants per room; average number of children per woman; the share of the population with electricity, water, and drainage in their home; the working age share of the population; and the share of households with internet at home. The omitted irregularities category is Q1/Q2. All segments are weighted equally and by the inverse probability of treatment assignment. Standard errors clustered by municipality are in parentheses.* denotes $p < 0.1$, ** denotes $p < 0.05$, *** denotes $p < 0.01$ from two-sided t tests. The one-sided coefficient tests at the foot of the table test whether the effect of high saturation is larger in a given direction in high saturation than in low saturation municipalities.

Table A16: Effect of Facebook ads on precinct-level municipal incumbent party vote share (share of registered voters), by binary operationalizations of irregularities

	Incumbent party vote (share of registered voters)											
	Higher malfeasance measure:											
	Pooling Q3 and Q4				>0% irregularities				>5% irregularities			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)
Facebook ads	0.044***	0.055***			0.042**	0.053**			0.033**	0.051***		
	(0.017)	(0.020)			(0.017)	(0.021)			(0.014)	(0.016)		
Facebook ads × 1.higher_malfeasance	-0.057**	-0.060*			-0.048*	-0.051			-0.068**	-0.094***		
	(0.025)	(0.033)			(0.025)	(0.032)			(0.027)	(0.034)		
Spillover	0.022	0.030			0.023	0.029			0.024*	0.032**		
	(0.015)	(0.018)			(0.015)	(0.018)			(0.012)	(0.014)		
Spillover × 1.higher_malfeasance	-0.029	-0.046			-0.028	-0.041			-0.071**	-0.090***		
	(0.025)	(0.031)			(0.025)	(0.031)			(0.030)	(0.033)		
Facebook ads in high saturation			0.054***	0.064**			0.050**	0.057**			0.037**	0.049**
			(0.019)	(0.025)			(0.019)	(0.025)			(0.016)	(0.020)
Facebook ads in high saturation × 1.higher_malfeasance			-0.076***	-0.057			-0.061**	-0.037			-0.079***	-0.054
			(0.029)	(0.041)			(0.029)	(0.041)			(0.029)	(0.044)
Facebook ads in low saturation			0.021	0.026			0.024	0.024			0.026*	0.018
			(0.019)	(0.020)			(0.019)	(0.020)			(0.015)	(0.016)
Facebook ads in low saturation × 1.higher_malfeasance			-0.030	-0.062*			-0.035	-0.052			-0.076**	-0.084**
			(0.028)	(0.034)			(0.028)	(0.034)			(0.034)	(0.038)
Spillover in high saturation			0.051***	0.060**			0.048**	0.055**			0.033**	0.043**
			(0.019)	(0.025)			(0.019)	(0.026)			(0.016)	(0.020)
Spillover in high saturation × 1.higher_malfeasance			-0.079***	-0.059			-0.067**	-0.044			-0.076**	-0.046
			(0.029)	(0.042)			(0.028)	(0.041)			(0.029)	(0.045)
Spillover in low saturation			0.012	0.019			0.015	0.016			0.024	0.020
			(0.018)	(0.020)			(0.018)	(0.020)			(0.015)	(0.015)
Spillover in low saturation × 1.higher_malfeasance			-0.012	-0.042			-0.017	-0.031			-0.074**	-0.080**
			(0.028)	(0.034)			(0.028)	(0.034)			(0.034)	(0.038)
Observations	13,251	13,251	13,251	13,251	13,251	13,251	13,251	13,251	13,251	13,251	13,251	13,251
R^2	0.56	0.62	0.60	0.69	0.56	0.62	0.59	0.69	0.58	0.64	0.61	0.70
Control outcome mean	0.18	0.18	0.18	0.18	0.18	0.18	0.18	0.18	0.18	0.18	0.18	0.18
Control outcome std. dev.	0.11	0.11	0.11	0.11	0.11	0.11	0.11	0.11	0.11	0.11	0.11	0.11
Test: null effect of Facebook ads with higher malfeasance (p value, 2-sided)	0.445	0.826			0.718	0.902			0.098	0.071		
Test: null effect of spillover with higher malfeasance (p value, 2-sided)	0.723	0.426			0.756	0.535			0.080	0.026		
Test: larger effect of Facebook ads in high (vs. low) with lower malfeasance (p value, 1-sided)			0.030	0.020			0.073	0.040			0.217	0.012
Test: larger effect of Facebook ads in high (vs. low) with higher malfeasance (p value, 1-sided)			0.038	0.550			0.174	0.676			0.460	0.232
Test: larger effect of spillovers in high (vs. low) with lower malfeasance (p value, 1-sided)			0.007	0.015			0.020	0.022			0.261	0.042
Test: larger effect of spillovers in high (vs. low) with higher malfeasance (p value, 1-sided)			0.996	0.693			0.032	0.649			0.478	0.211
Interactive covariates		✓		✓		✓		✓				

Notes: Each specification is estimated using OLS, and includes a lagged dependent variable and randomization block fixed effects and an interaction between treatment. Specifications including interactive covariates further include interactions between treatment conditions and the following municipal-level covariates: year of audit; amount of FISM funds received; population aged above 18; lagged incumbent party vote share; average years of schooling; share of the population that is illiterate; average occupants per room; average number of children per woman; the share of the population with electricity, water, and drainage in their home; the working age share of the population; and the share of households with internet at home. The omitted irregularities category is Q1/Q2. All segments are weighted equally and by the inverse probability of treatment assignment. Standard errors clustered by municipality are in parentheses.* denotes $p < 0.1$, ** denotes $p < 0.05$, *** denotes $p < 0.01$ from two-sided t tests. The one-sided coefficient tests at the foot of the table test whether the effect of high saturation is larger in a given direction in high saturation than in low saturation municipalities.

