

The Ethics of Electoral Experiments: Design Based Recommendations Supporting Information

Tara Slough*

March 31, 2021

Contents

A1 Vote Aggregation: A Framework	A-2
A2 Survey of Electoral Experiment Pre-Analysis Plans	A-3
A2.1 Justification	A-3
A2.2 Data Collection, Sample	A-3
A2.3 Intervention Classificaton	A-4
A2.4 Ethics Coverage of Pre-Analysis Plans	A-7
A3 Applications of Decision Rule from Existing Literature	A-9
A3.1 Application #1: Boas, Hidalgo, and Melo (2019)	A-9
A3.2 Application #2: Gerber and Green (2000)	A-12
A3.3 Application #3: Bond et al. (2012)	A-16
A3.4 Application #4: López-Moctezuma et al. (2021)	A-18
A4 Existing Information Experiments	A-19
A5 Supporting Information for Empirical Illustration	A-23
A5.1 Data and Data Sources	A-23
A5.2 Mapping the Framework onto Data	A-23
A5.3 Prediction Method	A-23

*Assistant Professor, New York University. taraslough@nyu.edu

A1 Vote Aggregation: A Framework

Consider the case of an K -candidate (or K -choice) election in which n_d registered voters choose from candidates $i \in \{1, 2, \dots, K + 1\}$ where abstention is denoted by $k + 1$. Vote totals absent the intervention are denoted v_d^i where $\sum v_d^i = n_d$. Without loss of generality, assume that:

1. Candidates 1 and 2 are the *ex-ante* marginal candidates.
2. $v_d^1 > v_d^2$ such that candidate 1 would win the last/only seat contested in the absence of intervention.¹

From the definition of the margin to pivotality, $\psi_d, \psi_d n_d \equiv v_d^1 - v_d^2$.

In response to an intervention, denote the *net* change in votes from party r to party s as Δ_{rs}^d where $r < s$. If $\Delta_{rs} > 0 (< 0)$, candidate r received more (less) votes from candidate s voters than candidate s received from candidate r voters. The post-intervention vote total for party i , \tilde{v}_d^i , can thus be calculated:

$$\tilde{v}_d^i = v_d^i + \sum_{r=i} \Delta_d^{rs} - \sum_{s=i} \Delta_d^{rs} \quad (1)$$

The difference in votes between candidate 1 and candidate i can thus be written:

$$\tilde{v}_d^1 - \tilde{v}_d^i = v_d^1 - v_d^i + \sum_{r=1} \Delta_d^{rs} - \left(\sum_{r=i} \Delta_d^{rs} - \sum_{s=i} \Delta_d^{rs} \right) \quad (2)$$

If the intervention does not change the election result, candidate 1 must still win the last/only seat. This implies that $\tilde{v}_d^1 > \tilde{v}_d^i$ for all $i \in \{2, \dots, n\}$.²

$$\tilde{v}_d^1 > \tilde{v}_d^i \Rightarrow v_d^1 - v_d^i > -2\Delta_d^{1i} - \sum_{r=1, s \neq i} \Delta_d^{rs} + \left(\sum_{r=i} \Delta_d^{rs} - \sum_{r \neq 1, s=i} \Delta_d^{rs} \right) \quad (3)$$

Given an interference assumption, Definition 1, and the definition of Δ_d^{rs} imply that:

$$n_d MAEI_d \geq \sum |\Delta_d^{rs}| \quad (4)$$

Equation 4 further implies that $2n_d MAEI_d \geq 2 \sum |\Delta_d^{rs}|$. It therefore follows that:

$$2n_d MAEI_d \geq 2|\Delta_d^{11}| + \sum_{r=1, s \neq i} |\Delta_d^{rs}| + \sum_{r=i} |\Delta_d^{rs}| + \sum_{r \neq 1, s=i} |\Delta_d^{rs}| \quad (5)$$

Equations 3 and 5 imply that if $v_d^1 - v_d^2 > 2n_d MAEI_d$, it must be the case that $\tilde{v}_d^1 - \tilde{v}_d^2 > 0$. Substituting $v_d^1 - v_d^2 = \psi_d n_d$, if $\psi_d > 2MAEI_d$, then the experimental intervention could not change who wins office.

¹In a PR election, it may be useful to think of v_d^i as a quotient or remainder on the last seat allocated. The logic follows equivalently.

²I assume that there is no minimal participation rule. Thus abstention (option $n + 1$) therefore cannot “win,” though this does not change the result.

A2 Survey of Electoral Experiment Pre-Analysis Plans

A2.1 Justification

To survey the scope of electoral experiments, two forms of data collection seemed possible: a survey of published articles or pre-analysis plans. Sample selection is a concern in both cases, as depicted in Figure A1. The use of pre-analysis plans admits some bias toward recent studies since electoral experiments proliferated before the practice of pre-registering most experiments (sometime between 2013 and 2015). This has the benefit of characterizing current practices without delays in publication. Further, it is useful to catalogue pre-analysis plans because we do not know the process via which the publication process selects on ethical concerns (or lack thereof). It may be that the most concerning studies are never written up or are rejected in the review process as a function of ethical concerns. Alternatively, preferences for large-scale experiments or large samples may yield a sample of published studies that overrepresents the most concerning studies.




Figure A1: Selection processes. The current dataset emphasizes all pre-registered experiments in the AEA and EGAP databases, including all studies under the red node.

I collect pre-analysis plans for electoral experiments that are registered in the Evidence in Governance and Politics (EGAP) and American Economic Association (AEA) registries. These registries are open to all and is costless for researchers (beyond the costs of preparing a pre-analysis plan). Reading registries and pre-analysis plans allows for systematic examination of how researchers addressed ethical questions when designing their experiments. This corresponds to the research stage at which the method developed in this paper can be most productively implemented. While not all experiments are pre-registered in these registries, these documents provide some evidence about how researchers made design decisions and why. Most of the cited experiments in the paper (even some of the earlier studies) are included in this dataset of pre-analysis plans.

A2.2 Data Collection, Sample

I draw my sample of studies from the two largest social science registries: EGAP and AEA. At the time of database construction, EGAP had 1,434 registrations and AEA had 4,308. To narrow these registrations, I define an electoral experiment as one that:

1. Intervenes in a “real” election outside of a laboratory setting. This excludes a large number of survey experiments (typically vignettes) and lab-based studies of voting.

2. The researcher has some control over the random assignment of the intervention. This excludes natural experiments.

I consider downstream experiments based on the classification of the treatment. When authors look for downstream consequences of an ostensibly non-election related intervention (i.e., education interventions or a development programs), I exclude them from the set of pre-analysis plans. If authors look at downstream outcomes of an election-related intervention I classify this as an elections experiment.

I describe the process through which the sample was sequentially narrowed in Table A1.

Step	EGAP		AEA	
	Criteria	n	Criteria	n
1	Complete database	1,434	Complete Database	4,308
2	Pre-2015: use of word “election” in database fields; post-2015 use of “election” keyword and “field experiment” methodology	209	Use of “electoral” keyword	181
3	Apply definition of electoral experiments (stated above)	77	Apply definition of electoral experiments (stated above)	57

Table A1: Selection of pre-analysis plans for inclusion from both databases.

For five relevant studies, authors registered the studies in both registries. I remove the duplicates from the dataset for a total of 129 pre-registered electoral experiments. All duplicates are counted in the EGAP category from hereon.

Figures A2 and A3 summarize the coverage of the data, both in terms of time and the countries where these experiments are conducted. While researchers pre-register experiments in US elections at a higher rate than in any other country, the majority of experiments in the dataset are conducted outside the US. These AEA and EGAP registries became much more commonly used between 2015 and 2020 as pre-registration norms were popularized (Figure A3). Note that the AEA registry, in particular, contains some pre-analysis plans for experiments that were fielded *before* the creation of the registry, between 1998 and 2013. One further note is that a number of studies fielded in 2020 are currently embargoed and are thus not present in the dataset. These embargoes, in addition to the pandemic-related constraints on field research, may account for the slower pace of registrations in 2020.

