

Mobilizing Pasadena Democrats: Measuring The Effects of Partisan Campaign Contacts

R. Michael Alvarez California Institute of Technology
Asa Hopkins Lawrence Berkeley Nation Laboratory
Betsy Sinclair University of Chicago

This paper examines the effect of an entire campaign using a randomized field experiment where the treatment consists of campaign decisions made by a campaign manager. In contrast to the majority of the field experiments found in the contemporary get-out-the-vote literature, this paper studies the actual behavior of a campaign within a particular election as opposed to studying particular mobilization tactics. Thus, the campaign itself chooses the method used to contact each individual within the randomly assigned treatment group. Contacts are made via face-to-face canvassing, phone calls, e-mails, and door hangers and consist of experienced volunteers making partisan appeals. We observe a large treatment effect of campaign contact despite a small number of face-to-face contacts, suggesting that the targeting strategy of the campaign manager is particularly effective.

Despite the proliferation of academic studies of voter mobilization campaigns, it is unclear whether recent research sheds light on how effective partisan voter mobilization is in the real world of practical politics. Most of the existing academic field experiments are conducted using non-partisan get-out-the-vote (GOTV) messages, but most real-world campaigns use partisan messages (e.g., Gerber and Green 2000a). Additionally, most studies have focused on the efficacy of different types of mobilization strategies—measuring the relative effectiveness of direct mail versus phone calls versus door-to-door canvassing, for example (Gerber and Green 2004). Finally, it is not clear that academic research has examined adequately the groups in the registered voter population that campaigns may target for different types of mobilization efforts, especially low or moderate propensity voters. Thus, few academic studies have examined what established political campaign managers and their respective staff do when left to their own devices to combine each of these components, devise the partisan messages, develop a target group in the electorate, and run a campaign.

The field experiment in this paper provided us with a unique opportunity to examine mobilization effects from a real-world campaign, conducted exactly as a real-world campaign manager would undertake various voter mobilization efforts. An evaluation

of a real campaign's get-out-the-vote efforts provides insights into the kind of turnout effects campaign managers might expect where the campaign consists of a Democratic political message which was delivered according to the decisions of the campaign manager. We find that the messages and strategies of the real-world campaign manager are highly effective in stimulating voter participation, more effective than might be extrapolated from previous academic studies using nonpartisan messages.

The existing voter mobilization literature has not yet resolved whether or not the content of the mobilization message affects turnout. Nickerson (2006) finds that when the message consists of a comparison of candidates as opposed to a comparison of parties, there is an increase in turnout. Similarly Bennion (2005) finds that for voters under the age of 30, appeals to civic duty increase turnout more than statements about a close election, but that for older voters there is no difference. Niven's (2006) field experiment found that negative campaign messages increased turnout, contrary to the earlier laboratory and observational studies of Ansolabehere et al. (1994) and Ansolabehere and Iyengar (1995), who found that negative messages depressed voter turnout.

Much of the literature has found that content does not matter. Green and Gerber (2000) found no variation in turnout according to a variety of messages

delivered (where they vary the message to include statements about civic duty, neighborhood solidarity, and whether it is a close election). These experiments explore direct mail, face-to-face canvassing, and phone calls. Nickerson (2006) compares professional versus volunteer phone banks and finds no variation in turnout by message (where the message includes statements about civil rights, terrorism, or generational solidarity). Arceneaux and Nickerson (2006) compare positive and negative messages with phone and face-to-face canvassing and again find no variation in turnout. Trivedi (2005) finds no variation in turnout with a direct mail experiment which makes different content appeals to ethnic identity. Finally, Green and Karlan (2006) find no variation in turnout with a robotic phone call experiment where they compare a get-out-the-vote message with an election protection message.

Most of the messages used in previous studies are nonpartisan. Our field experiment is different in that the content of the message is partisan and unrestricted: canvassers were encouraged to deviate from the script to expand upon the partisan message. Individuals may be more likely to vote based upon a partisan appeal—this may spark their interest in a particular contest or make a contest or issue particularly salient. Individuals may also be less likely to act based upon something abstract (a desire to participate in civic life) and more likely to act based upon a political fact which has a direct impact on that individual's life (that the candidate wishes to change a health care policy, for example). Very little work has been conducted on partisan campaign tactics—even less has analyzed the effect of an entire partisan campaign. Thus our experiment adds to the literature by determining the effect of a real campaign.

Our experiment allowed the campaign to choose how to contact each individual assigned to receive the mobilization message. While much of the get-out-the-vote literature has found many of the strategies employed by real-world campaigns inefficient or ineffective, these studies have typically used nonpartisan settings where the researcher has determined which individuals are eligible to receive a particular campaign treatment. This may or may not correlate with the patterns that a campaign would actually employ in choosing which tactics to use with particular individuals.

In our study the campaign decides which strategies to employ among the individuals randomly assigned to receive campaign treatment. Only one prior study has conducted a similar investigation: Nickerson, Friedrichs, and King (2006) examine the

effects of a variety of campaign tactics (door hangers, door-to-door canvassing, and volunteer phone calls) on turnout in the 2002 Michigan gubernatorial election. Like our study, their research examines the effect of partisan campaigning in an off-year election. However, in our study the process by which the campaign determined how to employ a particular campaign tactic is different. In their case, a campaign manager determined which strategy to employ within a congressional district; within that district, treatment and control groups were assigned for each strategy. In our study the campaign manager makes the campaign strategy decision on a household-by-household basis. While this makes it difficult for us to study the causality of the marginal effects of each distinct campaign tactic (because there is no randomization), we can estimate the total effect of allowing the campaign manager to make each decision (beyond the initial randomization). As a consequence, we can generalize our results to what would happen in a partisan off-year election campaign where the campaign manager was permitted to make all campaign decisions. Thus, our paper estimates the effect of an entire campaign on voter turnout, and we can evaluate the efficacy of the campaign strategy. Our study finds the campaign manager is highly effective; our study suggests that future experiments should evaluate the efficacy of many of the strategies employed by the campaign manager, to document the effectiveness of each component of the campaign.

This study tests several hypotheses. First, we hypothesize that partisan, manager-chosen campaign contacts will have greater effects than estimated in previous studies. We evaluate the overall effect of campaign contact using our randomized field experiment where contact is partisan and the method of contact is chosen by the campaign manager and staff. We then compare our results to those from previous field experiments involving nonpartisan contact. Second, we anticipate that, consistent with the habitual voting literature (e.g., Bendor, Diermeier and Ting 2003; Gerber, Green and Shachar 2003; Fowler 2006; Plutzer 2002), voters with higher propensities to vote will be more likely to vote in this election. As our data contains information about the prior voting history of all participants in our field experiment, we evaluate to what extent there is heterogeneity in the treatment effect due to prior propensity to vote. This test is particularly important given that recent studies have argued that certain voters may have a predisposition to participate, and that GOTV campaigns may work best among people who have previously demonstrated that they can be influenced

by environmental factors (Fowler, Baker, and Dawes 2008; Fowler and Dawes 2008). Finally, we are able to compare the effects of this campaign to other campaigns where particular tactics were tested; we hypothesize that the campaign manager, within the context of a campaign, will be able to make good decisions about how to allocate resources to different types of contact. Thus we evaluate the extent to which these different tactics are efficacious. Note that only the test of our first hypothesis will have the benefit of randomization; the second and third hypotheses will require tests on observational data with the exception of our ability to evaluate the effectiveness of phone contacts on turnout.

