DO PHONE CALLS INCREASE VOTER TURNOUT?
A FIELD EXPERIMENT

ALAN S. GERBER
DONALD P. GREEN
Yale University, Department of Political Science and Institution for Social and Policy Studies

Political campaigns in the United States have grown increasingly reliant on mass-marketing techniques as a means of mobilizing voters. Personal canvassing is gradually being replaced by impersonal forms of communication, such as television and direct mail. One of the most important developments has been the advent of inexpensive telemarketing services. Large telemarketing firms have the capacity to call hundreds of thousands of voters in a single day. Coupled with increasingly detailed data bases about the political and demographic profile of each voter, phone banks represent an attractive means by which to conduct large-scale get-out-the-vote campaigns.

The question is whether such efforts do in fact raise voter turnout. A great many nonexperimental studies have noted the positive correlation between "party contact" and voting (e.g., Kramer 1970; Rosenstone and Hansen 1993). Unfortunately, measures used to gauge party contact do not distinguish between personal and phone contact. More important, nonexperimental studies of voter mobilization may produce misleading results if parties are more likely to phone active voters. Experimental studies of phone mobilization are very rare. There appear to be only three previous experiments measuring the effect of phone calls encouraging voters to turn out. Eldersveld (1956, pp. 160–61) conducted a small experiment in which 33 Ann Arbor, MI, residents were contacted by phone prior to a local election. Miller, Bositis, and Baer (1981, p. 450) conducted a small experiment in the 1980 primary election in Carbondale, IL, contacting 39 people by phone. Both Eldersveld's nonpartisan appeal and Miller, Bositis, and Baer's partisan appeal appeared to increase turnout by more than 15 percentage points, although the small sample sizes cause the results to fall short of statistical significance. The only study of appreciable size is Adams and Smith's (1980) partisan get-out-the-vote experiment prior to a local Washington, DC, special election. Randomly dividing

Public Opinion Quarterly Volume 65:75–85.0033-362X/2001/6501-0004$02.50

2,650 voters into equal treatment and control groups, Adams and Smith find a statistically significant positive effect of phone mobilization. By our calculations, their experiment boosted turnout approximately 8 percentage points, from 23.7 to 31.9 percent (1980, p. 392).

Although these previous results are suggestive, our knowledge of the effects of phone appeals remains fragmentary. The small sizes of previous studies (Eldersveld 1956; Miller, Bositis, and Baer 1981) led to large standard errors and therefore imprecise estimates of the effect of the experimental stimulus. The only large experimental study (Adams and Smith 1980) employed a partisan appeal in a low-turnout special election, and special elections may be an unusual case where a reminder to vote is particularly effective. Despite these important earlier efforts, we still lack a sense of how large an effect phone calls have, under what conditions they might matter (e.g., high-turnout elections vs. low-turnout elections), and whether certain kinds of appeals (e.g., nonpartisan vs. partisan) are more effective than others.

The present study reports the results of a large-scale field experiment in which more than seventeen thousand registered voters were assigned to treatment and control conditions. In contrast to Eldersveld (1956) and to Miller, Bositis, and Baer (1981), the large size of our experiment permits much more accurate measurement of the marginal effect of a phone call. In contrast to Adams and Smith's (1980) large-scale study of a local special election, we measure the effect of a phone call during a relatively high-turnout election, the 1998 general election, which featured races for Congress, governor, and other state offices. An additional contribution of our work is that, in contrast to previous efforts, our statistical analysis guards against methodological errors that have led some previous scholars to overestimate the effects of phone contact. In the end, we find little indication that nonpartisan appeals communicated by phone increase turnout.