A2.3 Intervention Classification

Figure 1 plots the frequency of different classes of interventions in electoral experiments. Table A2 elaborates the definition of each class and disaggregates the use of these interventions in the US and outside the US.

Beyond the US/non-US distinction, Figure A4 depicts the differences in the distribution of treatments in upper- versus low-/middle-income democracies. We see substantial differences in the rates of application of most intervention types as a function of country wealth.




Figure A2: The distribution of experiments, by country, in the pre-analysis plans dataset.




Figure A3: Distribution of the the start dates of electoral experiments in the pre-analysis plans dataset (left). Note that AEA encouraged the submission of old pre-analysis plans after the repository was created, accounting for many of the pre-2013 studies (right). In this plot “A” stands for AEA and “E” stands for EGAP

Class	Description	Published Example	% of experiments	
			in US	outside US
Mobilization	Get out the vote (GOTV) interventions. These interventions may use information to get out the vote, but they are aimed at mobilizing turnout, not changing votes or voter opinions of candidates/ballot items.	Gerber and Green (1999)	61.8%	23.0%
Information	Treatment is some form of information or access to information about candidates on the ballot or ballot items (as in referenda). Information may be prospective or retrospective.	Dunning et al. (2019), Bidwell, Casey, and Glennerster (2020)	14.5%	36.5%
Persuasion	Treatments intended to change votes or behavior in one specific direction. These experiments frequently randomize campaign strategy.	Wantchekon (2003), Kendall, Nannicini, and Trebbi (2015)	14.5%	18.9%
Election administration	Treatments intended to promote clean election administration or deter fraud.	Asunka et al. (2017)	–	10.8%
Voter registration	Treatments intended to induce eligible voters who are not registered to register to vote.	Shineman (2020)	5.5%	1.4%
Candidate-level	Treatments targeted at candidates or prospective candidates running for office.	Casey, Kamara, and Meriggi (2021)	–	5.4%%
Campaign finance	Treatments intended to shape financial contributions by voters	Boatright, Green, and Malbin (2006)	3.6%	–
Other	Other treatments deployed in advance of elections.	Jha and Shayo (2019)	–	4.1%

Table A2: Classification of experiments in the pre-analysis plan dataset.

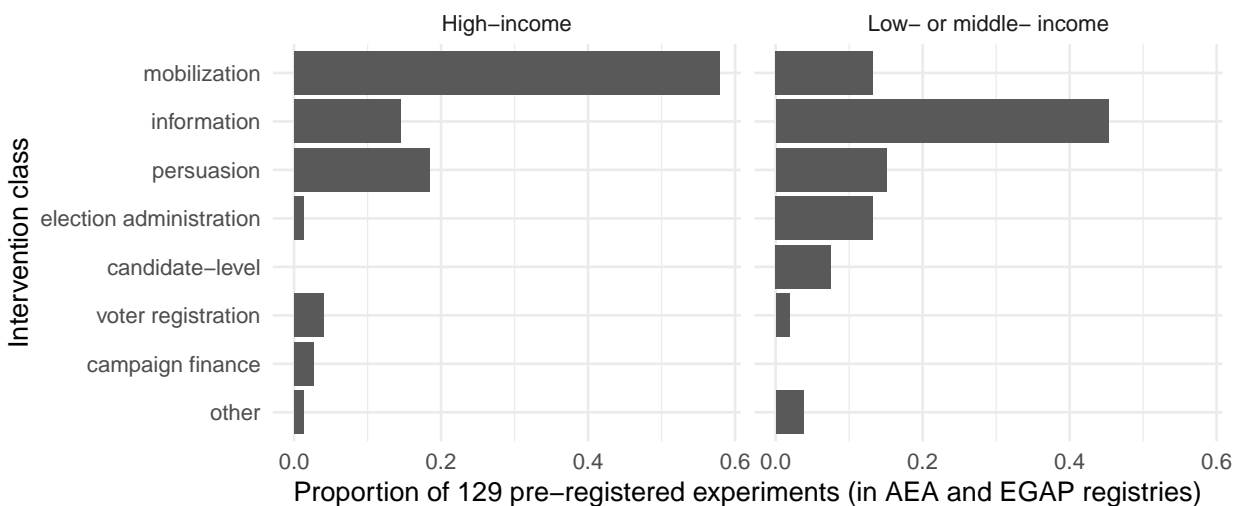


Figure A4: Distribution of treatments by wealth of “host” country.

Table A3 describes the elections in which the pre-registered experiments have been set. Note that some interventions, like mobilization interventions, do not generally correspond to a specific office. If an intervention increases turnout in a given election, it presumably increases turnout for all offices contested. These elections are coded as “non-specific.” In contrast, other interventions like information dissemination and persuasion campaigns are more likely targeted at a single race. The targeted offices are the rows in Table A3. Among studies targeted to specific offices, the most frequently targeted offices are: (1) national legislatures; (2) local executives (mayors); and (3) national executives.

A2.4 Ethics Coverage of Pre-Analysis Plans

Finally, I consider the coverage of ethical considerations in this subset of pre-analysis plans. In general, I consider three types of ethical considerations as they appear in pre-analysis plans or registry entries:

1. **IRB approval:** Does a registry entry or pre-analysis plan reports having IRB approval from at least one institution? Note that I do not assume that failure to include an IRB number does not necessarily mean that a project did not have IRB approval. For the AEA registry, this is a required entry, so the absence of an associated IRB protocol may be more indicative of a lack of approval in that dataset.
2. **Discussion of any additional ethical considerations:** Does the pre-analysis plan or registry discuss ethical concerns or risks beyond an IRB approval?
3. **Consideration of aggregate electoral impact:** Does the pre-analysis plan or registry discuss the possibility of changing election outcomes as a potential ethical concern?

For the EGAP registry, I reviewed pre-analysis plans for all registered experiments. For the AEA registry, I reviewed all public pre-analysis plans and registry entries (not all pre-analysis plans are public). For this reason, I disaggregate by repository before reporting the overall rate of inclusion of these items in Table A4.

Targeted office	National elections		State elections		Local elections		TOTAL
	General	Primary	General	Primary	General	Special	
Non-specific	43	1	5		9		58
National executive	7						7
National legislature	21	3				2	26
National referendum	3						3
Supranational legislature	3						3
State executive	3		2				5
State legislature	1		4				5
State judiciary				1			1
Local executive (mayor)					9		9
Local council					5		5
Local referendum					1		1
TOTAL	81	4	11	1	24	2	123

Table A3: Number of studies according to each office (where relevant) and election type. Six studies were could not be coded on at least one dimension.

Item	AEA registry (<i>n</i> = 52)	EGAP registry (<i>n</i> = 77)	Combined registries (<i>n</i> = 129)
IRB approval	73.1%	41.6%	54.6%
Additional ethical considerations	0%	10.4%	6.2%
Aggregate electoral impact	0%	2.6%	1.6%

Table A4: Ethical coverage of pre-analysis plans and registry entries.

A3 Applications of Decision Rule from Existing Literature

I walk through the application of the proposed framework and decision rule to four experiments documented in the existing literature. The goal is to illustrate four aspects of the current framework: individual vs. clustered random assignment and intervention without vs. with a partner (Case #1 vs. Case #2 of Table 3). The studies I employ are listed in Table A5.

Assignment is:	Without experiment:	
	No intervention occurs	Partner conducts intervention
Individually randomized	Boas, Hidalgo, and Melo (2019)	Bond et al. (2012)
Cluster randomized	Gerber and Green (2000)	López-Moctezuma et al. (2021)

Table A5: Studies used for illustration of framework, decision rule.

In addition to calculation of the MAEI and (to the best of my ability) ψ_d , I provide ex-post data on election results to highlight the relationship between *ex-ante* and *ex-post* voting.

A3.1 Application #1: Boas, Hidalgo, and Melo (2019)

Boas, Hidalgo, and Melo (2019) conduct a field experiment in 46 municipalities in Pernambuco, Brazil prior to the 2016 election. This experiment was one of seven experiments in EGAP’s Metaketa-I on electoral accountability. It is notable because of the considerations of electoral impact undertaken in the pre-analysis plan.