Experimental Approach

The Pasadena Area United Democratic Headquarters targets its efforts on California's 44th Assembly District, which includes Pasadena and its environs.¹ In 2006, the local State Assembly seat was open because of term limits. The district is generally considered to be a "safe" Democratic district, and the Assembly race was not close.² As a known Democratic area, however, local Democratic activists felt that the 44th District could play a significant role in statewide elections, where a large Democratic vote could aid in the effort to elect State Treasurer Phil Angelides to the Governorship, as well as the rest of the Democratic ticket.³

Prior to the November 2006 General Election, the Pasadena Area United Democratic Headquarters (UDH) decided to conduct a randomized experiment in order to evaluate their campaign tactics. Their hope was to increase Democratic turnout by focusing their campaign on Democrats with some history of voting or interest in politics, but who may not vote

¹This area includes Pasadena, the neighboring cities of La Canada Flintridge, South Pasadena, Arcadia, Temple City and Duarte; the district also includes a small section of the City of Los Angeles and some unincorporated areas of Los Angeles County. A map of the district can be viewed at <http://democrats.assembly.ca.gov/members/global/maps/map44.jpg>.

²The Democratic candidate, Anthony Portantino, won with 58% of the vote.

³The 2006 general election included the full slate of California state constitutional offices on the ballot (Governor, Lieutenant Governor, Secretary of State, Controller, Treasurer, Attorney General, and Insurance Commissioner). Also on the statewide ballot was U.S. Senator Dianne Feinstein, running for reelection against Richard "Dick" Montjoy. Finally, the statewide ballot included 13 state ballot measures. Voters in the 44th District also had on the ballot a U.S. House race, in some places a state senate race, and a variety of local races and local ballot measures.

without a reminder or request from the campaign. To this end they hoped to mobilize occasionally voting Democrats, where "occasionally voting" was defined as voting in two, three, or four of the last six statewide elections, or as having recently registered to vote. Individuals who had voted in the June 2006 primary election were not targeted, under the logic that they were already politically engaged in this election cycle and did not require an additional stimulus to turnout to vote.

The campaign selected a subset of 24,500 households within the 44th Assembly District which contained one or more "occasionally voting" Democrats. These households were then randomly assigned to treatment and control groups. The campaign, based on experience from previous campaign cycles, anticipated being able to attempt contact of approximately 17,000 households, and this determined the relative size of the treatment and control groups.⁴ Thus the randomly assigned treatment group contained roughly 17,000 households, a number chosen by looking at the number of targets identified in the UDH campaigns in 2002 and 2004. These households contained about 20,500 voters who met the campaign's voter history criteria, and a total of over 41,000 voters. In addition, roughly 7,500 households were randomly assigned to the control group, containing roughly 9,500 voters who met the voter history criteria, and a total of nearly 16,000 voters.

The UDH staff worked full work weeks, with a day off each week, and the office was open from about 10 AM to 10 PM each for the last nine weeks of the campaign. Weeknight phone banks (Sunday–Thursday) ran from approximately 6 PM to 8:45 PM (about two hours of active calling), and weekend afternoons primarily consisted of both walking and calling. A common week-night phonebank might have had between four and eight callers.

Several other Democratic campaigns operated in the same geographic area and targeted some of the same voters. None of these campaigns, however, targeted this experiment's treatment and control groups differently and these efforts were not part of the campaign studied in this experiment. First, Portantino for

⁴One advantage of working with an experienced campaign was that they recognized they were unlikely to have sufficient resources to contact every individual within the sample over the course of the campaign, and thus were more willing to designate a control group. Some campaigns are reluctant to do so, and the size of the control group must be large enough so that the experiment has sufficient statistical power at the $\alpha \leq .05$ level to be able to detect small campaign effects. In this experiment, we drew a large enough universe of voters so that the number of households per precinct would be similar to other campaigns conducted by the UDH.

Assembly sent mailers widely throughout the district. That campaign also organized separate precinct walking operations in the town of La Canada Flintridge, the candidate's hometown. Second, Congressman Adam Schiff ran successfully for reelection against both Republican and Green Party candidates. Schiff sent a small amount of election mail. Third, the California Democratic Party (CDP) selected a number of precincts within the 44th Assembly District to target with a direct mail campaign.⁵ The CDP sent direct mail to its target voters regardless of their presence in this experiment's treatment or control groups. We have no records of the contacts made by any of these organizations, and these contacts are not included in our experiment. The presence of these other campaign activities potentially reduces the impact of our campaign's effect by saturating both the treatment and control groups with campaign messages. Thus our estimated campaign effect is likely a conservative estimate of the true campaign effect. This strengthens our findings regarding the efficacy of the campaign.

Individual voters were aggregated to the household level, and then households were randomly assigned to treatment and control groups. Households within the control group would not be assigned to receive any contact during the course of the campaign. Households in the treatment group could be contacted in whichever way the campaign manager determined was most likely to increase turnout. This meant that the campaign was continually revising their strategies based upon their resources—the number of volunteers available, whether the weather or neighborhoods were conducive for face-to-face canvassing, the number of phone lines available to make phone calls, and other factors. This resulted in a variety of campaign treatments over the course of a nine-week campaign.

The campaign treatments consisted of a variety of methods, all of which contained a partisan get-out-the-vote message. For individuals assigned to receive treatment, the treatments were administered in the following way: if an individual's record had an email address, then that individual was assigned to receive an email during the final GOTV push.⁶ Similarly, if

an individual had a phone number, then that individual was assigned to receive a phone call. Note that not all individuals assigned to receive emails or phone calls did so. Some phone numbers and email addresses were incorrect, and additionally some individuals were not able to be reached by phone. When recording contacts, an email contact was recorded if it was attempted, regardless of whether the message was delivered successfully.

The phone call procedure was split into two phases—the initial campaign phase, which took place for the first eight weeks of the campaign, and the “get-out-the-vote” phase, which was conducted during the final week of the campaign. In the first phase, the procedure with phone calls was to attempt to call each available phone number until the point at which it was either possible to speak to an individual directly or to leave a message on their answering machine. The outcome (message or speaking to a person) was recorded at the point at which the contact was successful. Individuals who directly responded that they would be positively supporting the Democratic Party in the November 2006 election were first in line to receive a second phone call during the final get-out-the-vote campaign week. Individuals who had received messages or were not reached in the initial phase were called after all the positive responders had been attempted.⁷ Phoning was done first to permanent absentee voters, and then to the rest of the list, without regard to voting history, gender, ethnicity, etc. There was a small program to have young volunteers call younger voters, but it was not large or clearly delineated from the rest of the calling.

