The Effects of Large-Scale Campaigns on Voter Turnout: Evidence from 400 Million Voter Contacts

Ryan D. Enos and Anthony Fowler[1]

Abstract

To what extent do political campaigns mobilize voters? Despite the central role of campaigns in American politics and despite many experiments on campaigning, we know little about the aggregate effects of an entire campaign on voter participation. The 2012 presidential election provides a ripe opportunity to address this question, because the scale of mobilization was unprecedented and because campaigning varied discontinuously across state boundaries. We combine individual-level data on every eligible voter in the U.S. with inside information on the activities of the presidential campaigns to estimate the aggregate effects of a large-scale campaign. Even though single voter contacts only minimally increase turnout, an entire campaign consisting of many contacts increases participation substantially. We estimate that the presidential campaigns increased turnout by more than 10 percentage points among targeted subgroups, indicating that modern campaigns can significantly alter the size and composition of the voting population.

What are the consequences of political campaigns for mass participation? Despite extensive research on individual campaign tactics and despite cumulative campaign spending now totaling billions of dollars, we do not yet understand the aggregate effects of an entire campaign on voter participation. In this paper, we take advantage of geographic discontinuities in campaigning and use new data on every eligible voter in the United States to show that campaign efforts have large effects on voter turnout, significantly altering the size and shape of the voting population. Our findings demonstrate that modern campaigns are a vital force in American democracy, causing the participation of millions of voters who would otherwise not have participated.

Campaigns can be roughly divided into a "ground campaign," consisting of personal contact such as phone calls, direct mail, and door knocking, and an "air campaign," consisting of television and internet advertising. Personally-targeted techniques are now a major part of modern campaigns, and in 2012 alone, the two presidential campaigns attempted more than 400 million individual voter contacts.[2] Based on prior evidence, we expect the ground campaign to have a larger effect on voter turnout, so we focus on that effect in this paper, while also demonstrating that, consistent with previous research, the effects of television and internet advertising are small to negligible.

Researchers have repeatedly shown that individual voter contacts can increase political participation, but the substantive size of these effects is usually small. Green, McGrath, and Aronow (2013) pool evidence from more than 200 get-out-the-vote (GOTV) experiments—published and unpublished—to show that, on average, door-to-door canvassing increases turnout by 1.0 percentage point, direct mail increases turnout by 0.7 percentage points, and phone calls increase turnout by 0.4 percentage points.[3] Noting the miniscule size of these average effects, one might

---

[2] Data obtained from the Romney campaign indicates that the campaign attempted approximately 225 million voter contacts. If Obama's efforts were similar, there were more than 400 million voter contact attempts in the 2012 presidential campaign.

[3] These are estimates of the average intent-to-treat effects, obtained through private correspondence with Green, McGrath, and Aronow (2013).

conclude that campaigns play a negligible role in elections. However, these estimates arise from single-shot experiments where a single intervention is compared to no intervention. As such, we have little knowledge of what the total effect of a campaign might be, especially large-scale, modern campaigns which attempt many contacts for each voter. Few experiments have assigned subjects to receive multiple interventions (e.g. Cardy 2005; Gerber and Green 2000; Michelson, Bedolla, and Green 2007; Ramirez 2005), and none have come close to matching the level of a full-scale campaign.

Given the current evidence, we have virtually no way of knowing what the effects of a large-scale campaign may be. At one extreme, the effects of many voter contacts could be counterproductive by fatiguing potential voters and turning them away. At the other extreme, the effects of many interventions could be synergistic and multiplicative—perhaps some citizens are insensitive to single interventions but can be mobilized by several. Somewhere in between, multiple interventions could have additive effects or the returns could be positive but diminishing. Simply put, despite extensive research on campaigning, we still know little about the aggregate effects of campaigns on political participation.

In this paper, we exploit the 2012 presidential campaign to assess the aggregate effects of a large-scale campaign on the size and composition of the voting population. We take advantage of variation in ground campaigning across state boundaries, extensive information on Romney and Obama campaign tactics, and detailed information on every voter in the United States to estimate the effects of the entire campaign. Our results suggest that the aggregate mobilizing effect of a presidential campaign is quite large. We estimate that the 2012 campaign increased aggregate turnout by approximately 7 percentage points in the most heavily targeted states, and this effect is greater than 10 percentage points for targeted subgroups. In total, our estimates imply that 2.6 million

individuals voted who would have otherwise not participated in the absence of the ground campaign, suggesting that modern campaigns play a significant role in mass participation.