1. **Electoral context:** This experiment provided information on municipal government performance to voters in 46 municipalities in the Brazilian state of Pernambuco prior to the 2016 municipal elections. Municipal elections include elections for mayor (FPTP races for all municipalities in the sample) and city council (open list PR races).
2. **Experimental overview:** The experiment consists of 3,200 voters, of whom 2,400 are assigned to two treatment conditions that reveal information about (1) municipal financial management and (2) education. The allocation of these 3,200 voters across 46 municipalities is depicted in A5. The experiment randomizes individual respondents to these treatment arms.
3. **Calculating $MAEI_d$:** This experiment presents a very straightforward calculation of the $MAEI_d$. Begin by assuming SUTVA (Assumption 1). Because the intervention would not happen in the absence of the experiment, we know that $|S_{01}^{cd}| = 0$. (In other words, no voters are being denied the intervention because of the experiment.) In an individually randomized experiment, $n_c = 1 \forall c$. This implies that $a_i(0) \in \{0, 1\} \forall c$ since each voter either votes for the marginal winner ($a_i(0) = 1$) or does not ($a_i(0) = 0$). Further, we know that $\sum_{c \in d} |S_{10}^{cd}|$ is simply the number of treated voters in each district, which is depicted by the two darker bar segments in Figure A5. Further, we know that n_d is the number of registered voters in each municipality. In the Brazilian case, this data is easily accessible from the Brazilian Tribunal Superior Eleitoral (TSE).

The results of the $MAEI_d$ calculation for each district under SUTVA (potentially non-conservative) and assuming $E[a_i(0) = 0]$ for all treated voters (equivalent to $E[a_i(0) = 1]$ and conservative) is

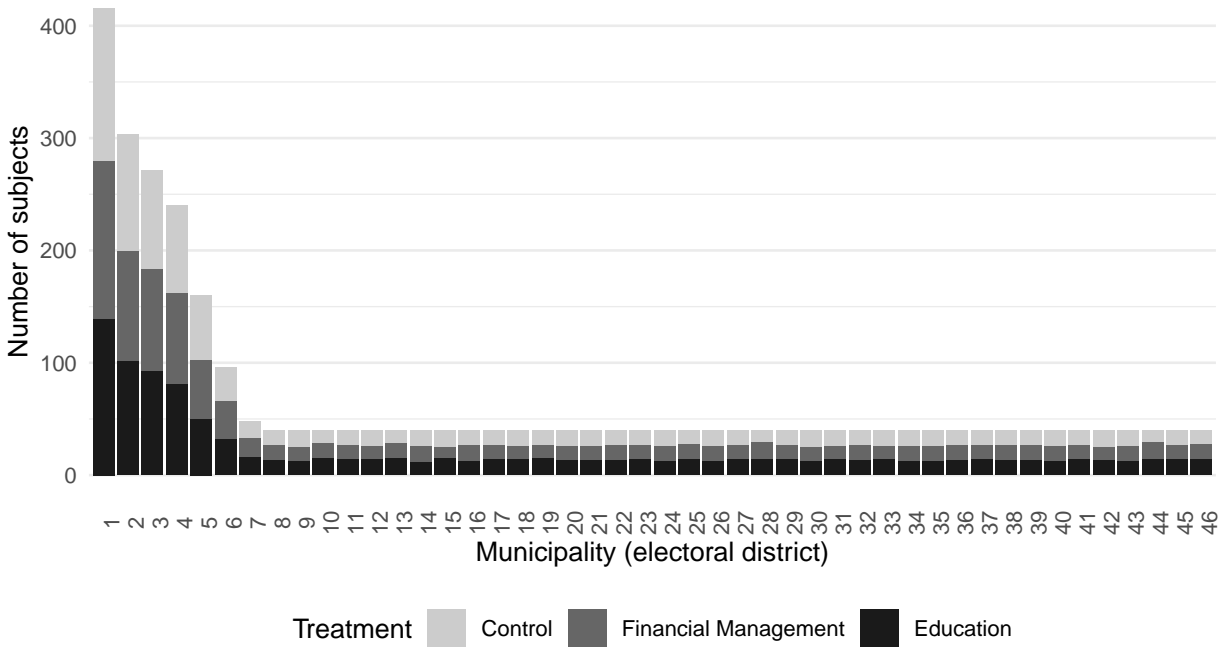


Figure A5: Sample allocation across districts/constituencies in Boas, Hidalgo, and Melo (2019).

graphed in Figure A6. Under these assumptions, the $MAEI_d$ collapses to the number of treated voters over the number of registered voters.

Figure A7 depicts how the $MAEI_d$ changes under less conservative assumptions about $E[a_i(0)]$ (left) and more conservative assumptions about interference in the form of modest (relative to treatment density) between-cluster spillovers (right). Given that the PAP does not anticipate spillovers or interference, I use the otherwise conservative estimates of $MAEI_d$ from Figure A6 for the following analysis.

4. **Calculating ψ_d :** I am unaware of systematic pre-election polling across municipalities leading up to the Brazilian elections of 2016. To demonstrate the approach, I rely simply on predicting the mayoral margin of victory of the 2016 election from the margin of victory in the preceding 2012 election following the method in Appendix A5.3 with a single predictor: 2012 margin of victory. I fit the data on changes in margin of victory from 2004-2008 and 2008-2012 in the state of Pernambuco (both inside and outside the experimental municipalities). Figure A8 depicts the resultant estimates of ψ_d on the y -axis against the $2MAEI_d$ on the x -axis. Points (municipalities) below the 45° line do not pass the decision rule. Points (municipalities) above the 45° line do pass the rule. Using these estimates, 11/46 municipalities pass the decision rule.

I encourage caution in interpreting Figure A8. A better predictive model would, in theory, be less conservative. The current estimates of ψ_d are penalized by the weak predictive model ($R^2 = 0.044$). The model is weakly predictive because of limited autocorrelation in margins of victory ($\sigma = 0.21$).

5. **Ex-post assessment** As depicted in Figure A9, I find that one of the mayoral races fell within $MAEI_d$. We certainly cannot ascertain from this finding that the treatment changed the election

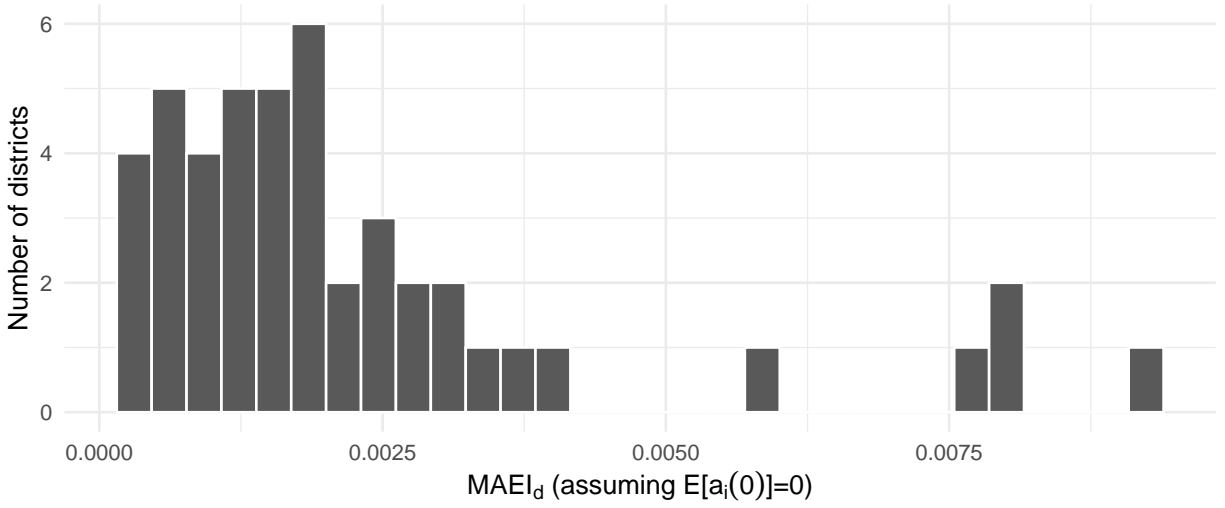


Figure A6: $MAEI_d$ for the 46 districts in Boas, Hidalgo, and Melo (2019) under SUTVA and an assumption that $E[a_i(0) = 0]$ for all treated voters.




