GET OUT THE VOTE

HOW TO INCREASE VOTER TURNOUT

By Donald P. Green and Alan S. Gerber

978-0-8157-4063-6 | Paperback | December 2023 | $29.95USD
978-0-8157-4064-3/e | Book | December 2023 | $28.00USD

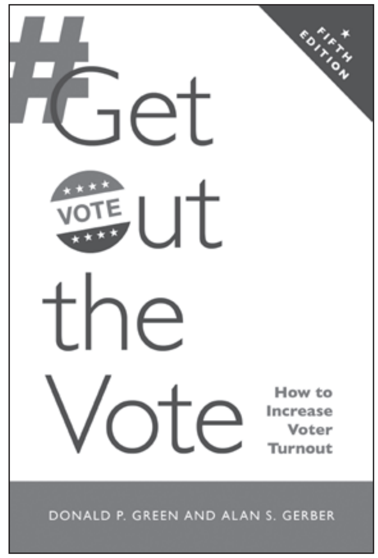

Conducting Your Own GOTV Experiment

The experiments described in the latest edition of *Get Out The Vote: How to Increase Voter Turnout* by no means exhaust the range of possible questions about voter mobilization tactics and how to make them more effective. While reading this book, you may have thought of some experiments of your own. You may be running for office, for example, and wondering how you can use the lessons learned from your current campaign to improve the efficiency of subsequent campaigns. Or perhaps you remain unpersuaded by the experimental results presented in this book and want to see for yourself whether they hold up.

With a bit of planning, you should be able to put a meaningful experiment of your own into place. You do not need an advanced degree in social science, but it is important to proceed methodically in order to guard against the problems that sometimes confront this type of research. Here is a brief overview of how experiments are designed, conducted, and analyzed.

Spell Out Your Hypothesis

A useful first step in any experimental project is to write in a single sentence the research question that you will be addressing. Doing so will force you to clarify what the treatment is and who the subjects are: for example, "Among voters with Latino surnames, Spanish-language mailers increase turnout more than English-language mailers," or "Candidates are better able to mobilize voters through door-to-door canvassing than are the activists who work for candidates," or "Election Day rides increase voter turnout among those who live more than a mile from their polling place."

Define Your Target List

Create a list of voters who will be targeted for the GOTV intervention you are planning. Depending on how your campaign is organized, the list of targets may be individual people or places (such as voting precincts). For example, your list may consist of Latino registered voters in Fresno, California. If there are some areas or people that your campaign absolutely must treat, exclude them from the list to be randomized. These must-treat observations fall outside your experiment.

Determine How Many People (or Places) You Wish to Assign to the Treatment and Control Categories

The more units you put into your experiment, the more precise your results will be. It's usually a good idea to split the target list in half so that half are assigned to the treatment group and the other half are assigned to the control group. However, do not assign more people to the treatment group than you have the resources to treat. It is more important that you apportion your experimental groups so that your contact rate in the treatment group will be as high as possible. Suppose, for example, that you are conducting an experiment to evaluate the effectiveness of GOTV lawn signs. Each treated precinct will receive forty signs. Your funds and people-power allow you to place signs in a total of ten precincts. If there are thirty precincts in your constituency, assign ten to the treatment group and twenty to the control group. Do not assign fifteen to the treatment group if you have the resources to only treat ten.

Divide the List into Treatment and Control Groups

Creating random treatment and control groups is easy to do with a spreadsheet program. Box A-1 walks you through the steps of randomly sorting and subdividing your target list.

Random sorting may also be useful when you are apprehensive about excluding a control group. Suppose you are planning to encourage turnout by calling voters. Sort your phone list so that it is in random order and then begin calling, starting at the top and working your way down. When your campaign is over, any names that you did not *attempt* to call represent your control group. (Note that the control group includes only the names that you did not attempt to call; the names that were called but not reached are still considered part of the treatment group.) Nothing about the phone bank's