The decision to walk a precinct was based on the number of individuals assigned to our treatment group that lived in that precinct. The campaign determined that walking in precincts with high numbers of individuals assigned to treatment would be most efficient in terms of spending resources. Additionally, walking was done only in precincts which were also targeted by the CDP. Most of the walkers had canvassed in previous elections, and thus were allowed to go off-script in administering the Democratic message. The new canvassers received simple training and were asked to follow a basic script outline until they were comfortable going off of it. Walkers were usually also phonebankers who knew that script and understood that the point of their contact with the voter was to identify and motivate Democratic voters to turn out to vote.

⁵The Pasadena Area UDH also made phone calls and walked door-to-door to CDP-targeted voters who were outside the treatment and control groups of this experiment. That is, the CDP targeted a different universe of potential voters, and the Pasadena Area UDH attempted to get out the vote from the union of its target universe (on which this experiment was run) and the CDP universe.

⁶The script for each email is included in this paper's Supplementary Materials, available at <http://home.uchicago.edu/~betsy/papers/index.html>.

⁷The script for the phone calls is included in paper's Supplementary Materials, available at <http://home.uchicago.edu/betsy>.

Some precincts were walked with door hangers, using an assignment mechanism similar to that used for precinct walking. Door hangers went to precincts targeted by the CDP, and covered them completely. This is a nonrandom sample of the treatment group. Precincts with higher numbers of households in the treatment group were assigned to receive a door hanger treatment. Due to the larger number of volunteers for GOTV weekend, as well as the relative speed of door hangers compared with precinct walking, more precincts were walked with door hangers than had been walked earlier in the campaign.

Individuals were classified as “treated” if they or any member of their household received any sort of treatment: an email, a phone call, an at-home visit, or a door hanger. Each treatment consisted of a Democratic message. Note that regardless of whether a household contained a single voter assigned to the treatment group, or several voters assigned to the treatment group, the household would receive a single phone call, doorhanger, or at-home visit. This was implemented by the campaign such that the canvasser would ask to speak to any of the individuals within the household who had the appropriate voter history. This is a very conservative definition of “treated” and the strength of our campaign effect should be considered in the context of this definition.

Finally, one important aspect of our field experiment is that all of the phone calls and door-to-door contacts were administered by volunteers. Nickerson, Friedrichs, and King (2006) suggest the possibility that volunteer phone calls may be more effective and findings from García-Bedolla, Michelson, and McConnell (2007) support this finding. We return to this point after presenting our results.

Data

We obtained a preelection dataset of registered voters from Political Data, Inc.⁸ This dataset includes the gender, partisan registration, ethnicity, and age of each registered voter. Additionally, this file documents the

⁸Political Data, Inc. (PDI) is currently the largest provider of voter information files in California, primarily providing voter information data to political consultants, campaigns, and polling firms. PDI obtains voter information directly from election officials and integrates other information not commonly found on the public voter history file. PDI updates their voter information database frequently, as often as once per month during an election. See <http://www.politicaldata.com> for additional information on PDI. Our preelection dataset used for our randomization was obtained in August 2006.

previous elections in which each individual has turned out to vote.

We used this dataset to conduct the randomization for our field experiment in September 2006, where each household was assigned to either the treatment or control group. The campaign successfully documented whether or not each individual was contacted, and additionally, recorded the type of stimulus and date of contact. After the election, Political Data, Inc. then provided us with an updated file which recorded whether or not each individual in our dataset had participated in the 2006 general election. Thus for each registered voter within each household, we know from the postelection voter history file whether or not they participated in the 2006 general election. Based upon the randomization, we are then able to determine the effect of campaign contact on voter turnout, and we are able to condition this statement on the treatment variable to get an accurate assessment of the size of this effect. Additionally, we are able to observe differing turnout rates by time of contact as well as by type of contact. The randomization, however, occurred prior to the beginning of the experiment. Thus we are not able to necessarily infer that the timing or type of contact is the causal mechanism for the turnout. We do however present our observational results on timing and type of contact to gain some insight into which strategies are most likely to be efficacious.

Table 1 documents the dynamics of household contacts. Each cell in the treatment group or control group category represents the number of individuals the campaign was able to successfully contact in that week (the row). Note that the campaign can select which individuals to contact in which weeks, so that while this classification has the potential to provide insight into the mechanism by which this campaign increased turnout, it is not possible to infer anything directly about the timing of the contact and the resulting vote from the analysis reported here. For example, the campaign called voters who are registered as permanent absentee Voters earlier in the cycle than voters who are not, on the assumption that these voters often vote before election day.

The information presented in Table 1 shows that the number of contacts steadily increased during the course the campaign. In the first two weeks of the experiment, the campaign contacted just over 400 households. But by the eighth week, they contacted over 3500 households; and in the GOTV week, they contacted over fifteen thousand households. The campaign was able to make more contacts closer to the election in part because volunteers are easier to

TABLE 1 Unique Individual Contacts by Week and Type of Contact

Week	Treatment Group	Control Group	Total
Week 1	335	12	347
Week 2	93	0	93
Week 3	425	0	425
Week 4	1,829	11	1,840
Week 5	674	50	724
Week 6	1,432	99	1,531
Week 7	735	64	799
Week 8	3,595	166	3,761
GOTV Week	15,140	1,453	16,593

Contact	Treatment	Control	Total
Phone	7,722	440	8,162
Walk	823	79	902
Email	1,426	183	1,609
Door hanger	7,313	700	8,013
Any Type of Contact	15,217	1,296	16,513

Note: Each cell entry represents the number of contacts the campaign was able to make successfully, either in that week or by that method. Individuals may have been contacted by multiple forms or multiple times within a single week and thus for clarification the any contact field represents the number of unique contacts made by the campaign and is not a sum of the values above.

find as election day nears. Individuals were often contacted more than once during the campaign; thus in Table 1, it is likely, for example, for an individual to be contacted both in the first week and the GOTV week, but not every single voter who was contacted in the early weeks was contacted again during GOTV week. The timing of contacts is a component of the strategy for the campaign that we do not explore in depth here but anticipate exploring in future work.