Experimental Design

The experiment was performed during the 1998 general election in West Haven, CT, a town of 54,000 people. Connecticut holds state elections in midterm years, so in addition to a congressional race, the ballot included races for several statewide offices (including governor), as well as races for state representative and state senator. From public records we obtained a list of all registered voters, which we sorted by household address. Excluded from the sample were all voters with post office box addresses and all addresses where more than two registered voters resided. We provided this list to Survey Sampling, Inc., which performed a telephone match and located phone num-

bers for approximately two-thirds of the names and addresses.[1] We then randomly assigned households on this list to the experiment and treatment groups. After the election, we determined from the voter cross-off sheets which registered voters had actually cast ballots. After excluding those registered voters who could not be found on the cross-off sheets, we were left with an overall sample of 17,100 registered voters of whom 6,843 were in the control group.[2]

The basic experimental treatment was a phone call reminding the voter of the upcoming election. All calls were made during the 3 days leading up to and including election day. The phone calls were made by a Washington, DC–area political consulting firm that specializes in political phone calls. The firm has done extensive work in U.S. Senate and House races, as well as state and local elections. We hired this firm in the spirit of building realism into our experiment: with the capacity to conduct more than one hundred thousand calls in a single day, this phone bank is typical of the large-scale firms that are becoming increasingly prevalent.

It should be noted, however, that professionalism in campaigning manifests itself not only in the quantity of calls that can be made but also in the quality of those calls. Having monitored several hours of phone calls, it is our impression that the calls were delivered in a routinized and at times rushed manner. In other words, the calls sounded as though they were made by a professional firm rather than local volunteers or neighbors. The telephone scripts were generally delivered competently, but sometimes hastily or mechanically.

The experimental get-out-the-vote (GOTV) phone calls each began with a reminder that the election was upcoming on Tuesday. The calls continued with one of three different short appeals designed to stimulate voting.[3] All three versions of the phone scripts (the scripts are provided in the appendix) began with: "Hi. This is [caller's name] calling from Vote '98, a nonpartisan group working with the League of Women Voters. We just wanted to remind you that elections are being held this Tuesday."

We also sought to estimate the effects of getting citizens to pledge to go to the polls, and so we prepared two versions of each of the three treatments.

1. This procedure excludes those registered voters residing in households with unlisted numbers, as well as households that have relocated but were not removed from the registration rolls. While randomization insures that we obtain an unbiased measure of the treatment effect for the population in our sample, it is theoretically possible that phone calls to those West Haven registered voters with unlisted numbers have a different effect. Turnout among the experimental control group was similar, though slightly higher, than the reported turnout rate for all voters on the registration rolls in West Haven (53 percent vs. 45 percent).

2. This adjustment, which involved only a very small percentage of the sample, stems from the fact that between the time our copies of the voter lists and the final election day lists were prepared, the registrar of voters removed from the cross-off lists voters who died or moved. Counting subjects who were removed from the list as not voting has no effect on our results.

3. These appeals were developed in collaboration with professional political consultants. The themes were selected and the wording crafted with the goal of maximum effectiveness. Each of the three appeals had approximately the same effect on turnout.

Table 1. Experimental Design

Experimental Group	Not-Asked Follow-Up Question	Asked Follow-Up Question	*N*
Control	N.A.	N.A.	6,843
GOTV message	3,595	5,229	8,824
Donate blood	N.A.	N.A.	1,433
Total			17,100

NOTE.—GOTV message = get-out-the-vote messsage; N.A. = not applicable.

In the first version the phone callers closed their appeal by saying, "We hope you'll come out and vote." For a random subset of each treatment group, this closing was replaced with the question "Can we count on you to vote this Tuesday?" The callers then waited for the respondent to answer before concluding the call.

Previous research suggests that asking the respondent to affirm a voting intention might boost the effectiveness of the appeal (Greenwald et al. 1985; Morwitz, Johnson, and Schmittlein 1993; Spangenberg and Greenwald 1999). As we are asking for a pledge of participation rather than simply soliciting a prediction about the respondent's intentions, our work finds a close analogy in Reams and Ray's recycling experiment, which demonstrates the importance of wringing a commitment out of would-be participants (Reams and Ray 1993). Although a verbal promise to a stranger is a long way from signing an enforceable contract, these previous findings imply that making a pledge might create or reinforce feelings of obligation.

A final treatment group received a phone call asking the respondent to donate blood to an upcoming Red Cross blood drive. (The script is provided in the appendix.) As described later in the article, the purpose of this appeal was to provide another benchmark for comparison.