Figure A7: $MAEI_d$ for the 46 districts in Boas, Hidalgo, and Melo (2019) varying $E[a_i(0)]$ (left) and allowing for between-cluster spillovers of magnitude $\pi_c = 1$ (right).




Figure A8: Visualization of the decision rule across the 46 experimental municipalities. Points below the 45° line do not pass the decision rule. Points above the 45° line do pass the rule.

outcome. This election was decided by five votes. Nevertheless, it is useful to think about how the $MAEI_d$ compares to the ex-post distribution of election results. One further note is that if we were to consider municipal council races, the margin between the two marginal candidates is certainly closer than the distribution of the ψ_d in mayoral races.

A3.2 Application #2: Gerber and Green (2000)

Gerber and Green (2000) is a seminal study of voter mobilization that was conducted in New Haven, CT in the 1998 midterm elections.

1. **Electoral context:** Gerber and Green (2000) report the results of a voter mobilization campaign in New Haven, CT conducted during the November 3, 1998 midterm election. In this election, New Haven voters voted for several statewide offices (Governor, Secretary of State, Treasurer, Comptroller, Attorney General, and US Senate); US House (in district CT-3); state legislative seats (two state Senate races and seven state House races); one county race (New Haven County Sheriff); and one city race (for probate court judge). As a city, New Haven voters predominantly supported Democratic in 1998. Several races were uncontested including 1/2 state Senate races; 3/7 state House races; and the probate court race.
2. **Experimental overview** Gerber and Green (2000) randomly assigned voters in New Haven, Connecticut to various get out the vote treatments prior to the November 1998 (midterm) election. According to the Connecticut Secretary of State, there were 52,761 registered voters in New Haven as of Novem-

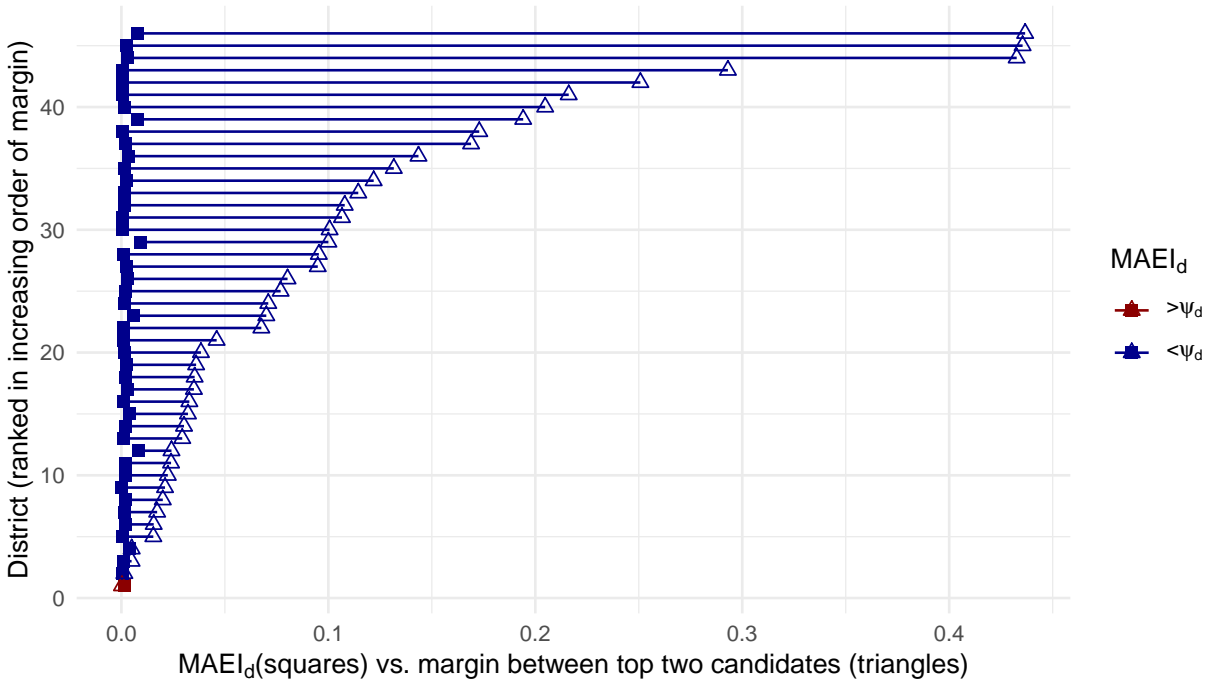


Figure A9: Comparison of $MAEI_d$ projections to the outcome of mayoral races in 2016. Here ψ_d corresponds to the actual (not the predicted margin of victory.)

ber 3, 1998.³ The researchers reduced the experimental sample by excluding individuals with PO boxes, households with more than two registered voters, and a university ward, for a total of 30,544 voters residing in 23,008 households.⁴

Households were cluster-randomized to sixteen treatment arms that varied: (1) the number of direct mailings sent (0-3 mailings); (2) in-person canvassing (yes/no); and telephone canvassing (yes/no). Treatments were administered by researchers and graduate student assistants. In the absence of this research, voters would not have been contacted by mail, phone, or in-person canvass *by the resesarch team for the purpose of research*. This does not preclude other interventions. I assume that the control group represents the “status quo” at least respect to the absence of interventions by the researchers. As such, I consider a control group of 10,800 voters and a treatment group of 18,580 voters.

3. **Calculating the $MAEI_d$:** The concurrent elections with different districts complicates the calculation of the $MAEI_d$. Consider two types of districts:

- Districts that contain all experimental households. This includes all statewide races, the US House race (to the best of knowledge), and the city and county races. For the purposes of this analysis, I consider the largest and smallest electorates of these districts: the statewide ($n_d = 1,806,750$) and city-wide electorate ($n_d = 52,761$).
- Districts that contain a subset of experimental households. Replication does not include the mapping to the relevant State Senate and State House districts (likely for privacy reasons). Nor is public registration compiled at this level. I roughly approximate the number of voters in State House and Senate districts as the number of registered voters statewide divided by the number of districts statewide. I consider the smaller districts – those for State House – with approximately $n_d = \frac{1,806,750}{133} = 13,585$ voters.

Given that clusters (households) were fully saturated (all voters were considered to be assigned to treatment), Assumption 2 is subsumed by SUTVA (Assumption 1). Given that this research was implemented by the researchers, $|S_{01}^{cd}| = 0$. In other words, no voters are being denied the intervention because of the experiment. We can construct conservative bounds (under SUTVA) by assuming $E[a_i(0)] = 0 \forall c$. In this case, the $MAEI_d$ is simply the number of treated voters as a fraction of the registered voters in the district $\frac{\sum_{c \in d} |S_{10}^{cd}|}{n_d}$. This calculation is straightforward when the experiment is entirely contained within a district (first case above).

When the experiment is not fully contained within a district, we cannot ascertain the $MAEI_d$ without further information relating experimental subjects to districts. We can, however, bound the plausible range of $MAEI_d$'s under the same assumptions. First, to bound the maximum, imagine that all subjects of a district were experimental subjects. Here, given random assignment of treatment, the proportion of voters assigned to treatment should be equivalent (in expectation) to the probability of assignment to (any) treatment. At a minimum, there may be no experimental subjects in a district. New Haven had 52,761 registered voters. It is plausible that any district smaller than the difference between the number of registered voters and the number of experimental subjects or 52,761-30,544

³Data retrieved from <https://portal.ct.gov/SOTS/Election-Services/Statistics-and-Data/Statistics-and-Data>.

⁴These numbers come from a revised household-level replication dataset published by the authors in 2005. They are slightly higher than those reported in Gerber and Green (2000). No substantive conclusions below are affected by differences in these numbers.