Based on the data collected by the campaign, we also know the type of contact: by phone, by face-to-face canvassing (here called “walk”), by email, or by a door hanger.⁹ Table 1 presents the type of contacts the campaign was able to make successfully, breaking this down further into the contacts made in the treatment and control group. Contacts made within the control group are inadvertent, but we document these contacts regardless so as to account for them when measuring the effect of the campaign. The majority of the campaign contacts in the treatment group were made by door hanger and by phone (over 7,000 contacts for each of these types), while only 1,426 individuals were contacted by email and a

⁹A door hanger is a printed message from the campaign which is left on an individual’s front door.

mere 823 were contacted by face-to-face canvassing. The split for the inadvertent control group contacts is similar but much smaller—a total of 1,296 individuals were contacted from the control group while 15,217 individuals were contacted from the treatment group. We have a slightly higher number of inadvertent contacts for the control group—8.10%—than in a typical mobilization experiment.¹⁰ Note that some individuals were contacted multiple times, so that the row sum for the type of contact is not equivalent to the “any contact” cell entry. This “any contact” category will be useful to determine the effect of the campaign.

Methods and Results

In order to determine whether or not partisan, manager-chosen campaign contacts will have greater effects than previously estimated in other studies, we estimate the effect of the campaign contacts on turnout. We consider the individual to have been treated if that person received any campaign contact whatsoever and determine the effect of that contact on the likelihood that the person would turn out to vote. We begin by estimating the intent-to-treat effect (*ITT*) and the treatment-on-treated effect (*TOT*). The intent-to-treat effect is defined as the observed difference in turnout between those assigned to the treatment and control groups. If the contact rate is 100% of the treatment group then the intent-to-treat effect is identical to the treatment effect. Generally this is not the case, however, and to calculate the treatment effect we must adjust for the contact rate. The actual treatment (*TOT*) effect is defined as the intent-to-treat effect divided by the contact rate.¹¹ Again we define treatment as any campaign contact.

Formally, let Z be an indicator which defines whether or not the individual is assigned to the treatment or control. Let x denote whether the individual actually receives the treatment. Let y indicate whether or not the individual voted. Then the intent-to-treat effect is defined as: $ITT = E(y|Z = 1) - E(y|Z = 0)$.

¹⁰Sets of volunteers in both South Pasadena and La Canada, two areas contained within our experimental universe, printed some contact materials before the control group had been selected and flagged in the database. In addition, it is possible that campaign staff or volunteers accidentally printed contact materials for the wrong group of voters—those who were designated to be part of the control group—by mistakenly selecting to print materials for the “control group” category from the database. This strengthens our need to use a 2SLS estimator, which we discuss later in the text.

¹¹For additional discussion of these quantities, see Rosenbaum (2002) or Morgan and Winship (2007).

Following from this, then $TOT = \frac{ITT}{E(x|Z=1) - E(x|Z=0)}$. Note again that the difference between the intent-to-treat and the total-treatment-effect is that the *TOT* accounts for the contact rate. In our particular example, some fraction of the control group also received an inadvertent contact, thus the voting rate of the control group may be higher—and thus the *ITT* is a particularly conservative estimate of the treatment effect. We can control for the contact rate of both the control and treatment groups, however, by estimating *TOT* after regressing *y* on *x* and using *Z* as an instrument (e.g., Gerber and Green 2001).

We present the results for *ITT* and *TOT* in Table 2. The contact rate for the campaign is 28%, which is consistent with previous studies in the literature, in particular given the campaign’s emphasis on phone contacts. We cluster the standard errors by household to account for potential spillover effects (one individual speaking to another within the same household).¹² We estimate an *ITT* of 1.59 percentage points—the difference in turnout between the percent voting in the treatment group and the percent voting in the control group—which is statistically significant at the $\alpha \leq .05$ level. Once we condition on the contact rate, we find an estimated *TOT* of 5.6 percentage points, which is again significant at the $\alpha \leq .05$ level. Relative to previous studies, this is a large effect. Recall again that this effect is due to campaign contact that consisted mostly of phone calls made by volunteers. In the existing literature, the highest estimate of the increase in turnout due to phone contact by volunteers is 5.0 percentage points with the inclusion of covariates (Green 2004), with the average size of an effect of 2.4 percentage points. This suggests that partisan campaign contacts, especially those spread out over the a 10-week campaign and chosen by a campaign manager, may produce greater turnout effects than are implied by prior studies of nonpartisan campaign contact efforts, although this estimate is not statistically distinguishable from most others.¹³

¹²As we anticipate that there will be correlation in the voting decisions within a household, we want to account for this correlation in estimating robust-cluster standard errors. Thus, we can consider the variance of our estimator to be: $Var(\hat{\beta}) = (X'X)^{-1} \sum_{j=1}^{n_c} u'_j u_j (X'X)^{-1}$ with *j* clusters where $u_j = \sum_i e_j x_i$.

¹³A fixed-effect meta-analysis conducted on the coefficients from the previous literature produces an average effect size of 1.57 with a 95% confidence interval of (.55,2.58) from the previous literature and our estimate of 5.6 has a 95% confidence interval of (2.07, 9.13); however, given the size of the standard error from our analysis the 95% confidence interval from our estimate includes all but the Nickerson (2007) and McNulty (2005) estimates on the list.

TABLE 2 Summary: Campaign Effect on Turnout

N in Treatment Group	41,683
N in Control Group	15,999
N in Treatment Group who are Actually Treated	15,217
N in Control Group who are Inadvertently Treated	1,296
N who Voted in Treatment Group	24,805
N who Voted in Control Group	9,266
Percent Voting in Treatment	59.51%
Percent Voting in Control	57.92%
Percent Contacted in Treatment	36.51%
Percent Contacted in Control	8.10%
Contact Rate	0.28
Estimated Intent-to-Treat Effect	1.59
Standard Error of Intent-to-Treat Effect	0.46
Estimated Treatment Effect	5.61
Standard Error of Estimated Treatment Effect	1.80

Note: Standard errors are robust and clustered by household.

We calculate the *TOT* using two-stage least-squares (2SLS), where we estimate the effect of a canvass contact on individual turnout in the election where the application of the contact treatment may be endogenous. It may be the case that the individuals who are able to be contacted by the campaign are also those individuals who are most likely to vote—this is true for both the individuals in the treatment group who are contacted as well as for the inadvertent contacts in the control group. In order to implement this analysis we construct an instrumental variable for campaign contact using the randomization assignment. Note that we are assured that there is no correlation between the random assignment and the error terms, as the assignment was completed randomly prior to conducting the experiment, and this process assures us that the instrument is statistically independent of the observables and unobservables. Our instrument here is fairly good: the correlation between the treatment assignment and any contact is .28. The presence of this instrument allows us to estimate the effect of the campaign contact.

We compare the 2SLS estimate of the contact effect with the 2SLS estimate of the contact effect with all available control variables as well. These include partisan registration (included here as an indicator for whether or not the individual was a registered Democrat), gender, age, ethnicity (included here as an indicator for whether or not the individual was non-Anglo; this field was enhanced by Political Data, Inc.

using a database of ethnic last names and other census data), and the variable we calculated as the individual's propensity to vote. This allows us to extract the impact of some of the variables which we know tend to influence voter turnout from other models and to focus on the contact effect alone. We present these results in Table 3. We again calculate robust-cluster standard errors by household to account for potential spillover effects (one individual speaking to another within the same household). Additionally, so that we can directly compare the two models, we create indicator variables for the instances where there is a missing covariate.¹⁴ This also allows us to observe the effect of the the indicators for the missing data; in our analysis we find that many of the missing indicators are statistically significant, which suggests that their absence from the voter file may be correlated with the voter's turnout decision.