Further analyses suggest that the effects of multiple voter contacts are approximately additive. Despite the small marginal effect of individual voter contacts, the aggregate effect of multiple contacts is quite large. The effects are largest among predicted partisans who were more likely to have been targeted by the campaigns, suggesting that campaigns may produce a more partisan, polarized set of voters. Additionally, the GOTV effects are nearly identical for Democrats and Republicans, suggesting that accounts of Barack Obama's advanced campaign relative to Mitt Romney's may have been overstated. Lastly, because ground campaigning has increased significantly in the 21st century, we may be witnessing a "Victory Lab Revolution" whereby the nature of presidential campaigns and voting populations has changed dramatically over a short time period.

In separating the ground campaign from other factors that influence participation, we rely on the fact that the presidential campaigns targeted individuals in battleground states, while ignoring demographically-similar individuals immediately across the state line who received the same television advertisements and media coverage. We also look at the specific content of television advertisements and argue that they are unlikely to have been more effective in battleground states than non-battleground states. We further control for the effects of internet-based advertising by using original data on digital campaign efforts across states. Importantly, we anticipate objections that living in a battleground state may have an independent effect on turnout even in the absence of campaigning. Multiple, independent sources of evidence indicate that these concerns are minimal. For example, placebo tests using survey data confirm that battleground and non-battleground residents were similar in terms of their interest in politics and their intentions to vote before the general election campaign. We discuss all of these issues in more detail later in the paper.

The subsequent sections describe our data and empirical strategy. Next, we present placebo tests assessing the validity of our identifying assumptions. Then, we present our main estimates of the aggregate effect of ground campaigning in 2012. The next two sections explore variation in the effect of ground campaigning across different types of individuals and previous presidential elections. Then, we separate the effects of ground campaigning from those of internet advertising. Lastly, we present a detailed discussion of challenges and limitations, paying particular attention to concerns about the independent effects of battleground states in the absence of campaigning. We conclude by discussing the implications of our results for the study of electoral politics and for future elections and campaigns.

Data

Our primary data comes from Catalist, a for-profit data vendor that maintains records of every eligible voter in the United States. We obtained detailed tabulations which allow us to reconstruct an anonymous version of their database. Therefore, for virtually every eligible voter in the U.S., we know their gender, predicted race, predicted income range, state, media market, and validated voter turnout in 2012, 2010, 2008, 2006, and 2004.[4] We analyze this data in conjunction with information that we obtained through interviews and private correspondence with numerous high-level operatives and strategists from the Obama and Romney campaigns. Through these interviews, we obtained specific information regarding the campaigns' overall strategies, the extent of voter mobilization efforts, the places that were targeted, and the types of individuals that were

---

[4] We do not know the voters' names, addresses, or any other identifying information. Catalist generates and maintains their database through state voter files and consumer records. There are surely some eligible but unregistered voters who are missed by Catalist and some ineligible individuals who are mistakenly included in their database. However, the numbers in their database closely match estimates of the eligible voting population in each state, suggesting that these concerns are minimal. For more information on Catalist data and examples of political science research employing this data source, see Ansolabehere and Hersh (2012, 2013) and Hersh (2011, 2013).

targeted. We also obtained data from the Obama campaign on all internet-based advertising, including the geographic location, type of advertisement, cost, volume, and advertising provider.

Empirical Strategy: Comparisons within Media Markets but across State Boundaries

Our basic empirical strategy relies on the discontinuous changes in presidential campaigning across state boundaries. Because of the Electoral College, the returns to campaigning vary across states. For example, the presidential campaigns made no effort to mobilize voters in Wyoming and Massachusetts, where Romney and Obama were essentially guaranteed to win, respectively, regardless of campaign effort. However, both teams campaigned heavily in Colorado and New Hampshire, where the outcome of the race was uncertain and these states could potentially tip the nationwide election result. In this sense, we follow previous studies that have exploited this variation across states in studying campaigns or voting behavior (e.g., Gerber et al. 2009; Huber and Arceneaux 2007; Kim, Petrocik, and Enokson 1975; Krasno and Green 2008; Hill and McKee 2005; Wolak 2006; Lipsitz 2008).