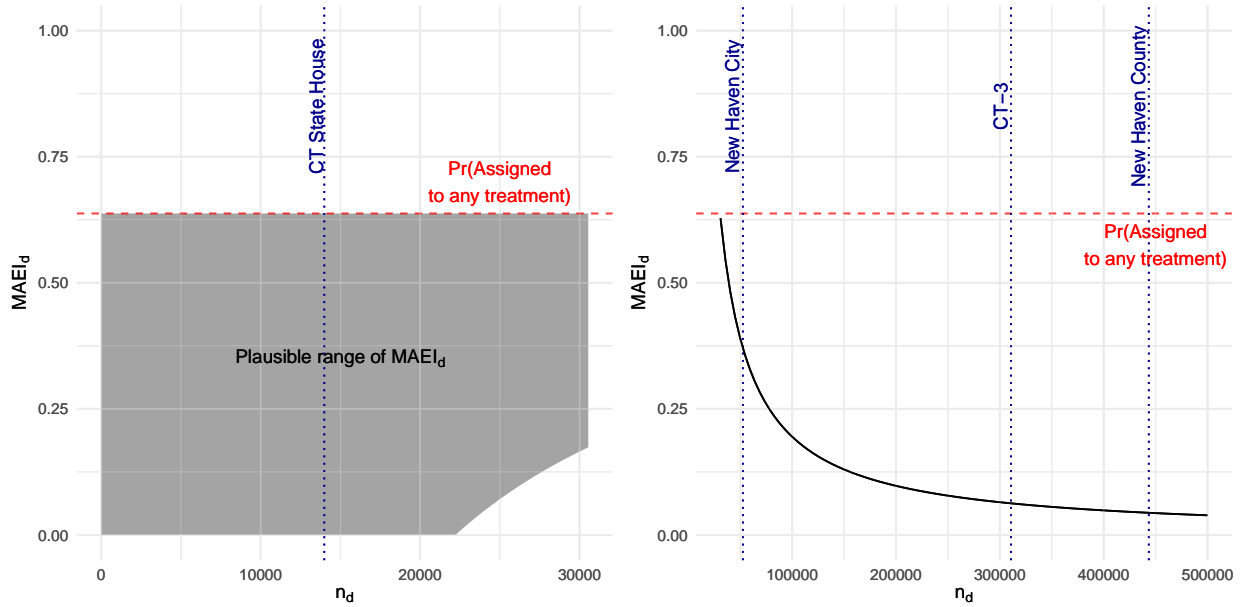


Figure A10: The left panel gives the bounds on the $MAEI_d$ for districts that do not contain all experimental households. The right panel gives the $MAEI_d$ as a function of n_d . Note that both plots assume SUTVA and that $E[a_i(0)] = 0$.

= 22,217 registered voters could be completely excluded from the experiment. Above that threshold, we can derive a minimum bound on the plausible $MAEI_d$ based on the probability of assignment to treatment and the number of registered voters in a district. Figure A10 proves a visualization. Note that these bounds are exceedingly wide in the absence of information about the relationship between voters and districts.

4. **Predicting ψ_d :** I am unable to predict many of the down-ballot races. However polls compiled by the Cook Political Report allow for prediction of ψ_d for several state-level races. I report predictions from the last poll prior to the election in Table A6.

Race	Favored party	Survey	n	$\underline{\psi}_d$
Governor	R	Mason-Dixon PMR (Oct 9-12, 1998)	629	0.24
Senator	D	Mason-Dixon PMR (Oct 9-12, 1998)	629	0.20

Table A6: Estimates of $\underline{\psi}_d$ for two races for which polling data remains available.

Two notes are necessary. First, these estimates represent margin as a share of voters, not as a share of registered voters. Recall that ψ_d is constructed as a ratio to the number of registered voters. I normalize the rates in Table A6 by the 1997 elections (42.9%). Because 1997 was an off-cycle election, this is likely a conservative estimate (lower turnout than predicted in the midterm election in 1998). The points representing these $\underline{\psi}_d$ estimates in Figure A11 are equivalent to multiplying the predicted margin by this turnout rate. Second, several of the races are uncontested. In these cases, imputation of $\underline{\psi}_d = 1$ seems appropriate, particularly when there exists no turnout requirement. In Figure A11, I distinguish contested from uncontested races.

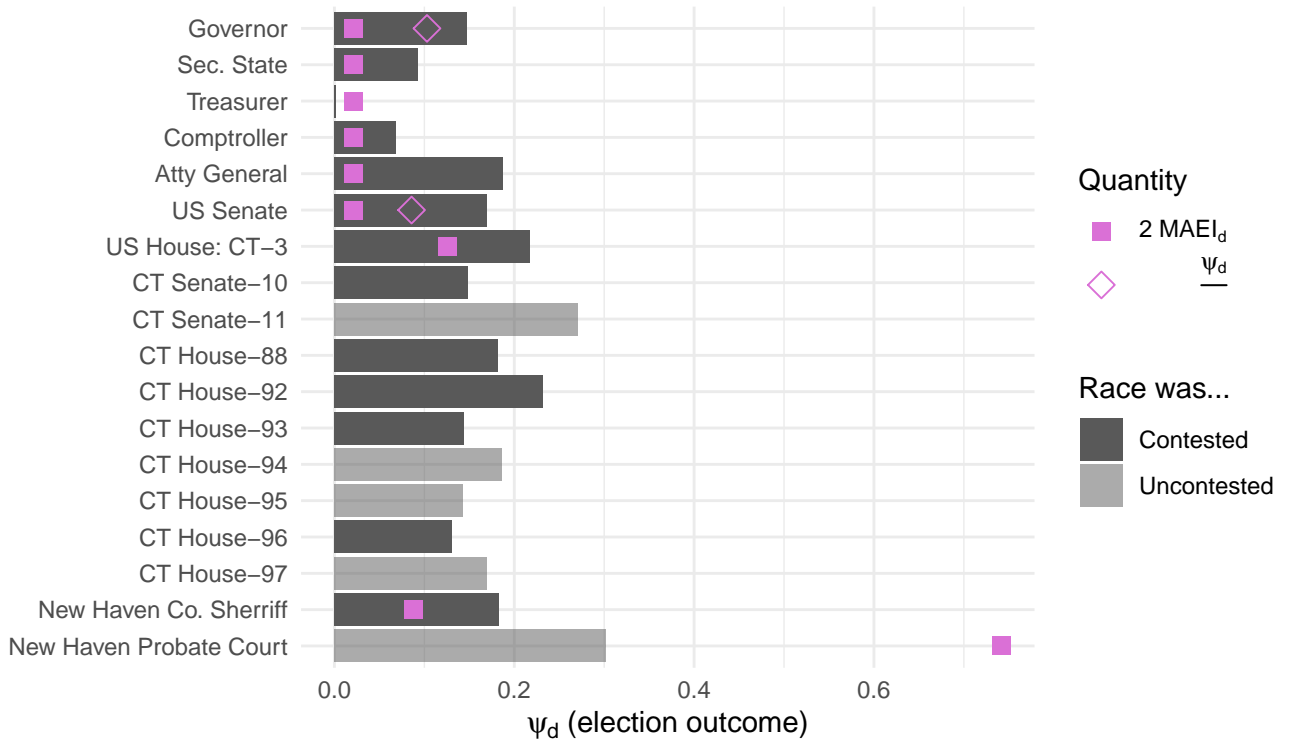


Figure A11: Comparison of estimable $MAEI_d$'s, $\underline{\psi}_d$'s, and electoral outcomes (ψ_d). See #3 and #4 for the issues that preclude estimation of some $MAEI_d$'s and $\underline{\psi}_d$'s.

- Ex-Post Assessment:** Figure A11 compares (1) the election results (stated in terms of ψ_d , the margin of victory), by race; (2) the point estimable $MAEI_d$'s; and (3) the estimable $\underline{\psi}_d$'s. Two observations of note. First, both races for which we have all three quantities pass the decision rule. $2MAEI_d < \underline{\psi}_d$. Further, in the city race, the race was uncontested. Here $2MAEI_d$ presents less of a concern given the uncontested nature of the race. In contrast, the attorney general's race was very closes (decided by a mere 2,684 votes, or $\psi_d = 0.003$). This means that the intervention *could* have changed the outcome, though we can never know way.⁵

A3.3 Application #3: Bond et al. (2012)

Bond et al. (2012) is a very large-scale field experiment conducted in collaboration with Facebook during the 2010 midterm elections.