Table 1 shows the number of registered voters assigned to the control and treatment groups.[4] Roughly seven thousand voters were assigned to the control group, and approximately nine thousand were assigned to the GOTV messages. Of those encouraged to vote, approximately 60 percent were asked the follow-up question "Can we count on you to vote this Tuesday?" The remaining approximately fourteen hundred voters were assigned to receive the blood donation message.

4. The slight differences in the size of the treatment groups results from chance variation introduced by our random assignment process and chance variation in the number of registered voters in each household.

Table 2. Voter Turnout Rates by Experimental Group

Experimental Group	Turnout Rate (%)	*N*	Turnout Rate among Those Contacted (%)	*N*	Turnout Rate among Those Not Contacted (%)	*N*
Control	53.5	6,843				
GOTV message	52.9	8,824	67.1	3,840	42.0	4,984
Donate blood	54.2	1,433	67.8	631	43.5	802

NOTE.—GOTV message = get-out-the-vote message.

Data Analysis

In our analysis, a voter is considered contacted if the phone message was delivered in its entirety to someone in the treatment household who identified themselves as the voter listed on the phone list.[5] Only a subset of the treatment group was actually contacted. The contact rate was 43.5 percent for the get-out-the-vote message and 44.0 percent for the blood drive message.

Table 2 shows the effect of the phone calls on voting rates. The table lists the turnout rates for those registered voters who were contacted, as well as the turnout rates among those who were not reached. It is clear that the turnout rate for those who received our treatment is much higher than the turnout rate of either the control group or those in the treatment group who were not contacted. However, two important complications in measuring the true marginal effect of the phone calls are that those who are easy to contact, and therefore easy to "treat," are expected to have higher turnout rates than those who could not be contacted, and that this difference stems from factors independent of our get-out-the-vote message.[6] Previous scholars have overestimated the effects of phone contact by failing to correct this source of bias. Studies have calculated the effect of a get-out-the-vote treatment by comparing the voting rates of those who are actually contacted with those who are not contacted; this second group was formed by combining the original control group and those individuals who were in the treatment group but, for some

5. The definition of a "contact" implies a different thing for voters in a one- vs. a two-person household. For households with only one registered voter, a contact means that this registered voter was identified and received the treatment. In the case of two-voter households, one of the two registered voters in the household was identified and received the treatment. For simplicity, in the data analysis that follows, we assume the effect of a contact is the same in one- and two-person households. The results are substantively equivalent when the treatment groups are broken into one- and two-voter households.

6. For example, people who are out of town in the days before the election will be hard to reach and are also more likely to be out of town on election day. The results in table 2 provide empirical confirmation that turnout is higher among those who are easier to contact.

Table 3. Estimated Effects of Treatment, Controlling for Contact Rates

Experimental Group	Experimental Effects
GOTV message	(52.9–53.5)/.435 = −1.3(SE = 1.8)
Blood donation	(54.2–53.5)/.440 = 1.6(SE = 3.3)

NOTE.—GOTV message = get-out-the-vote message.

reason, were not contacted.[7] Nonexperimental analysis in which individual turnout is regressed on a variable that indicates whether the individual has been contacted (e.g., by a political party) involves a similar calculation. Ordinary least squares regression of turnout on a dummy variable that equals one if there is a voter contact is algebraically equivalent to lumping together those the party does not try to contact with those the party did not successfully reach (e.g., Kramer 1970; Rosenstone and Hansen 1993).

Fortunately, there are statistical procedures that permit us to generate accurate estimates of the treatment effects. In the appendix we detail a statistical technique that isolates the treatment effect, patterned after the work of Angrist, Imbens, and Rubin (1996). In the statistical derivation in the appendix, we show that subtracting the turnout rate of the control group from the turnout rate of the treatment group and then dividing this difference by the observed "contact rate" isolates the effect of the experimental treatment.[8] Table 3 shows the estimated effects of the experimental treatments. We find phone calls to be ineffectual; the higher voting rates among those who were contacted stem from preexisting differences between those contacted and those not contacted, rather than the effect of the experimental treatment. The turnout rate for the experimental group is 52.9 percent, and the contact rate is 43.5 percent. The estimated effect of the phone calls is −1.3 percent.[9] The effect is not statis-