Here we still find a large effect of contact when we include the covariates—6.56 percentage points—which is statistically significant at the $\alpha \leq .05$ level, and a statistically significant effect of contact without the covariates as well (5.6 percentage points). This implies that the covariates are themselves not driving the results and that the campaign is effective at increasing turnout. Note that here we have included an individual's propensity to vote as one of the covariates. We find a positive and statistically significant effect; individuals who are likely voters are again likely to turnout to vote. We observe a positive and statistically significant effect of being a registered Democrat (although the effect size is very small), a positive and statistically significant effect of age (which again is consistent with our expectations that older voters are more likely to go to the polls), and a negative and statistically significant effect for non-Anglo ethnicity, which is also consistent with our expectations.

In the paper's supplementary materials, we provide a table giving the estimated mobilization effects of volunteer phone banks from previous studies. Our effect, 6.56 percentage points with covariates, is slightly larger than the magnitude of campaigns with volunteer phone banks, which is 5.0 percentage points with covariates (Green 2004). However, given the size of the standard error from our analysis the 95% confidence interval from our estimate does include this effect. One possibility is that the volunteers here are of particularly high quality, and their

TABLE 3 2SLS Coefficients: Campaign Effect on Turnout

Variable	Without Covariates	With Covariates
Contact	.056* (.018)	.0657* (.0165)
Democratic Registration		.01** (.0052)
Missing Registration		.52* (.014)
Female		-.003 (.0036)
Missing Gender		-.04* (.014)
Age		.002* (.0001)
Missing Age		.11* (.011)
Propensity to Vote		.601* (.007)
Missing Propensity to Vote		.525* (.008)
Non-Anglo		-.018* (.005)
Missing Race		.02* (.006)
Constant	.575* (.006)	.156* (.008)
N	57,682	57,682
Root MSE	.4924	.45352
Clusters	38,591	38,591

* = $\alpha = .05$.

** = $\alpha = .10$.

Note: The instrument used here is the random assignment. Standard errors are robust and clustered by household.

ability to go off script and engage voters in conversation is particularly efficacious. This would correspond with findings in Nickerson (2007), where the quality of the phone call is key in determining turnout. However, our study does not consist entirely of phone calls.

At the very least, this is a campaign which has generated large turnout effects relative to those estimated in the earlier GOTV literature. Our effect is particularly surprising given the fact that there were other campaign activities taking place in Pasadena in this election cycle, making our estimate of the campaign effect a conservative one. Therefore, although our randomization is not conducted with respect to type of contact, vote history, or timing, we investigate each of these campaign strategies to determine to what extent the effect of the campaign may be attributed to each. Although we are unable to draw

¹⁴That is, for each instance where the covariate is missing in our data, we generate a variable with a value of 1, otherwise the variable takes on a value of 0.

causal claims in all of these instances due to the lack of randomization, we are able to document which characteristics are most highly correlated with the increased turnout. We first determine to what extent the treatment effect may be ascribed to the type of contact and then go on to examine each of the different campaign strategies in turn. The mechanisms—type of contact, timing, and who is contacted—may influence the campaign’s success but are not randomly assigned and thus cannot provide causal inferences, however, the correlations between turnout and the time of contact or turnout and the type of contact can be useful in ascertaining insight into the campaign’s success.

We first present Table 4 which describes the method(s) by which each individual was contacted by the campaign. Although these methods are not randomly assigned, we are able to compare the effects of this campaign to other campaigns where particular tactics were tested; we hypothesize that the tactics selected by the campaign manager are particularly efficacious. Some individuals were contacted with multiple methods; a very small number were contacted with three or four different approaches. Again, it is

not possible to draw inferences from any analysis presented here on the turnout resulting from these methods because they were not randomly assigned. However, it is possible to observe which methods are most correlated with an increase in turnout and potentially ascribe the campaign’s success to either the choice of those methods for those voters or for the efficacy of the method itself. The method of contact was determined independently of the particular individual’s characteristics; households received a door-to-door canvasser if they lived in easy-to-reach precincts, e-mails went to everyone with an e-mail address, and individuals who had already been reached once received a second phone call during the GOTV weekend. The campaign manager’s skill was not in microtargeting a set of voters with particular strategies, but instead in choosing how to allocate scarce resources to contact the greatest number of voters with motivated volunteers.

The majority of the contacts are made by a single type of campaign strategy—by either a door hanger, a phone call, an e-mail, or a door-to-door canvasser. The remaining contacts mostly consist of a combination of two methods (predominantly phone call and door hanger). We classify the treatment into five categories—one for each of the individual contact methods (door hangers, e-mails, phone calls, or door-to-door canvassing) and one if an individual was contacted by any mix of the prior four.

We present a summary of turnout rates by treatment type in Table 5, where we look at the differences in the rates at which individuals turned out to vote by the type of treatment conditional upon having received a campaign contact. Door-hanger contacts were made the most frequently but had the lowest observed turnout (51.00%), while phone calls were made almost as frequently and had the highest turnout (66.66%). This is surprising in contrast to the rest of the literature which tends to ascribe the highest turnout rates to door-to-door canvassing (here an observed 63.62% turnout), but could be due to the selection of households. Email contact and multiple types of contact had similar rates of turnout as well (approximately 59%), which are both interesting results: e-mail contact is associated with a higher level of turnout than expected, but this may simply be due to characteristics of the population which was able to receive an e-mail and thus merits future analysis. In particular, this result is not in agreement with the existing field experiment literature where randomized experiments have yet to find that e-mail generates positive turnout (Bennion and Nickerson 2008; Nickerson 2007), and this e-mail result does not account for a

TABLE 4 Type of Contact Received by Treatment Assignment

Contact	Control Group	Treatment Group
Door hanger	632	5,892
Email	171	963
Phone Call	346	5,961
Door-to-Door Canvass	41	418
Door hanger + Email	9	153
Door hanger + Phone Call	53	1,137
Door hanger + Door-toDoor Canvass	0	43
Email + Phone Call	3	244
Door-to-Door Canvass + Email	0	16
Door-to-Door Canvass + Phone Call	35	292
Door hanger + Email + Phone Call	0	35
Door hanger + Email + Door-to-Door Canvass	0	1
Email + Phone Call + Door-to-Door Canvass	0	12
Door-to-Door Canvass + Door hanger + Email + Phone Call	0	2
Total	1,290	15,169

Note: Each category is exclusively defined. Each cell entry refers to the number of individuals who were contacted by the method(s) displayed in the first column.