- Electoral Context:** In 2010, US midterm elections included election of 35 Senators, the entire House of Representatives, 34 Governors, and countless state and local officials. At the national and state levels, Republicans picked up many seats.
- Experimental Overview:** This experiment was presumably conducted at national scale. There are some vagaries due to the proprietary Facebook data. It included three treatment arms:

⁵Given the modest estimates on turnout, it seems highly unlikely that this was the case in the absent of large, unmeasured persuasion effects.

- A *social message group* in which users saw a banner with a clickable “I voted” button at the top of their “Newsfeeds” ($n = 60,055,176$). This is consistent with Facebook’s banner that a large number of users saw in 2008.
- An *informational message group* without the button on their “Newsfeed” ($n = 611,044$).
- A *control group* in which users did not see anything on the top of their “Newsfeed” ($n = 613,096$).

An author interview with a researcher on this study indicates that in the absence of the experiment, all subjects would have seen the *social message group*. This implies that $\sum_{c \in d} |S_{10}^{cd} \cup S_{01}^{cd}| = 1,224,140$ (the sum of the subjects in the control and informational treatment arms). Obviously, this is a very small proportion ($\approx 2\%$) of the experimental sample.

The US Census Bureau estimates that there were 137,263,000 registered voters in 2010.⁶ As such, the experiment included 43.8% of registered voters with 0.9% assigned to one of the treatment conditions that existed *because of the experiment*.

3. **Calculating $MAEI_d$:** There are two important sources of ambiguity relevant to calculation of the $MAEI_d$. First, we do not know with certainty whether Facebook conducted the experiment across all possible US voters. The rate of treatment saturation is consistent with (outside Facebook) estimates of adult Facebook usage in 2010, so it is plausible that all states were included. Below, I make the simplifying assumption that Facebook intervened in all states/districts. Second, I do not know the distribution of subjects – or Facebook users – across the United States. To my knowledge, this data is not available. For the purposes of exposition, suppose that the proportion of Facebook users as a share of registered voters in an arbitrary district, d , is $FB_d \in (0, 1]$. Given the above calculation, the population-weighted average of FB_d is ≈ 0.438 .

One aspect of treatment assignment is straightforward. The experiment did not use blocked random assignment and maintained constant probabilities of assignment to treatment across the experimental sample.⁷ As above, the probability of assignment to treatment implies that $E[\frac{\sum_{c \in d} |S_{01}^{cd} \cup S_{10}^{cd}|}{n_d / FB_d}] = \frac{1,224,140 \times FB_d}{60,055,176} = 0.02 \times FB_d$.

I consider three estimates of the $MAEI_d$, under different interference assumptions. First, consider the calculation with *SUTVA*. In principle, it would be useful to know $E[a_c(0)]$ for the Facebook user population in the experimental sample. This may be approximated by data on “I voted” use in the 2006 election. For the purposes of the analysis, make the conservative assumption that $E[a_c(0)] = 0 \forall c \in d$.⁸ Under *SUTVA*, therefore:

$$MAEI_d = 0.02 \times FB_d$$

Because the experiment is randomized at the individual level, Assumption 2 holds trivially when we invoke *SUTVA* (Assumption 1). Hence, I consider relaxing *SURVA*. In this case, we would need to specify π_c , the extent of spillovers. I simply consider two cases. First, consider the case in which

⁶See <https://www.census.gov/data/tables/2010/demo/voting-and-registration/voting-registration-2010-election.html>.

⁷This point was confirmed in an interview with a researcher on the Bond et al. (2012) paper.

⁸Obviously, the least conservative assumption of $E[a_c(0)] = \frac{1}{2} \forall c \in d$ would at most double $MAEI_d$ in the absence of interference.

spillovers only occur among the population of Facebook users. Second, consider the case in which spillovers might occur among all registered voters. We can then write $MAEI_d^{bw}$ as:

$$MAEI_d^{bw} = \begin{cases} FB_d(0.02 + \pi_c 0.98) & \text{if spillovers are confined to Facebook} \\ 0.02FB_d + (1 - 0.02FB_d)\pi_c & \text{if spillovers are offline} \end{cases}$$

4. **Discussion of ψ_d and experimental design:** There were thousands of races contested within the Facebook experiment. I do not quantify or predict ψ_d or assess the predictions *ex-post* given the number of races involved in this massive nation-wide experiment. However, note that ultimately only $\frac{1,224,140}{137,263,000}$ or 0.89% of the national electorate was assigned to a treatment condition *because* of the experiment. At least for congressional races, these experimental conditions could have been re-allocated to districts with safe House and Senate races with little risk of changing outcomes. Note however, that this would eliminate common support in any swing districts, reducing the size of experimental sample. This occurs because districts with no probability of assignment to the experimental intervention arms (information or no banner) would be non-experimental.

A3.4 Application #4: López-Moctezuma et al. (2021)

1. **Electoral Context:** The 1987 Filipino Constitution mandates that 20% of the House of Representatives (lower chamber) be composed by members of marginalized groups. Since 1995, this representation mandate has been fulfilled using closed-list proportional representation within a single national district to fill 58 seats of the 289-seat chamber. Seats are apportioned according to a Hare quota with a 2% threshold. Note that major parties are prohibited from running in these races that seek to provide representation to marginalized groups. There are many party lists that compete for the 58 seats. In the 2013 election studied by López-Moctezuma et al. (2021), 122 parties contested the party list contest. There were 52,006,910 registered voters in this election.⁹

Note that both of the parties with whom researchers collaborated were relatively small. However, improving the prospects of a small party can remove votes from a marginal seat winner, changing election outcomes. As such, I employ symmetric considerations regardless of party size.

2. **Experimental overview:** López-Moctezuma et al. (2021) randomize the campaign strategy of two party lists. The collaborations with each party list were conducted in different barangays (neighborhoods/villages). Treatment was assigned at the level of the barangay (the cluster) and the treatment is the use of deliberative town-hall meetings whereas control is “business as usual” campaigning. Per an interview with a researcher on this project, in the absence of the intervention, parties would have campaigned in treatment barangays in the “business as usual” fashion. Here, we can assume that there are no “if non-experiment assigned” communities, i.e. $|S_{01}| = 0$.

One point worth discussion is that there was two-sided non-compliance in the administration of this intervention. Two barangays assigned to treatment were untreated and two barangays assigned to control were ultimately treated. The two-sided non-compliance occurred because researchers sought to assign additional barangays to treatment after the two treatment barangays did not participate in the debates. As such, two (very similar) quantities are relevant. While planning, the number of registered voters in treatment barangays would be the relevant measure of $|S_{10}| = 98,547$. During

⁹National electoral outcome data comes from https://en.wikipedia.org/wiki/2013_Philippine_general_election.

Assignment	Treatments delivered		$ S_{10} $ (as assigned)
	“Business as usual” (Control)	Deliberative town halls (Treatment)	
Control	24 barangays 178,635 RV	2 barangays 15,141 RV	98,948 RV
Treatment	2 barangays 15,542 RV	11 barangays 83,406 RV	
$ S_{10} $ (as delivered)	98,547 RV		

Table A7: Non-compliance with treatment assignment. RV stands for registered voters.

the administration, that number shifted to $|S_{10}| = 98,948$ when the replacement barangays were assigned. Fortunately these numbers are very similar.