Of course, battleground states differ from non-battleground states in many ways other than the extent of voter mobilization. Most notably, voters in battleground territory receive more television advertising and more intense news coverage of the election. In order to separate the effects of television advertising and other mass media from ground campaigning, we take advantage of media markets that span state boundaries and conduct all of our comparisons within media market, so that we can be sure that the "treated" and "control" groups received the same level of television ads and news coverage. For example, we compare voters within the Denver media market that live in either Colorado or Wyoming, and we compare voters within the Boston media market that live in either New Hampshire or Massachusetts. In some specifications, we also condition on individual characteristics—gender, race, and predicted income—to improve precision and account

6

for further differences between citizens in different states. In additional specifications, we control for whether there is a competitive Senate or gubernatorial election in a state, the particular election laws and voting methods in a state, and the level of internet-based advertising in each state. The results are nearly identical across each specification. Our basic strategy has similar elements to those of Huber and Arceneaux (2007) and Krasno and Green (2008) who compare battleground and non-battleground media markets within the same state to assess the impact of television advertising. Our design is essentially the inverse; we compare battleground and non-battleground states within media markets to assess the impact of ground mobilization.[5]

This strategy provides unbiased estimates of the aggregate effects of the 2012 ground campaign under the assumptions that, on average, voters in the same media market in battleground and non-battleground states 1) receive the same boost in turnout from mass media, 2) have the same underlying propensity to vote and 3) are subject to the same boost in turnout from forces other than campaigning such as psychological and social pressures not induced by the campaign. Assumption 1 is defensible because voters in the same media market see the same television, because we are able to control for internet advertising, and because few television ads contain state-specific content. Assumption 2 is supported by idiosyncratic factors determining which states are targeted. We further justify assumptions 2 and 3 by referencing previous research, conditioning on observable covariates, running placebo tests, and showing that our estimates vary in predictable ways. We discuss each of these assumptions and their potential violations in more detail later in the paper.

Our design requires us to categorize the intensity of campaign activity across states, and we take three different approaches. First, we code a binary variable, *Battleground*, indicating whether intense campaigning occurred in a particular state. This variable takes a value of 1 in Colorado,

---

[5] In the Appendix, we replicate the basic design of Krasno and Green (2008) to estimate the effect of television advertising on turnout. Consistent with evidence from previous elections, we find that television ads had no detectable effect on turnout in 2012.

Florida, Iowa, Nevada, New Hampshire, North Carolina, Ohio, Virginia, and Wisconsin. These states received the vast majority of media attention, television ads, and campaign field offices (Sides and Vavreck 2013). These were also the states identified by news organizations and political consultants as the key swing states.

For our second approach, we divide the states into four categories according to an internal categorization used by the Obama campaign.[6] The first tier, receiving the greatest degree of campaign effort, includes Colorado, Iowa, Nevada, New Hampshire, Ohio, and Virginia. The second tier, also receiving significant campaign effort but less than the first tier, includes Pennsylvania, Florida, and North Carolina. The third tier, receiving some campaigning but less than the first two tiers, includes Wisconsin, Minnesota, Washington, Oregon, Michigan, and New Mexico. The remaining states in none of the three tiers received virtually no attention.

For our third approach, we utilize data on the number of voter contacts attempted by the Romney campaign in each state. Specifically, we obtained an internal document from the campaign summarizing the number of attempted phone calls, direct mail deliveries, and door-to-door canvassing efforts in each state.[7] Using these numbers, we construct a continuous scale indicating the average effort level of the Romney campaign in each state. However, not all types of contacts indicate equal effort. In determining the relative weights that we should give to each type of voter contact, we relied upon estimates of the cost of each mobilization method. Green and Gerber (2008) estimate that the cost of phone calls, direct mail, and door knocks are approximately 30 cents, 75 cents, and 3 dollars, respectively, per attempt. These estimates allow us to approximate the amount of effort per state in terms of the hypothetical cost of those campaign efforts per eligible voter.