TABLE 5 Vote Rates by Type of Campaign Contact

Type of Contact	N Treated	N Voted	Observed Turnout
Email	1,134	680	59.96%
Walk	459	292	63.62%
Phone	6,307	4,204	66.66%
Door hanger	6,524	3,327	51.00%
Mixed	2,035	1,206	59.26%
Total	16,459	9,709	58.99%

Note: That individuals in this sample were all those assigned to receive a campaign contact.

failure-to-treat (e-mails that bounced, for example) and thus is actually a conservative estimate on the effect of e-mail. Yet it is possible that an e-mail message with a partisan frame is simply more efficacious than previous nonpartisan academic research and again, given that this treatment is not randomized, it may also be a feature of this population. That multiple contacts produced a similar effect on turnout indicates that two types of campaign contact may not have a simple additive effect on participation rates.

We next examine the effect of phone calls in greater depth as we can take advantage of the fact that we know the mechanism by which individuals in the experiment were assigned to receive a phone call from the campaign. If it is the case that an individual had a phone number in the dataset, they were assigned to receive a phone call. Therefore, we are again able to rely upon a 2SLS estimate of the effect of contact where we create a new instrument using both an indicator for whether or not an individual had a phone number in the dataset and the random assignment variable. We can then examine the effect of a phone call on turnout. Here we find that there was a 8.1% treatment-on-treated effect from phone calls, with a standard error of 1%.¹⁵ The campaign manager's choices appear to have been very good—phone call contact was highly effective at getting out the vote, and these contacts produced turnout rates greater than suggested in previous research on the effect of volunteer phone calls on voter mobilization.

When we dug deeper into the data we observed higher percentages of individuals voting who actually spoke to a volunteer as opposed to those who were

¹⁵This calculation incorporates all the 43,562 observations who have a phone number and relies upon conditional random assignment. We have 6,307 individuals who received a phone call in this group and of those, 4,204 voted. Including all the control variables into this regression results in a 7.0% treatment-on-treated effect, with a standard error of 1%.

simply contacted by leaving a message on an answering machine: 68.74% relative to 58.28%, respectively.¹⁶ This again suggests that phone calls completed by motivated volunteers are likely to be very effective. When we compare our finding to the existing literature, this also suggests that by allowing the campaign manager to choose whom to call, the effect of volunteer phone calls increases.¹⁷

We next examine the extent to which vote history or timing determines to what extent the campaign is effective. We first present results where we separate individuals who were contacted during the campaign by category—new registrants, unlikely voters, low-propensity voters, and high-propensity voters. We hypothesize that voters with higher propensities to vote will be more likely to vote in this election. We calculate a voter's propensity to vote by looking at the number of statewide California elections in which they participated out of the number for which they were eligible based upon their registration date over the last six elections.¹⁸ Note that this campaign was particularly interested in the population of occasional voters and hence chose a selection of individuals (from which the randomization occurred) who had some history of participating in politics but were not the individuals who were likely to always participate. Recall that the universe of voters studied consists of voters selected to meet certain voting history criteria, plus voters who share a household with such a voter. This strategy was designed to maximize Democratic turnout; within a population of individuals who are already likely to vote, the campaign's efforts would not likely result in large numbers of additional votes. A by-product of this strategy is that it is easier to measure the effect of additional stimulus to get out to the polls.

Here we have a total population of 57,682 individuals. We define an unlikely voter as an individual who has not voted in any of the elections for which they were eligible, a low-propensity individual as someone who has voted in at least one, but less than 50% of the elections for which they were eligible

¹⁶732 of the 1,256 who were contacted by an answering machine message voted, while 3,472 of the 5,051 who spoke with a volunteer on the phone voted.

¹⁷Note also that turnout was higher among the population with a working phone line and an answering machine or live contact than it was among the voting population as a whole, an indication of the kind of selection bias we cannot control for in this experiment.

¹⁸These included the November 2002 General Election, October 2003 Special Election, June 2004 Primary Election, November 2004 General Election, November 2005 Special Election and the June 2006 Primary Election.

(since November 2002), and a high-propensity individual as someone who has participated in 50% or more of the elections for which they were eligible (since November 2002). We classify an individual as newly registered if they had registered after November 2005, but not vote in the June 2006 election. (Recall that the campaign did not target voters who voted in the June 2006 election.) Note that Political Data, Inc. obtained their dataset from the Los Angeles County Registrar-Recorder's office in August 2006. This implies that this classification of newly registered voters may be somewhat different than the population of individuals who might register in the days and weeks immediately before an election, and thus these results may not apply to those voters if they are different than the population of newly registered voters in this study.

Looking at the efficacy of the campaign for the new registrants in Table 6, we find a positive and statistically significant (at the $\alpha \leq .05$ level) large effect of campaign contact. Our estimated *ITT* is 6.26 percentage points and our estimated *TOT* is 24.24 percentage points. This is a sizeable effect and contrasts starkly against the treatment effects for the other groups. The unlikely voters have no statistically significant effect of contact, while the low-propensity voters have a large and statistically significant effect (14.63 percentage points for *TOT*) and the high-propensity voters have a statistically significant, but smaller effect (3.34 percentage points for *TOT*). However, it is important to draw cautious inferences about these differences in treatment effects across voter propensity. The type of treatment each individual received was not randomly assigned and thus merits further investigation.

We also examine when the campaign chose to contact each of these individuals. Here we produce a classification for whether an individual was contacted early in the campaign (in weeks 1–5), in the middle of the campaign (in weeks 6–8), or during the final get-out-the-vote drive (week 9). Thus we have 3,353 individuals who were contacted in the early campaign, 5,784 individuals who were contacted in the middle campaign, and 16,593 individuals who were contacted during the final GOTV week. Note that individuals can be contacted multiple times over the course the campaign so that these classifications are not exclusive. We then look at the marginal effect of each of these classifications on the turnout decision, conditional upon being contacted. We use early contact and phone call contact as our two base comparison categories. We anticipate that later contact will be more effective (and thus we will observe positive and statistically significant coefficients for middle and GOTV contact).

In Table 7 we present the ordinary least-squares coefficients for the variables described above.¹⁹ The timing variables (middle and GOTV) have coefficients which are consistent with our expectations—it appears that the marginal effect of contacting a voter later in the campaign is positive.²⁰ Like the type of treatment discussed above, the timing of treatment in our field experiment was not randomly assigned, and thus producing a better estimate for how the timing of treatment influences voter mobilization should be the subject of future research. It is possible that in fact there are additional factors which influenced the voters' mobilization which are exogenous to the point at which they were contacted by the campaign. The campaign manager set out a strategy to contact as many individuals as possible each day and each week, based upon the availability of volunteers. Thus the effects of timing are highly correlated to the particular strategies implemented, as some campaign strategies required additional volunteers to complete.