3. **Calculating $MAEI_d$:** Here, there is a single district in which $n_d = 52,006,910$. Both parties were both small: in the preceding election, the larger party won 3% of the vote (as a share of valid votes) and the other party won 0% of the vote. We can approximate $E[a_c(0)] = 0$, without a large loss of accuracy. We know from above that $|S_{10} \cup S_{01}| = S_{10}$ given the definition of treatment. As such, $MAEI_d = \frac{|S_{10}|}{n_d}$. With either measure of $MAEI_d$, this simplifies to approximately 0.0019.
4. **Predicting ψ_d :** In a proportional representation race with 58 seats and >100 party lists, it is very difficult to predict ψ_d . Analytically, we know that the upper bound on ψ_d is $\frac{1}{58} = 0.017$, which is substantially larger than the $2MAEI_d$ calculated above. However, this is a strict upper bound. I lack the tools to predict this quantity, though it is evident that even if $\psi_d \sim U[0, \frac{1}{58}]$ an admittedly very rough approximation, $F_{\psi_d}(2MAEI_d) = 0.22$, outside the decision rule. There may be much better ways to calculate ψ_d with more electoral information.
5. **Ex-Post Assessment:** The realized $\psi_d = 0.0002$ was calculated from the marginal party earning 1 seat (Append) and the marginal party earning 0 seats (Alif). This margin is smaller than between the marginal party winning 3 seats and the marginal party winning 2 seats as well as the marginal party winning 2 seats and the marginal party winning 1 seat. This is clearly smaller than $2MAEI_d$. Note that of course that this does not indicate that the experiment changed the aggregate outcome, only that it had the potential to.

A4 Existing Information Experiments

I focus on published experiments on the provision of incumbent performance information to voters before elections, adapting the list of studies from Enríquez et al. (2019). Note that all calculations are back-of-the-envelope. I cannot estimate $E[a_{jc}(0)]$ in the case of cluster-randomized experiments. For this reason, I show the full range of $MAEI_d$ over the possible domain of $E[a_{jc}(0)] \in [0.5, 1]$.

Table A8 describes studies in the framework described in this paper.

Article	Country	Mapping to Framework	Calculation Details	$MAEI_d$ Est.	$ D $
Adida et al. (2017)	Benin	d : Commune* c : Village (or urban quarters) j : Individual	Treatment (five variants) was assigned to 195 of 1498 villages. The saturation (density) of treatment in a village varies by treatment arm (below 100% in all villages) and village population is unclear without data on cluster size. Because of the noted distinction in the <i>de-jure</i> vs. <i>de-facto</i> characterization of parliamentary electoral districts, more information needed to clarify n_d .	–	30
Arias et al. (2019)	Mexico	d : Municipality c : Precinct j : Individual	At most $\frac{1}{3}$ of precincts per municipality were sampled, albeit non-randomly. Treatment was assigned to 200 households in each of 400 precincts (T1-T4). Precincts contain a maximum of 1,750 registered voters. Non-random sampling of precincts prevents calculation of $MAEI_d$.	–	26
Banerjee et al. (2011)	India	d : State leg. district c : Polling station j : Individual	20 treated polling stations and average of 57.5 control polling stations per district. All households were treated.	[0.129, 0.258]	10
Bhandari, Larreguy, and Marshall (2020)	Senegal	d : Department c : Village j : Individual	9 individuals were sampled per village. 450 villages were (non-randomly) sampled from the 859 villages in the 5 experimental departments. 375 villages received some treatment (non pure-control). Without further information on the distribution of villages (experimental and non-experimental) and population by district, the $MAEI_d$ cannot be calculated.	–	5
?	Brazil	d : Municipality c : Individual j : Individual	I assume $\frac{2}{3}$ of experimental sample was assigned to treatment (T1 or T2). The most over-sampled municipality had 416 voters in experimental sample and a population (not registered voters) of 45,503. If 70% of population were registered (mandatory in Brazil), upper bound (for any district) is given by $\frac{2}{3} \cdot \frac{416}{.7 \times 45,503}$.	0.008	47

Article	Country	Mapping to Framework	Calculation Details	$MAEI_d$ Est.	$ D $
Buntaine et al. (2018)	Uganda	d : District c : Individual j : Individual	Study includes 16,083 subjects (T or placebo) in 111 districts. The subjects per district and registered voters per district are not provided so $MAEI_d$ cannot be calculated.	–	111
Chong et al. (2015)	Mexico	d : Municipality c : Precinct j : <i>Individual</i>	450 of 2360 precincts were treated (selected randomly). No information is provided on saturation within precinct so I assume all voters were treated.	[0.095, 0.191]	12
Cruz et al. (2019)	Philippines	d : Municipality c : Village j : Individual	All households in 104 treatment villages (T1 or T2) across 7 municipalities were visited. Each municipality has “20-25 villages.” I assume 25 villages/municipality and that the experimental villages were randomly sampled.	[0.279, 0.594]	7
Cruz, Keefer, and Labonne (2018)	Philippines	d : Municipality c : Village j : Individual	All households in 142 treatment villages in 12 municipalities were visited. The average number of villages/municipality not reported. I assume 25 villages/municipality per Cruz et al. (2019) (which is consistent with 284 villages in the experimental sample). Villages were randomly sampled from the municipality.	[0.237, 0.473]	12
de Figueiredo, Hidalgo, and Kasahara (2011)	Brazil	d : Municipality c : Precinct j : Individual	≈ All households were visited with flyers in 200 treatment (T1 or T2) precincts of 1,759 precincts in the municipality. The precincts were selected randomly subject to a set of constraints.	[0.057, 0.114]	1
George, Gupta, and Neggers (2018)	India	d : Assembly constituency c : Village j : Individual	The intervention treated 500,000 voters (T1-T4) in 1,591 villages. Villages have ≈ 1,200 registered voters, so saturation rate in treatment villages was averaged 26%. Non-random sampling of villages within constituencies prevents estimation of $MAEI_d$.	–	38

Article	Country	Mapping to Framework	Calculation Details	$MAEI_d$ Est.	$ D $
Humphreys and Weinstein (2012)	Uganda	d : Parliamentary constituency c : Polling station j : Individual	2 polling stations in selected constituencies and all households visited with flyers. Number of polling stations/constituency not reported so $MAEI_d$ cannot be calculated. The total number of constituencies where experiment occurred (known to be <147) is not reported.	-	-
Lieri and Holmlund (2019)	Burkina Faso	d : Village* c : Individual j : Individual	12 individuals were assigned to treatment (T or placebo) per village. Information about village population (n_d) is not reported.	$\frac{12}{n_d}$	146
Sircar and Chauchard (2019)	India	d : Assembly constituency c : Polling booth area j : Individual	16 polling booth areas per precinct assigned to treatment (T1 or T2) with $\frac{2}{3}$ of households in each polling booth area assigned to receive flyer. While selection of experimental polling booths is random, the total number of polling booths per constituency is not reported so $MAEI_d$ cannot be calculated	-	25

Table A8: Survey of experiments on information disclosure about incumbent performance. * indicates that there may be distinctions between the *de-jure* electoral system and the *de-facto* vote aggregation rule, indicating some uncertainty about how to determine the electoral district.

A5 Supporting Information for Empirical Illustration

A5.1 Data and Data Sources

I simulate different research designs on electoral data from the state of Colorado. Because statewide data is sufficient to simulate all but presidential elections (and the Electoral College renders states the first unit of aggregation in presidential elections), I randomly selected the state of Colorado. As such data comes from:

- Colorado:
 - Precinct-level electoral returns voter registration from Colorado Secretary of State <https://www.sos.state.co.us/pubs/elections/VoterRegNumbers/VoterRegNumbers.html>
 - 2018 House of Representative seat predictions from The Crosstab <https://www.thecrosstab.com/project/2018-midterms-forecast/>

District type	Year	Registered Voters		Precincts	
		Mean	Std. Dev.	Mean	Std. Dev.
State House	2018	55,472	9,489	43.55	15.67
US House	2018	505,812	61,654	404.43	89.46

Table A9: Summary statistics on State and US House districts in terms of registered voters and precincts. Note that past electoral data from 2012, 2014, and 2016 is also collected for use in prediction.

A5.2 Mapping the Framework onto Data

To clarify how the data is used, I map the parameters expressed in the paper onto variables in the data/simulation in Table A10.