7. Adams and Smith (1980), e.g., calculated the effect of a phone call by comparing the turnout rates of those who were actually contacted (a subset of the treatment group) and those who were not contacted (the control group plus the portion of the treatment group not contacted). Eldersveld (1956) also used a similar technique in his experimental analysis. Reanalyzing the data from the former study shows that this methodology exaggerates the effect of phone calls, though the conclusion that phone calls did in fact have a statistically significant effect on turnout rates remains intact. Miller, Bositis, and Baer (1981) do not report contact rates.

8. In theory, if we were able to collect data for all the variables that differentiate those who are contacted from those who are not, a regression of voting behavior on actual contact would yield unbiased estimates of the effect of contact. In practice, however, this condition is impossible to achieve and verify.

9. The point estimate for the effect of the turnout call is negative. While this is unexpected, it is conceivable that the phone call irritated some people and made them slightly less likely to vote. As we discussed earlier, while the firm making the phone calls is a professional political telemarketer with many clients, the calls were not executed with great flair. At times, the callers mispronounced words in the scripts or voters' names, delivered the messages in an unenthusiastic tone of voice, and sounded more like employees rather than neighbors and volunteers. These problems occur in actual campaigns as well. In the 1999 mayoral election, the firm hired to make

tically significant; the 95 percent confidence interval for the estimated effect runs from +2.2 percent to −4.8 percent.[10]

An alternative approach to analyzing the effect of receiving a political message is to compare the turnout rates among those who were contacted and given a political message and those who were contacted and given some other message. If there is a difference in observed turnout rates, this could be ascribed to the political message rather than the ability of the experimenter to contact the voter, since in both cases the contact is successful. Recall that table 2 reports voting rates for those in the treatment group who were contacted and given the script describing an upcoming Red Cross blood drive. Comparing the voting rates of those who were contacted and got the Red Cross script to those who received the treatment scripts shows the get-out-the-vote calls produced no statistically significant increase in voting rates. In fact, those given the Red Cross script were slightly *more* likely to vote than the typical voter who got the political treatment. The "donate blood" script was associated with 1.6 percent higher turnout. This boost in turnout is small and, given a standard error of 3.3 percent, probably meaningless. Comparing the turnout rates among those contacted and given the donate blood script rather than a political script shows that the difference in voting rates across the political and nonpolitical treatments is very small and that the hypothesis of no difference across messages cannot be rejected ($p > .10$, two-sided test).

Discussion

Unlike previous experimental studies, the experiment reported here finds that a phone call before the election did not increase turnout rates. This contrast may be explained in various ways. In terms of sample size, our experiment is many times larger than previous studies, combined. Given that some previous findings were statistically equivocal (Eldersveld 1956; Miller, Bositis, and Baer 1981), the finding in these earlier studies that phone messages raised turnout may have resulted from sampling variability. The only large-scale experimental demonstration of the effectiveness of phone canvassing is Adams and Smith (1980). In that study, the phone message encouraged respondents to turn out, but the majority of the message text provided reasons to support one of the candidates. Although Adams and Smith found that the calls did not affect which candidate voters preferred, their partisan calls were more effective at mobilizing voters than the nonpartisan calls examined here. A

phone calls locally (New Haven) used a calling house with callers who had detectable southern accents. It is possible that getting a call of this sort might reinforce the cynicism of some voters, who might think they are being manipulated in some way.

10. We also examined whether appeals including the follow-up question ("Can we count on you to vote this Tuesday?") were more effective. There is some weak evidence that including the follow-up question boosted the effectiveness of the treatment, but the improvement is small and not statistically significant. For details, see Gerber and Green (2000a).

potentially significant difference between our study and Adams and Smith's work is that while Adams and Smith studied a low-turnout special election (in which only 24 percent of the control group voted), we studied an election where turnout was over 50 percent for the control group.[11] Further experimentation is needed to determine whether partisan appeals are, in general, more effective than nonpartisan appeals, or if phone calls are more effective in low-turnout elections.[12]