Table A17: Effect of Facebook ads spillovers on precinct-level municipal incumbent party vote share (share of registered voters), by quartile of the sample irregularities distribution and access to 3G+ cell phone coverage

	Incumbent party vote (share of registered voters)	
	(1)	(2)
Spillover	0.034*	
	(0.018)	
Spillover × Q3	0.025	
	(0.045)	
Spillover × Q4	-0.098**	
	(0.045)	
Spillover × 3G+	-0.007	
	(0.020)	
Spillover × Q3 × 3G+	-0.066	
	(0.050)	
Spillover × Q4 × 3G+	0.067	
	(0.045)	
Spillover in high saturation		0.056
		(0.037)
Spillover in high saturation × Q3		0.074
		(0.070)
Spillover in high saturation × Q4		-0.109**
		(0.044)
Spillover in high saturation × 3G+		0.019
		(0.038)
Spillover in high saturation × Q3 × 3G+		-0.221***
		(0.070)
Spillover in high saturation × Q4 × 3G+		0.055
		(0.047)
Spillover in low saturation		0.024
		(0.024)
Spillover in low saturation × Q3		0.016
		(0.040)
Spillover in low saturation × Q4		-0.120**
		(0.057)
Spillover in low saturation × 3G+		-0.013
		(0.024)
Spillover in low saturation × Q3 × 3G+		-0.034
		(0.044)
Spillover in low saturation × Q4 × 3G+		0.099*
		(0.056)
Observations	8,683	8,683
R^2	0.65	0.71
Control outcome mean	0.18	0.18
Control outcome std. dev.	0.11	0.11
Mean 3G+	0.816	0.823

Notes: Each specification is estimated using OLS, and includes a lagged dependent variable and randomization block fixed effects and an interaction between treatment. Specifications including interactive covariates further include interactions between treatment conditions and the following municipal-level covariates: year of audit; amount of FISM funds received; population aged above 18; lagged incumbent party vote share; average years of schooling; share of the population that is illiterate; average occupants per room; average number of children per woman; the share of the population with electricity, water, and drainage in their home; the working age share of the population; and the share of households with internet at home. The omitted irregularities category is Q1/Q2. All segments are weighted equally and by the inverse probability of treatment assignment. Standard errors clustered by municipality are in parentheses.* denotes $p < 0.1$, ** denotes $p < 0.05$, *** denotes $p < 0.01$ from two-sided t tests. The one-sided coefficient tests at the foot of the table test whether the effect of high saturation is larger in a given direction in high saturation than in low saturation municipalities.

Table A18 further examines cross-border spillovers to municipalities that should not have received Facebook ads. Specifically, we created a sample of precincts in non-experimental municipalities that are within 5 kilometers of an experimental (i.e. treated or control) segment. In order to avoid treatment correlating with the number of segments a cross-border spillover precinct is close to, we further restrict this sample to precincts for which only one experimental segment was within 5 kilometers. Consequently, we leverage the same experimental sources of variation used for the main analysis. This yielded a sample of 642 precincts within 5 kilometers of 184 different experimental segments from 88 of the experimental municipalities. Due to this smaller sample size, randomization block fixed effects are excluded (although the results are robust to their inclusion).

The results in columns (1) and (2) of Table A18 indicate that proximity of precincts in non-experimental municipalities to nearby segments that were (directly or indirectly) treated in experimental municipalities does not affect the precinct vote share of the incumbent party in the nearby experimental municipality. Columns (3) and (4) further show that the saturation level in the municipality of the nearby segment also does not affect this vote share. Although this test leverages only around 30% of the segments used for the main analysis, which is why we do not interact direct and indirect treatment with saturation, we again find little evidence to suggest that the mistargeting of Facebook ads explains our main findings.

A.14 Belief updating in the WhatsApp survey

One mechanism through which the spillover and saturation effects could influence voting behavior is via belief updating. The preceding findings that the Facebook ads induced greater support for the best-performing incumbents and potentially somewhat reduced support for the worst-performing incumbents is consistent with voters updating their relatively pessimistic prior beliefs. While underpowered by the low treatment compliance rate in the survey (see Table 2), a more direct test of the belief updating mechanism, in a context that is unlikely to activate coordination mechanisms, leverages the WhatsApp version of the treatment within our panel survey.