Box A-1. Random Assignment

Random assignment to treatment and control groups is easily accomplished with a spreadsheet program, for example, the freeware available at OpenOffice.org. First, open the spreadsheet containing the list of voters. Second, place your cursor on an empty cell in the spreadsheet and click the equal sign to call up the equation editor. Enter the formula RAND() and hit enter. A random number between 0 and 1 should appear in the cell. Third, copy and paste this random number into a column beside the columns in your data set. Every row in your data set should now contain a random number. Fourth, highlight all the columns in your data set and click DATA > SORT. A box should appear asking you to indicate which column(s) to sort by. Choose the column that corresponds to the column of random numbers just generated and click OK. Finally, having sorted the data in random order, add a new column to your data set. This new column will indicate whether each person is assigned to the treatment group. If you want five hundred people in your treatment group, put the number 1 into the first five hundred rows and the number 0 into the remaining rows.

execution has changed (so long as they make the calls in order), but now the phone bank's efforts become amenable to rigorous evaluation if it is unable to call all of the names on your target list. This approach is ideal for evaluating campaigns whose goals outstrip their resources.

Check the Randomization

Random assignment should, in principle, create treatment and control groups that have similar background characteristics. To ensure that you have conducted the randomization properly, check to see that the treatment and control groups have approximately the same rate of voter turnout in some prior election (before your intervention). If this information is not available, check to see that the average age in both the treatment group and the control group are approximately the same. If the treatment and control groups differ appreciably, you may have made a computer error. Check your work and redo your randomization.

Administer the Treatment to the Treatment Group Only

Be vigilant about adhering to the treatment and control assignments. Do not contact anyone on the control list! The easiest way to keep the experiment from going awry is to release the treatment names only to the phone bank or direct mail vendor. Send leafleteers and door-to-door canvassers out with the names of people in the treatment group only and remind them not to knock blindly on every door.

Maintain Records of Who Was Contacted

Maintain records of who was contacted, even if that person is someone in the control group who was contacted by mistake. You will need this information to calculate the effect on those contacted of your intervention.

Archive Your Campaign Materials

Once you have successfully carried out this protocol, archive your campaign materials and write a brief description of the experimental procedures. While you wait for the registrar of voters to furnish voter turnout information (and before you forget the details), write up the description of the experiment and how it was conducted. Create a physical or electronic archive of your campaign materials—your scripts, recorded messages, mailings, and so forth. Make sure to collect whatever data you need from canvassers (for example, walk sheets with contact information) and callers before these materials disappear.

Calculate Turnout Rates

When voter turnout data become available, calculate the turnout rate of the people in the treatment and control groups. Remember, your treatment group consists of those individuals assigned at the outset to the treatment group, regardless of whether you were able to contact them.

Analyze the Results

Analyze the experimental results. The difference in turnout between the original treatment and control groups—ignoring for the moment whether persons in the treatment group were actually treated—tells you quite a lot. (Note: Do not discard people in the treatment group who weren't home or were found to have moved. Stick to your original list of random assignments. The only people who can be deleted from an experiment are those in the treatment and control groups who voted before your treatment launched.) If the turnout rate in the assigned treatment group is higher than in the control group, your intervention seems to have worked. This difference is known as the intent-to-treat effect, because it compares those you intended to treat with those you intended to leave alone.

Next, calculate the contact rate. Divide the number of people contacted in the treatment group by the number of people assigned to the treatment group. If your contact rate is less than 100 percent, divide the intent-to-treat effect by the contact rate. The resulting number indicates the effect of the treatment on those who were contactable. For example, if your control group votes at a rate of 42 percent, your treatment group votes at a rate of 47 percent, and your contact rate is 50 percent, the effect of the treatment on contactable subjects is estimated to be (47 – 42)/0.5, or 10 percentage points.

Perform a Statistical Analysis

A bit of statistical analysis can indicate how likely it is that the difference you are seeing was produced by chance rather than your intervention. Why worry about chance? Because sometimes, just by luck of the draw, you may underestimate or overestimate the effectiveness of the treatment. We have developed some (free) web tools to walk you through this process as painlessly as possible and give you some pointers about how to interpret the numbers (see box A-2).

Present Your Results

Once you have conducted your experiment, try to present it in a manner that enables it to contribute to the accumulation of scientific knowledge about campaigns. Here is a brief checklist of what to report to your audience.