The ineffectiveness of nonpartisan phone calls obviously does not imply that other efforts to mobilize voters are doomed to fail. On the contrary, if it is true that phone calls are relatively ineffective at increasing turnout, one intriguing implication is that there may be a link between long-term declines in voter turnout and the transformation of campaign mobilization tactics. As Gerber and Green (2000b) note, personal canvassing, which they find to be a highly effective means by which to mobilize voters, has waned in recent decades with the decline in participation in partisan and nonpartisan groups. Instead, campaigns are more reliant on mobilization campaigns based on direct mail and phone banks, and over time the latter activity has increasingly been conducted by large organizations. If our findings are correct, the switch from personal to more impersonal mobilization may have a net demobilizing effect on the electorate.

It may strike the reader as disappointing that so much effort has gone into producing what might be characterized as a negative finding—phone canvassing did not seem to affect voter turnout in our experiment. Based on our reading of the extant experimental literature, which uniformly reports large increases in turnout in the wake of phone contact, we were surprised to discover that our calls did nothing to boost participation rates. We have suggested some reasons why our findings might differ from previous studies, and further experimentation is needed to isolate the conditions under which phone banks mobilize voters.

Appendix

SCRIPTS

1. Red Cross

After identifying someone on the list, the caller proceeds with:

11. A referee noted that a further difference between our study and the Adams and Smith (1980) study is that Adams and Smith had a higher contact rate than our study (72 percent vs. 44 percent). This suggests that one logical possibility for why our results differ from those of Adams and Smith: those who are difficult to contact react strongly to contact.

12. It is also possible that the effectiveness of a phone appeal varies across population groups. For example, our appeal makes passing reference to the League of Women Voters. One reviewer suggested that this may have made the appeal more effective for Democrats, women, and liberals, though the failure to detect any mobilization effect despite our large sample size implies that, unless the voter information actually discouraged Republicans, men, and conservatives from voting, the phone call had no mobilizing effect for any significant population subgroup.

"Hi. This is [caller's name] calling on behalf of the American Red Cross to invite you to donate blood at an upcoming blood drive in your community. Each day volunteer blood donors are needed to support patients in Connecticut's hospitals. Your blood donation could save someone's life. Can a representative from your local blood drive call you to schedule an appointment?"

2. Three Versions of the Treatment Message

After the text reported in the article, the phone call proceeded with one of three appeals. In the "close election" condition, the script continued with "each November, significant elections are decided by a small number of votes. The races this year are very close, so please vote on Tuesday." In the "civic duty" appeal, the phone call instead went on to say, "The success of our democracy depends on whether or not we exercise our right to vote, so we hope you'll come out and vote this Tuesday." In the "neighborhood solidarity" condition, callers stated that "politicians sometimes ignore issues in a neighborhood when its people don't vote, so we hope you'll come out and vote this Tuesday."

ISOLATING THE EFFECT OF THE EXPERIMENTAL TREATMENT

Suppose that an experiment is performed and only a subset of the treatment group is contacted. Let α be the probability that a voter is contacted. If a voter is successfully contacted, the voter is called "reachable," otherwise she is called "nonreachable." Let p_{nr} be the probability that a "nonreachable" citizen votes, let p_r be the probability that a "reachable" citizen votes without receiving the experimental treatment, and let $p_r + t$ be the probability that a "reachable" type votes after being exposed to the experimental treatment. Our aim is to estimate the value of t.[13] The probability that a randomly selected member of the experiment group votes equals

$$P_E = \alpha(p_r + t) + (1 - \alpha)p_{nr}. \tag{A1}$$

Since the control group and the treatment group were formed by random assignment, the probability that a randomly selected member of the treatment group is "reachable" is the same as in the control group. This implies that the probability that a randomly selected member of the control group votes equals

$$P_C = \alpha p_r + (1 - \alpha)p_{nr}, \tag{A2}$$

where the only difference between equations (A1) and (A2) results from the effect of the experimental treatment. Using the data produced by the experiment, we can form a statistic that estimates the treatment effect. Consider the statistic