---

[6] Internal Obama campaign document from May 2012, "Developing State Priorities," obtained through personal communication with Obama campaign staff.

[7] Internal Romney campaign document, "Romney Voter Contact Summary," obtained through personal communication with Romney campaign staff.

Twenty-eight states plus Washington, D.C. received no effort whatsoever, while many other states received significant effort. Nevada received the highest level of effort with approximately 5 phone calls, 0.7 pieces of direct mail, and 0.5 door knocks attempted for every eligible voter in the state. This amounts to an imputed cost of $3.40 per person in Nevada. New Hampshire, Virginia, Iowa, Ohio, and Colorado were next, with costs exceeding 2 dollars per person. Wisconsin, Florida, and North Carolina followed with costs exceeding 1 dollar per person. We rescale this variable to range from 0 to 1 and call it *RomneyEffort*. Comparing the emphasis put on different states by the Romney campaign to the internal categorization used by the Obama campaign, the campaigns appear to have prioritized the states in roughly the same way.

Figure 1 presents a map of *RomneyEffort* across states and media markets. Media markets spanning states with different levels of *RomneyEffort* are surrounded by thick black lines. Portions of each media market are shaded by the level of Romney effort in that state. So, for example, the Reno media market is much darker in Nevada than California, representing the variation in *RomneyEffort* across those states. States with no effort or only minimal effort are the lightest shade. The media markets that do not span state lines with varying levels of *RomneyEffort* are indicated by diagonal lines. The voters in these areas do not enter our sample. This map highlights the sample of voters and sources of variation that we utilize in the regressions below to estimate the effects of campaigning.

[Figure 1]

We specify three approaches to estimating the average effect of the ground campaign on voter turnout, using each of the measures of campaign effort described above. Employing OLS, we estimate the following three regressions:

(1) $\text{Turnout}_{ijk} = \beta_0 * \text{Battleground}_j + \gamma_k + \delta X_i + \varepsilon_{ijk}.$

(2) $\text{Turnout}_{ijk} = \beta_1 \text{Tier1}_j + \beta_2 * \text{Tier2}_j + \beta_3 * \text{Tier3}_j + \gamma_k + \delta X_i + \varepsilon_{ijk}.$

(3) $Turnout_{ijk} = \beta_4*RomneyEffort_j + \gamma_k + \delta X_i + \varepsilon_{ijk}.$

Turnout$_{ijk}$ is an indicator for voter turnout for individual i in state j and in media market k. As described previously, *Battleground*$_j$, *Tier1*$_j$, *Tier2*$_j$, and *Tier3*$_j$ are binary variables indicating whether state j falls into that particular category of campaign effort, and *RomneyEffort*$_j$ is a continuous variable ranging from 0 to 1 indicating the extent of voter mobilization effort by the Romney campaign in state j. γ_k represents media market fixed effects which account for the fact that individuals in different media markets received different levels of news coverage and television advertising. These fixed effects allow all of our inferences to be made within media markets but across state boundaries. *X*$_i$ represents a vector of individual characteristics—race, income, and gender—that we include to account for the potential that even within media markets, there may be some systematic demographic differences across state boundaries.

β_0 in Equation 1 represents the average difference in turnout between battleground individuals and demographically similar non-battleground voters within the same media market. Under our assumptions, this quantity represents the average effect of the campaigns' ground mobilization efforts for all citizens in the battleground states. Similarly, β_1, β_2, and β_3 from Equation 2 represent the average effect of being in the first, second, or third tier of campaign effort, respectively, relative to receiving virtually no campaign effort. β_4 from Equation 3 represents the average effect of going from a state with virtually no mobilization effort to the state with the greatest voter mobilization effort. To be clear, all three specifications are intended to measure the aggregate effects of all GOTV campaigning by Obama, Romney, and independent groups. While we use information on mobilization efforts from the Obama and Romney campaigns in separate regressions, both campaigns prioritized states in virtually the same way, so all regressions inform us about the average effects of all campaigning. These differences between different regression results

tell us nothing about the relative effectiveness of Obama and Romney. Later, we separately examine different subgroups of individuals in order to compare the effectiveness of the two campaigns.