Thus we find the campaign has a large effect on turnout relying upon a range of techniques and delivering a partisan message. The campaign manager's strategies and choices appear to do better than a typical nonpartisan door-to-door canvassing experiment. We found that phone calls are particularly effective at increasing turnout. We present observational evidence with respect to the timing of the delivery of the campaign contact (closer to the election appears better) and with respect to the type of voter who should be targeted (new registrants have the largest effect). We observe positive turnout from multiple types of campaign contacts, from multiple time periods, and across multiple types of registrants. This campaign was effective at generating votes.

Conclusion

The Pasadena Area United Democratic Headquarters runs a campaign generally similar to the one described in this paper every two years, typically targeting "occasional voting Democrats." One of their motivations in undertaking this experiment was to evaluate the efficacy of the campaign's chosen

¹⁹We also analyze this model in a binary choice setting and a logistic framework but present the OLS coefficients here for ease of interpretation. There is no qualitative difference between the binary choice and OLS results.

²⁰We replicate this model including all the covariates for each voter but find that the results do not change substantively. This results are available from the authors upon request.

TABLE 6 Treatment, Contact, and Turnout by Four Categories of Propensity-to-Vote

Propensity	Treatment	Control	Contact (Treat)	Contact (Control)	Vote (Treat)	Vote (Control)	ITT	TOT
New Registrant (0/0)	1,593	376	589	42	947	200	6.26 (2.86)	24.24 (11.06)
Missed all eligible (0/N)	4,068	1,546	1,311	146	1,038	390	.29 (1.3)	1.27 (5.7)
Low propensity (< .5)	13,029	4,940	4,737	519	6,100	2,126	3.78 (.83)	14.63 (3.21)
High propensity (> = .5)	22,993	9,137	8,580	589	16,720	6,550	1.03 (.56)	3.34 (1.8)

strategy and methods. In particular, they wanted to know whether their selection criteria for targeting voters was effective and whether particular kinds of voters were more or less affected by the campaign's efforts.

We have shown that the Pasadena Area UDH increased turnout in the 2006 general election by roughly 852 votes.²¹ The campaign was effective in increasing turnout among newly registered voters. (Recall that these were not voters registered by the campaign, but rather voters who had newly registered before the campaign began.) These voters, while a small population (less than 4% of the treatment group), accounted for about 12% of the campaign's total increase in voter turnout. The campaign reaped its greatest rewards among low-propensity voters who had voted at least once, a population made up largely of voters who had voted two or three times in the preceding six elections. This population accounted for over 60% of the campaign's net effect, while making up 31% of the treatment group. While the campaign's effect on higher propensity voters was small when expressed as a percentage point increase, the size of this universe meant that it accounted for a larger number of additional votes than the newly registered voter group. In terms of general campaign strategies, however, it should be noted that the campaign was effective with voters with a variety of voting histories, thus indicating that the targeting should not be limited to any one of the classes we considered.

We have also shown that campaign contact in the last few weeks before the election, and contact with a volunteer by telephone (relative to other means of contact), are associated with increases in voter turnout. Given the relative ease of finding volunteers for the final weeks of a campaign, and the low cost of volunteer phone calls relative to other methods of voter contact, this experiment indicates that partisan get-out-the-vote campaigns can be highly effective

without exorbitant costs. Again we are cautious about extrapolating these results to other settings and elections, but for this local election and partisan context, the campaign was effective at mobilizing the base.

This campaign had total expenses of about \$48,000 in 2006.²² With a treatment effect of 5.6%, this implies that each of the additional 852 votes that the campaign successfully turned out cost approximately \$57.34 per vote. Note that this cost includes all campaign overhead, and as such is not directly comparable to many of the "cost per vote" calculations described in much of the academic literature. Subtracting overhead costs of \$36,300 from the campaign's total expenses, we calculate a cost per vote of only \$13.73. The campaign's volunteer phone banks with experienced and enthusiastic volunteers were very cost effective.

We note that the presence of these other campaign activities in the Pasadena area during this time period potentially reduces the impact of our campaign's effect and that our estimated campaign effect is likely a conservative estimate of the true campaign effect. This emphasizes the power of our findings, as it is possible that in an area where no other campaign activities were taking place, we would have observed a larger treatment effect.

The principal finding of this paper is that the campaign manager was able to make reasonable choices and turn out voters at levels exceeding those seen in many academic studies. This suggests that there may be an aspect of campaign craft that goes into the creation and development of mobilization efforts which accounts for their varying levels of success. This also suggests that academic work may underestimate the actual effect of real-world voter mobilization efforts. In particular, our definition of household treatment is such that our campaign effect is likely a very conservative estimate of the treatment effect—by defining contact for everyone in the household our *TOT* effect

²¹The campaign effect is 5.6% of the 15,217 voters in the treatment group who were contacted.

²²Of that, \$9,000 was rent, \$18,500 was staff/consultants, \$3,800 was fundraising, and \$5,000 was other overhead (insurance, utilities, etc.). The total phone bill was about \$4,100. Additionally all literature printing costs were covered by the California Democratic Party.

TABLE 7 OLS Coefficients: Marginal Effect of Campaign Strategy and Timing on Turnout

Variable	Coefficient
Middle Contact	.031* (.010)
GOTV Contact	.033* (.013)
N	16,513
Root MSE	.48688

* = $\alpha = .05$.

The comparison category here is early contact. The dependent variable is an individual's turnout decision. The population includes all individuals who were successfully contacted.

is probably smaller than it actually is, as we are estimated that more individuals received treatment than actually did. In our study, we also explored the potential effect of the timing of voter contact, the type of voter contact, and how contact interacts with voting propensity. We found that volunteer phone contact to low propensity voters immediately before the election appears to have been particularly efficacious. In particular however, this study may have high rates of turnout because of the freedom allocated to the campaign manager to make choices about how to allocate resources. Future research should determine the mechanisms by which this happens, which may provide future insights into the intricacies of voter behavior.

We suggest that future research focuses on two types of field experiments. The first type of experiment would better evaluate the efficacy of partisan frames and specific campaign tactics. The literature has seen the role of framing in mobilization experiments dramatically change the evaluations and preconceptions we have about existing tools, such as the efficacy of direct mail (Green, Gerber, and Larimer 2008). These framing effects may be key to generating an effect of email contact with a partisan frame, for example, and merits future analyses. The second type of field experiment should seek to better understand actual campaign strategies. Campaign managers may hold interesting keys to unraveling the mystery as to why individuals turn out to vote, and their innovations can be tested by comparing the impact of a dynamically evolving campaign strategy on a randomly assigned group of individuals to the impact of a particular tactic on another group of randomly assigned individuals. These strategies can not only give us insight into which techniques matter for mobilization, but also help us as political scientists understand the degree to which people let politics and campaigns impact their lives. Campaigns do affect voter behavior, and experiments

which study campaigns as a whole can give us new insights into why people participate in politics. This experiment highlights the benefits of academics working closely with practitioners; future experiments could work to strengthen these relationships. For example, small randomization experiments could be conducted during these dynamic campaign decisions that allow both the practitioner to know which tactics are most efficacious and for the academic to have the advantage of realism by using a real campaign as a laboratory.