13. This setup imposes the assumption of a "common treatment effect" for all individuals. This assumption is common in the literature on program evaluation. If the treatment effect varies across individuals within the treatment group, the estimator described at the conclusion of the discussion should be interpreted as the average treatment effect for the treated group.

$$\hat{t} = \frac{V_E - V_C}{\alpha^*} = \frac{X}{Y}, \tag{A3}$$

where V_E is the percentage of the experiment group that votes, V_C is the percentage of the control group that votes, and α^* is the percentage of the treatment group that was contacted. We define X as the difference in the average turnout rate among the experiment group and the treatment group, and Y as the observed contact rate for the experimental group. From equations (A1) and (A2), the expected value of X, denoted by μ_x, equals $P_E - P_C$, and the expected value of Y, denoted by μ_y, equals α. The statistical properties of the ratio (eq. [A3]) can be established using conventional approximation techniques for nonlinear functions (Rice 1995, pp. 206–8). The expected value can be approximated by dividing the expected value of the numerator by the expected value of the denominator:

$$E(\hat{t}) = \frac{P_E - P_C}{\alpha} = \frac{\mu_x}{\mu_y} = t. \tag{A4}$$

The variance is approximately equal to

$$\mathrm{Var}(\hat{t}) = \frac{1}{\mu_y^2}\left(\sigma_x^2 + t^2\sigma_y^2 - 2t\sigma_{xy}\right), \tag{A5}$$

where σ_x^2 is the variance of X, σ_y^2 is the variance of Y, and σ_{xy} is the covariance of X and Y. For purposes of statistical tests, the central limit theorem can be used to show that the distribution of t is approximately normally distributed, with mean and variance given by equations (A4) and (A5). Sample values are substituted into equations (A4) and (A5) to obtain an estimate of the expected value and variance.

References

Adams, William C., and Dennis J. Smith. 1980. "Effects of Telephone Canvassing on Turnout and Preferences: A Field Experiment." *Public Opinion Quarterly* 44:389–95.

Angrist, Joshua D., Guido W. Imbens, and Donald B. Rubin. 1996. "Identification of Casual Effects Using Instrumental Variables." *Journal of the American Statistical Association* 91 (June): 444–55.

Eldersveld, Samuel J. 1956. "Experimental Propaganda Techniques and Voting Behavior." *American Political Science Review* 50 (March): 154–65.

Gerber, Alan S., and Donald P. Green. 2000a. "Do Phone Calls Increase Turnout? A Field Experiment." Working Paper no. 2000.02. Yale University, Institution for Social and Policy Studies.

———. 2000b. "The Effects of Canvassing, Telephone Calls, and Direct Mail on Voter Turnout: A Field Experiment." *American Political Science Review* 94:653–63.

Greenwald, Anthony G., Catherine G. Carnot, Rebecca Beach, and Barbara Young. 1987. "Increasing Voting Behavior by Asking People If They Expect to Vote." *Journal of Applied Psychology* 72:315–18.

Kramer, Gerald H. 1970. "The Effects of Precinct-Level Canvassing on Voting Behavior." *Public Opinion Quarterly* 34 (Winter): 560–72.

Miller, Roy E., David A. Bositis, and Denise L. Baer. 1981. "Stimulating Voter Turnout in a Primary: Field Experiment with a Precinct Committeeman." *International Political Science Review* 2:445–60.

Morwitz, Vicki G., Eric Johnson, and David Schmittlein. 1993. "Does Measuring Intent Change Behavior?" *Journal of Consumer Research* 20:46–61.

Reams, Margaret A., and Brooks H. Ray. 1993. "The Effects of Three Prompting Methods on Recycling Participation Rates: A Field Study." *Journal of Environmental Systems* 22(4): 371–79.

Rice, John A. 1995. *Mathematical Statistics and Data Analysis*. Belmont, CA: Duxbury.

Rosenstone, Steven J., and John Mark Hansen. 1993. *Mobilization, Participation, and Democracy in America*. New York: Macmillan.

Spangenberg, Eric R., and Anthony G. Greenwald. 1999. "Social Influence by Requesting Self-Prophecy." *Journal of Consumer Psychology* 8(1):61–89.