- ✓ Describe the experimental setting. When and where was the campaign conducted? What other kinds of campaign activity or public debate might voters have been exposed to?
- ✓ Describe the experimental treatments. What were the interventions? When were they deployed and by whom? Present the phone or canvassing scripts. Show photos of the mailings or emails.
- ✓ What were the experimental groups, and from what population were they drawn? Describe the randomization procedure used to assign the groups. Describe the number of observations assigned to each experimental group. Show whether, as expected, the treatment and control groups have similar background attributes, such as average age or past voting rates.
- ✓ Were the treatments successfully administered to all of the people who were supposed to receive them? If not, what proportion of the treatment group actually received the treatment? Did the control group inadvertently receive your treatment? If so, what proportion of the control group was inadvertently treated?
- ✓ Briefly explain the source of your voter turnout information and how you dealt with cases in which people in the original treatment groups were not found on the rolls of voters. (Sometimes those not found are classified as nonvoters, and sometimes they are excluded from the analysis altogether.) Verify that the treatment and control groups have similar rates of voter registration, and present the voting rates for each experimental group, including the number of observations used to calculate each voting rate.

Box A-2. Simplified Web-Based Software for Analyzing Experimental Data

For the convenience of first-time experimenters, we helped to create free web-based software that reads experimental results and generates a statistical analysis. Go to https://egap.shinyapps.io/gotv-app or google "GOTV shiny app." You supply six numbers: the number of people that you (1) assigned to the treatment group, (2) assigned to the control group, (3) successfully treated in the treatment group, (4) inadvertently treated in the control group, (5) found to have voted in the treatment group, and (6) found to have voted in the control group.

To see how the program works, suppose you wish to analyze results from Melissa Michelson's door-to-door canvassing experiment in Dos Palos, California. Prior to the 2001 election, she assigned 466 people with Latino surnames to the treatment group and 297 Latinos to the control group. Of the people in the treatment group, 342 were successfully contacted. No one in the control group was contacted. In the treatment group, 86 people voted, whereas 41 people voted in the control group. The six inputs are, therefore, 466, 297, 342, 0, 86, and 41.

After entering these numbers in the appropriate boxes, you will see output that summarizes the research findings and estimates the size and precision of the treatment effects. Check the statistical summary that appears in the middle of the page to ensure that you have entered the data correctly. The computer will summarize the voting rates and contact rates based on the numbers you provided. Next, examine the intent-to-treat estimate. This number is calculated by subtracting the voting rate in the control group from the voting rate in the treatment group. In this example, the intent-to-treat estimate is 4.7, suggesting that assignment to the treatment group raised turnout 4.7 percentage points. Below this figure is the standard error of the estimated intent-to-treat effect. The larger this number, the more uncertain the intent-to-treat estimate. The average treatment effect among those who were reachable by canvassers is estimated by dividing the intent-to-treat estimate (4.7) by the contact rate (0.73), which produces the number 6.3. Door answerers who received the treatment became 6.3 percentage points more likely to vote. The uncertainty of this estimate is measured by its standard error, 3.7.

Finally, the statistical software makes three useful calculations. The first is the 95 percent confidence interval, which spans from –0.9 to 13.5. This type of interval has a 95 percent chance of bracketing the true average treatment effect among door answerers. The second calculation is the one-tailed significance of the estimated treatment effect. When conducting GOTV experiments, it is conventional to expect turnout to rise as a result of the treatment. The so-called null hypothesis is that the treatment failed to increase turnout. The one-tailed significance level states the probability of obtaining an estimate as large as the observed intent-to-treat effect simply by chance. When this probability is below 0.05, as is the case here, the estimate is conventionally dubbed "statistically significant." Naturally, if the experiment were repeated, the results might come out differently. The "power" of an experiment describes the probability that it would produce a statistically significant estimate assuming the observed intent-to-treat effect were the true effect of assignment. In this case, Michelson's experiment has a 51 percent probability of rejecting the null hypothesis given that the intent-to-treat effect is 4.7 percentage points.

These are the essential ingredients of any experimental write-up. Obviously, if you have the statistical skills, you can go into much more detail, but do not skip directly to a complex statistical analysis without first walking the reader through the fundamental facts of the experiment mentioned above. One of the most attractive features of experimentation is that it lends itself to a straightforward and transparent style of presentation.

Defending Your Experimental Design

Crafting an experiment requires planning and supervision, yet experience has shown that any energetic and well-organized researcher can learn to do it. For decades, we have taught courses and workshops on experimental design, and many of the participants have gone on to conduct clever and well-executed voter mobilization studies. One participant got a team of friends together to make GOTV calls on behalf of a candidate for governor; another designed and distributed nonpartisan direct mail; another organized a large-scale precinct walking campaign; and another mobilized her Facebook friends. One of the most gratifying aspects of writing a book like this is hearing about the imaginative experiments that readers have conducted.