Acknowledgments

The authors would like to thank the Pasadena Area United Democratic Headquarters for providing this data and the opportunity to conduct a randomized experiment. We also thank Alexander Trechsel, Mark Franklin, Thad Hall, Guido Schwerdt, and other participants in the "Innovations in Research on Political Behaviour" workshop at the European University Institute for their comments about this research project.

Manuscript submitted 23 January 2009

Manuscript accepted for publication 31 March 2009

References

- Ansolabehere, Stephen, and Shanto Iyengar. 1995. *Going Negative: How Political Ads Shrink and Polarize the Electorate*. New York: Free Press.
- Ansolabehere, Stephen, Shanto Iyengar, Adam Simon, and Nicholas Valentino. 1994. "Does Attack Advertising Demobilize the Electorate?" *American Political Science Review* 88 (4): 829–38.
- Arceneaux, Kevin, and David Nickerson. 2006. "Even if You Have Nothing Nice to Say, Go Ahead and Say It: Two Field Experiments Testing Negative Campaign Tactics." Prepared for presentation at the annual meeting of the American Political Science Association, Washington D.C.
- Bendor, J., D. Diermeier, and M. Ting. 2003. "A Behavioral Model of Turnout." *American Political Science Review* 97 (2): 261–80.
- Bennion, Elizabeth A. 2005. "Caught in the Ground Wars: Mobilizing Voters During a Competitive Congressional Campaign. The Science of Voter Mobilization." Special eds. Donald P. Green and Alan S. Gerber. *The Annals of the American Academy of Political and Social Science* 601: 123–41.
- Bennion, Elizabeth A., and David W. Nickerson. 2008. "Email is Cheap, but is it Effective?" Presented at the annual meetings of the Indiana Political Science Association, Purdue University North Central.
- Fowler, James H. 2006. "Habitual Voting and Behavioral Turnout." *Journal of Politics* 68 (2): 335–44.
- Fowler, James H., Laura A. Baker, and Christopher T. Dawes. 2008. "Genetic Variation in Political Participation." *American Political Science Review* 102 (2): 233–48.

- Fowler, James H., and Christopher T. Dawes. 2008. "Two Genes Predict Voter Turnout." *Journal of Politics* 70 (3): 579–94.
- Garcia-Bedolla, Lisa, Melissa R. Michelson, and Margaret McConnell. 2007. "Heeding the Call: The Effect of Targeted Phonebanks on Voter Turnout." Unpublished manuscript. <http://www.hss.caltech.edu/mmcconnell>.
- Gerber, Alan S., and Donald P. Green. 2001. "The Effects of Canvassing, Telephone Calls, and Direct Mail on Voter Turnout: A Field Experiment." *American Political Science Review* 94: 653–63.
- Gerber, Alan S., and Donald P. Green. 2004. *Get Out the Vote!* Washington, DC: Brookings Institute Press.
- Gerber, Alan S., Donald P. Green, and Christopher W. Larimer. 2008. "Social Pressure and Voter Turnout: Evidence from a Large Scale Field Experiment." *American Political Science Review* 108 (February): 33–48.
- Gerber, A. S., D. P. Green, and R. Shachar. 2003. "Voting May Be Habit-Forming: Evidence From a Randomized Field Experiment." *American Journal of Political Science* 47 (3): 540–50.
- Green, Donald P. 2004. "The Effects of an Election Day Voter Mobilization Campaign Targeting Young Voters." Unpublished manuscript. Yale University.
- Green, Donald and Dean Karlan. 2006. "Effects of Robotic Calls on Voter Mobilization." Unpublished Manuscript. Yale University.
- McNulty, John E. 2005. "Phone-Based GOTV-What's on the Line? Field Experiments with Varied Partisan Components, 2002–2003." *The Science of Voter Mobilization*. Special eds. Donald P. Green and Alan S. Gerber. *The Annals of the American Academy of Political and Social Science* 601: 41–65.
- Morgan, Stephen L., and Christopher Winship. 2007. *Counterfactuals and Causal Inference: Methods and Principles for Social Research*. New York: Cambridge University Press.
- Nickerson, David. 2006. "Volunteer Phone Calls Can Increase Turnout: Evidence from Eight Field Experiments." *American Politics Research* 34: 271–92.
- Nickerson, David. 2007. "Quality is Job One: Professional and Volunteer Voter Mobilization Calls." *American Journal of Political Science* 51 (2): 269–82.
- Nickerson, David. 2005. "Partisan Mobilization Using Volunteer Phone Banks and Door Hangers." *Annals of Political and Social Science* 601 (September): 10–27.
- Nickerson, David W. 2006. "Does Email Boost Turnout?" *Quarterly Journal of Political Science* 1: 369–79.
- Nickerson, David W., Ryan D. Friedrichs and David C. King. 2006. "Partisan Mobilization Campaigns in the Field: Results from a Statewide Turnout Experiment in Michigan." *Political Research Quarterly* 59 (1): 85–97.
- Niven, David. 2006. "A Field Experiment on the Effects of Negative Campaign Mail on Voter Turnout in a Municipal Election." *Political Research Quarterly* 59 (2): 203–10.
- Plutzer, E. 2002. "Becoming a Habitual Voter: Inertia, Resources, and Growth in Young Adulthood." *American Political Science Review* 96 (1): 41–56.
- Ramirez, Ricardo. 2005. "Giving Voice to Latino Voters: A Field Experiment on the Effectiveness of a National Nonpartisan Mobilization Effort." *The Science of Voter Mobilization*. Special eds. Donald P. Green and Alan S. Gerber. *The Annals of the American Academy of Political and Social Science* 601: 66–84.
- Rosenbaum, Paul R. 2002. *Observational Studies*. 2nd ed. New York: Springer.
- Trivedi, Neema. 2005. "The Effect of Identity-Based GOTV Direct Mail Appeals on the Turnout of Indian Americans." *The Annals of the American Academy of Political and Social Science* 601 (September): 115–22.
- Wong, Janelle S. 2005. "Mobilizing Asian American Voters: A Field Experiment." *The Science of Voter Mobilization*. Special eds. Donald P. Green and Alan S. Gerber. *The Annals of the American Academy of Political and Social Science* 601: 102–14.

R. Michael Alvarez is Professor of Political Science, California Institute of Technology, Pasadena, CA 91125.

Asa Hopkins is a post-doc at Lawrence Berkeley National Laboratory, Berkeley, CA 94720.

Besty Sinclair is Assistant Professor of Political Science, University of Chicago, Chicago, IL 60637.

Copyright of Journal of Politics is the property of Cambridge University Press and its content may not be copied or emailed to multiple sites or posted to a listserv without the copyright holder's express written permission. However, users may print, download, or email articles for individual use.