Although they are not statistically significant, the results in Table A19 are broadly in line with the precinct-level vote share outcomes. The 1 percentage point decrease in perceived irregularities in Q1/Q2 is broadly in line with the electoral returns data, but is small in magnitude. However, this reduced form effect would grow considerably after accounting for only 9% of voters recalling the WhatsApp message and video ad. These results thus suggest that belief updating may account for some of the observed effects on incumbent support, but are not definitive. The magnitude of the effects also suggests that other mechanisms could also be at play.

Table A18: Effect of Facebook ads spillovers on precinct-level vote share of nearby municipal incumbent parties (share of registered voters), by quartile of the sample irregularities distribution

	Vote for incumbent party in nearby experimental municipality (share of registered voters)			
	(1)	(2)	(3)	(4)
Facebook ads	0.008	-0.003		
	(0.040)	(0.026)		
Facebook ads × Q3	-0.014	0.013		
	(0.047)	(0.044)		
Facebook ads × Q4	0.024	0.090*		
	(0.052)	(0.054)		
Spillover	-0.011	-0.015		
	(0.044)	(0.034)		
Spillover × Q3	0.052	0.043		
	(0.062)	(0.052)		
Spillover × Q4	0.088	0.106*		
	(0.066)	(0.060)		
High saturation			0.022	0.013
			(0.041)	(0.026)
High saturation × Q3			-0.002	0.004
			(0.045)	(0.044)
High saturation × Q4			-0.006	0.097*
			(0.054)	(0.051)
Low saturation			-0.024	0.013
			(0.045)	(0.030)
Low saturation × Q3			0.039	0.014
			(0.069)	(0.061)
Low saturation × Q4			0.117*	0.132**
			(0.066)	(0.062)
Observations	642	642	642	642
R^2	0.16	0.46	0.17	0.57
Control outcome mean	0.16	0.16	0.16	0.16
Control outcome std. dev.	0.11	0.11	0.11	0.11
Test: null effect of Facebook ads in Q3 (p value, 2-sided)	0.825	0.744		
Test: null effect of Facebook ads in Q4 (p value, 2-sided)	0.379	0.036		
Test: null effect of spillover in Q3 (p value, 2-sided)	0.362	0.439		
Test: null effect of spillover in Q4 (p value, 2-sided)	0.117	0.064		
Test: larger effect in high (vs. low) saturation in Q1/Q2 (p value, 1-sided)			0.128	0.503
Test: larger effect in high (vs. low) saturation in Q3 (p value, 1-sided)			0.463	0.566
Test: larger effect in high (vs. low) saturation in Q4 (p value, 1-sided)			0.947	0.727
Interactive covariates		✓		✓

Notes: Each specification is estimated using OLS, and includes a lagged dependent variable. Specifications including interactive covariates further include interactions between treatment conditions and the following municipal-level covariates: year of audit; amount of FISM funds received; population aged above 18; lagged incumbent party vote share; average years of schooling; share of the population that is illiterate; average occupants per room; average number of children per woman; the share of the population with electricity, water, and drainage in their home; the working age share of the population; and the share of households with internet at home. The omitted irregularities category is Q1/Q2. All observations are weighted by the inverse probability of treatment assignment and weight each experimental segment equally. Standard errors clustered by municipality are in parentheses.* denotes $p < 0.1$, ** denotes $p < 0.05$, *** denotes $p < 0.01$ from two-sided t tests. The one-sided coefficient tests at the foot of the table test whether the effect of high saturation is larger in a given direction in high saturation than in low saturation municipalities.

Table A19: Effects of receiving Facebook ads via WhatsApp on perceptions of municipal incumbent party irregularities among panel survey respondents

	Perceived % irregularities (endline)		
	(1)	(2)	(3)
WhatsApp ad	1.667	2.314	2.350
	(1.204)	(1.529)	(1.678)
WhatsApp ad × Q3		-0.455	1.846
		(3.937)	(4.462)
WhatsApp ad × Q4		-2.221	-3.900
		(2.781)	(3.318)
Perceived % irregularities (baseline)	0.374***	0.374***	0.378***
	(0.025)	(0.025)	(0.025)
Observations	1,360	1,360	1,360
R^2	0.260	0.261	0.267
Outcome range	[0,100]	[0,100]	[0,100]
Control outcome mean	58.89	58.89	58.89
Control outcome std. dev.	24.53	24.53	24.53
Interactive covariates			✓

Notes: Each specification is estimated using OLS, and includes municipality fixed effects. Specifications including interactive covariates further include interactions between treatment conditions and the following municipal-level covariates: year of audit; amount of FISM funds received; population aged above 18; lagged incumbent party vote share; average years of schooling; share of the population that is illiterate; average occupants per room; average number of children per woman; the share of the population with electricity, water, and drainage in their home; the working age share of the population; and the share of households with internet at home. All observations are weighted by the inverse probability of treatment assignment. Robust standard errors are in parentheses. * denotes $p < 0.1$, ** denotes $p < 0.05$, *** denotes $p < 0.01$ from two-sided t tests.