A5.3 Prediction Method

While much has been invested in predicting the results of national elections (in some countries), much less effort has been invested in predicting lower-level (state- and local-level) elections and elections in developing countries. In particular, there is a general lack of public opinion polling in these races. I consider what is possible to ascertain through registration data and past electoral returns alone. I propose estimate the predictive distribution of each ψ_d , following the steps:

1. Estimate a model of the form: $y_i = f(\beta \mathbf{X}_i)$, where \mathbf{X}_i is a matrix of predictors. Note that the unit of analysis is the district.
2. Generate many draws from the joint distribution of β . For each draw:
 - (a) Estimate $\widehat{\psi}_d$ from the model (possibly by aggregating over precincts). Then calculate $\widehat{\epsilon} = \psi_d - \widehat{\psi}_d$, the residuals, denote the pdf of residuals by $f_{\widehat{\epsilon}}$.
 - (b) Randomly sample $x \sim f_{\widehat{\epsilon}}$ and calculate $\widehat{\psi}_d + x$.
3. These estimates form the empirical distribution $f(\psi_d | \widehat{\theta})$.

Variable	Mapping	Notes
j	Individual voter	
c	Simulation varies for: {Individual, precinct}	Implies n_c
d	Given by the electoral district for a context	Implies n_d
S_{10}	Set of treated voters. Implied by specification of c and assignment of treatment.	
S_{00}	Set of untreated voters. Implied by specification of c and assignment of treatment.	
$E[a_c(0)]$	Bound on possible change in turnout.	Predicted from available data or set to maximum (1) or minimum ($\frac{1}{2}$) possible values for all precincts.
ψ_d	Predicted margin of victory in district d .	Predicted from available data or third-party prediction algorithm (in US Congressional elections only).

Table A10: Mapping of parameters of the model onto variables in the data and simulation. I assume that, as in Case #1, no intervention would happen in the absence of the experiment, i.e. $|S_{11}| = 0$ and $|S_{01}| = 0$.

Supplementary Appendix: References

- Adida, Claire, Jessica Gottlieb, Eric Kramon, and Gwyneth McClendon. 2017. "Reducing or Reinforcing In-Group Preferences? An Experiment on Information and Ethic Voting." *Quarterly Journal of Political Science* 12 (4): 437–477.
- Arias, Eric, Horacio Larreguy, John Marshall, and Pablo Querubin. 2019. "Priors Rule: When do Malfeasance Revelations Help or Hurt Incumbent Parties?" Available at https://scholar.harvard.edu/files/jmarshall/files/mexico_accountability_experiment_v13.pdf.
- Asunka, Joseph, Sarah Brierly, Miriam Golden, Eric Kramon, and George Ofofu. 2017. "Electoral Fraud or Violence: The Effect of Observers on Party Manipulation Strategies." *British Journal of Political Science* 49 (1): 129–151.
- Banerjee, Abhijit, Selvan Kumar, Rohini Pande, and Felix Su. 2011. "Do Informed Voters Make Better Choices? Experimental Evidence From India." Available at https://scholar.harvard.edu/files/rpande/files/do_informed_voters_make_better_choices.pdf.
- Bhandari, Abhit, Horacio Larreguy, and John Marshall. 2020. "Able and Mostly Willing: An Empirical Anatomy of Information's Effect on Voter-Driven Accountability in Senegal." *American Journal of Political Science* Forthcoming. Available at https://scholar.harvard.edu/files/jmarshall/files/accountability_senegal_paper_v5.pdf.
- Bidwell, Kelly, Katherine Casey, and Rachel Glennerster. 2020. "Debates: Voting and Expenditure Responses to Political Communication." *Journal of Political Economy* 128 (8): 2880–2924.
- Boas, Taylor, F. Daniel Hidalgo, and Marcus André Melo. 2019. "Norms versus Action: Why Voters Fail to Sanction Malfeasance in Brazil." *American Journal of Political Science* 63 (2): 385–400.
- Boatright, Robert G., Donald P. Green, and Michael J. Malbin. 2006. "Does Publicizing a Tax Credit for Political Contributions Increase Its Use?: Results From a Randomized Field Experiment." *American Politics Research* 34 (5): 563–582.
- 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.
- 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.
- Casey, Katherine, Abou Bakarr Kamara, and Niccoló Meriggi. 2021. "An Experiment in Candidate Selection." *American Economic Review* Forthcoming.
- 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): 51–77.
- Cruz, Cesi, Philip Keefer, and Julien Labonne. 2018. "Buying Informed Voters: New Effects of Information on Voters and Candidates." Available at <https://static1.squarespace.com/static/>

- 58c979fad1758e09d030809c/t/5c048e82898583120b1f73cc/1543802523246/
buying_informed_voters_web.pdf.
- Cruz, Cesi, Philip Keefer, Julien Labonne, and Francesco Trebbi. 2019. “Making Policies Matter: Voter Responses to Campaign Promises.” Working paper available at https://static1.squarespace.com/static/58c979fad1758e09d030809c/t/5cfed616d6104500019dff1b/1560204824899/making_promises_matter_6102019.pdf.
- de Figueiredo, Miguel F.P., F. Daniel Hidalgo, and Yuri Kasahara. 2011. “When Do Voters Punish Corrupt Politicians? Experimental Evidence from Brazil.” Available at https://law.utexas.edu/wp-content/uploads/sites/25/figueiredo_when_do_voters_punish.pdf.
- Dunning, Thad, Guy Grossman, Macartan Humphreys, Susan D. Hyde, Craig McIntosh, and Gareth Nellis, eds. 2019. *Information, Accountability, and Cumulative Learning: Lessons from Metaketa I*. New York: Cambridge University Press.
- Enríquez, José Ramón, Horacio Larreguy, John Marshall, and Alberto Simpser. 2019. “Information saturation and electoral accountability: Experimental evidence from Facebook in Mexico.” Working paper.
- George, Siddharth, Sarika Gupta, and Yusuf Neggars. 2018. “Coordinating Voters against Criminal Politicians: Evidence from a Mobile Experiment in India.” Available at https://scholar.harvard.edu/files/siddharthgeorge/files/voter_mobile_experiment_181126.pdf.
- Gerber, Alan S., and Donald P. Green. 1999. “Does Canvassing Increase Voter Turnout? A Field Experiment.” *Proceedings of the National Academy of Sciences* 96 (14): 10939–10942.
- Gerber, Alan S., and Donald P. Green. 2000. “The Effects of Canvassing, Telephone Calls, and Direct Mail on Voter Turnout: A Field Experiment.” *American Political Science Review* 94 (3): 653–663.
- Humphreys, Macartan, and Jeremy M. Weinstein. 2012. “Policing Politicians: Citizen Empowerment and Political Accountability in Uganda - Preliminary Analysis.” IGC Working Paper S-5021-UGA-1.
- Jha, Saumitra, and Moses Shayo. 2019. “Valuing Peace: The Effects of Financial Market Exposure on Votes and Political Attitudes.” *Econometrica* 87 (5): 1561–1588.
- Kendall, Chad, Tommaso Nannicini, and Francesco Trebbi. 2015. “How do Voters Respond to Information? Information from a Randomized Campaign.” *American Economic Review* 105 (1): 322–353.
- Lierl, Malte, and Marcus Holmlund. 2019. *Information, Accountability, and Cumulative Learning: Lessons from Metaketa I*. Cambridge University Press chapter Performance Information and Voting Behavior in Burkina Faso’s Municipal Elections: Separating the Effects of Information Content and Information Delivery, pp. 221–256.
- López-Moctezuma, Gabriel, Leornard Wantchekon, Daniel Rubenson, Thomas Fujiwara, and Cecilia Pe Lero. 2021. “Policy Deliberation and Voter Persuasion: Experimental Evidence from an Election in the Philippines.” *American Journal of Political Science* Forthcoming.
- Shineman, Victoria. 2020. “Restoring Rights, Restoring Trust: Evidence that Reversing Felony Disenfranchisement Penalties Increases Both Trust and Cooperation with Governmetn.” Working paper, available at https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3272694.

Sircar, Neelanjan, and Simon Chauchard. 2019. *Information, Accountability, and Cumulative Learning: Lessons from Metaketa I*. Number 10 New York: Cambridge University Press chapter Dilemmas and Challenges of Citizen Information Campaigns: Lessons from a Failed Experiment in India, pp. 287–311.

Wantchekon, Leonard. 2003. “Clientelism and Voting Behavior: Evidence from a Field Experiment in Benin.” *World Politics* 55 (3): 399–422.