Experimental research requires a fair amount of effort, but the biggest hurdle in conducting a successful experiment is conceptual. To design and execute a randomized experiment, you must first understand why a randomly assigned control group is essential. At every stage in the process of executing an experiment, people will complain about your insistence on extracting a control group from the target list of voters. They will propose instead that you simply consider the people the campaign happened to treat as the treatment group and the people the campaign failed to contact as the control group. In other words, they will propose that you forget about random assignment and just let the campaign workers treat whomever they please. You will have to convince these skeptics that a randomly assigned control group is indispensable.

When you balk at their alternative research proposal, the skeptics may try to reassure you by suggesting that you can make the research design serviceable by comparing a select group of people who did or did not receive the call. The select group they have in mind are people who are similar in age, party attachment, and voter turnout rate in previous elections. Do not give in. What they are suggesting is a flawed research design. Here's why. Even if the people who were treated have the same observed attributes as the people who were not treated, there is no guarantee that their unobserved attributes are the same. Learning that a person is home when you call reveals something about his or her likelihood of voting (the person has not moved, is not dead, and is willing to pick up the phone when a stranger calls). Even if your phone calls had no effect on turnout, the people you reached will vote at higher rates than the folks you could not or did not reach.

To drive this point home to our academic colleagues, we conducted an experiment a couple of days before the November 2004 election. We called 15,000 registered voters in Illinois and urged them to buckle their seat belt when driving during the holiday season. Not a word was said about voting or politics. Just seat belts. Obviously, the true effect of this call on voter turnout was zero. But sure enough, the people who received the buckle-up message voted at a rate that was 5 percentage points higher than those who were not called, even though the two groups shared exactly the same pattern of voter turnout during the previous decade. (See Kevin Arceneaux, Alan S. Gerber, and Donald P. Green, "A Cautionary Note on the Use of Matching to Estimate Causal Effects: An Empirical Example Comparing Matching Estimates to an Experimental Benchmark," *Sociological Methods and Research,* vol. 39 (2010): 256–82.)

Why did an irrelevant "buckle up for safety" message seem to increase turnout? Because this comparison excluded the dead, the moved, and the unfriendly from the treatment group but did not exclude them from the control group. The point here is that if your research design is flawed, you risk generating deeply misleading conclusions. Don't be fooled into thinking that you can patch up a flawed design by focusing on people who share the same observed characteristics.

What is the right way to test the mobilizing effects of a phone call? Randomly assign voters to receive a call or not. Keep track of the fraction of the treatment group that actually received the calls. Compare the voting rates among those originally assigned to the treatment and control groups and divide by the fraction of the treatment group that received the calls. (Or just use the web-based software described in box A-2.) Because you are comparing the randomly assigned treatment and control groups, the groups have the same expected voting rate, and there is no bias in favor of finding that phone calls work (or don't work).

Another common concern among those who resist experimental research is that a campaign cannot risk extracting a control group from its list of targeted voters. This concern is often based on an unrealistic sense of how effective their mobilization campaign is likely to be (or how close the election is likely to be). If, for example, a state legislative campaign has resources to target 15,000 voters and you remove a control group of 1,500, two things happen. First, the campaign can reallocate its resources to the 13,500 and treat some of them more intensively. Second, your control group will cost the campaign votes in proportion to its effectiveness. Suppose the campaign does an outstanding job of mobilizing and persuading its base of supporters. Suppose the control group, if left untreated, votes at a rate of 45 percent and votes four to one in favor of your candidate. Had the control group been treated, this group would have voted at a rate of 50 percent and supported your candidate by a five-to-one margin, which translates into an additional ninety-five votes. In the worst-case scenario, the campaign reaps no additional votes from the extra resources it reallocates to the treatment group while losing those ninety-five votes. That is probably a much smaller number than the campaign envisioned when it worried about extracting a control group. These kinds of calculations help campaign managers realize that the small, short-term risks associated with an experiment are outweighed by the long-term benefits of acquiring knowledge about mobilizing and persuading voters.

Brookings Institution Press®