Placebo Tests Using 2010 Turnout and Prior Political Interest

Before proceeding to the results, we first present placebo tests in order to assess the validity of our empirical design. We implement our designs described above in Equations 1-3, using three different placebo outcomes. In the first case, we use the same data set but code voter turnout in 2010 as the outcome variable. Because there was no presidential campaigning in 2010, there should be minimal differences between battleground and non-battleground voters in the same media market if these two groups are comparable for our purposes. We also utilize survey data from the Cooperative Campaign Analysis Project (CCAP) to conduct two additional placebo tests using measures of political interest and intention to vote before the general election campaign got underway. These three placebo outcomes assess the comparability of battleground and non-battleground voters within the same media markets, and they help to rule out many potential alternative explanations and potential sources of bias.

Table 1 shows the results of our placebo tests using voter turnout in 2010. Columns 1 and 2 show results for Equation 1 where we compare battleground and non-battleground states. Columns 3 and 4 show the results for Equation 2 comparing the 4 different tiers of states. Columns 5 and 6 show results for Equation 3 examining turnout across *RomneyEffort*. All regressions include media market fixed effects, and columns 2, 4, and 6 include individual-level controls for gender, race, and income. The sample sizes vary across the three different estimation strategies because we only include media markets for which there is variation in the treatment variable. For example, the regressions associated with Equation 1 (Columns 1 and 2) only include media markets that span both a battleground and non-battleground state. This includes approximately 50 million individuals

living in 38 media markets spanning 34 states. For Equation 2, we include all individuals living in media markets which span states in different tiers (88 million individuals in 49 media markets spanning 44 states). For Equation 3, we include all individuals living in media markets which span states receiving different levels of effort from the Romney campaign (119 million individuals in 74 media markets spanning 48 states—Washington, DC is included as a state).

[Table 1]

Examining our specifications with no individual-level control variables (Columns 1, 3, and 5), we see no meaningful differences between individuals in targeted and non-targeted states in their turnout levels in 2010. All point estimates are statistically indistinguishable from zero, suggesting that individuals within the same media market but spanning state boundaries appear to have similar underlying propensities to vote. When we add individual control variables—gender, race, and predicted family income—to these regressions (Columns 2, 4, and 6), the point estimates are virtually unchanged because these covariates are well balanced across targeted and non-targeted states.[8] Most of the point estimates are small relative to our main results using 2012 turnout, so these differences are unlikely to meaningfully bias our subsequent results. Moreover, our subsequent results are unchanged by the inclusion of 2010 voter turnout as a control variable, further bolstering confidence in our design and results.[9] These placebo tests support the plausibility of our empirical

---

[8] All individual controls are included as dummy variables, so linearity assumptions are not a concern. For example, when we say that we include controls for race, we include binary indicators for white, black, Hispanic, and Asian individuals, leaving other races and missing race information as the omitted category. Predicted family incomes are divided into 9 categories (<5k, 5-12.5k, 12.5-20k, 20-30k, 30-40k, 40-60k, 60-100k, >100k, and missing). We also have some data on age, but we do not incorporate it into our analyses, because age data is missing for many individuals and the rate of missing data varies significantly across states.

[9] Only 1 of the 10 coefficients in Table 1 is statistically significant (the coefficient associated with Tier 2 when individual controls are included), and given the number of tests, at least one significant coefficient is expected to arise by chance. Furthermore, the statistical significance of Tier 2, rather than Tier 1, despite Tier 1 receiving greater campaign activity has no clear interpretation.

design. Voters within the same media market but spanning states that received differential levels of presidential campaigning are similar in their baseline levels of participation.

The results in Table 1 alleviate concerns that our comparison units are not fundamentally comparable in their underlying propensity to participate in elections. However, these results do not address additional concerns about factors specific to presidential elections that may lead voters in battleground states to participate more. For example, battleground voters may participate more, even in the absence of campaigning because they believe their vote is more likely to be pivotal. To address these issues, we present additional placebo tests using survey data from 2012 to measure differences between battleground and non-battleground voters in their underlying interests in voting. The 2012 CCAP surveyed approximately 1,000 different individuals in each week of 2012 before the November election. We focus on the surveys from the first 16 weeks—before the RNC declared Romney as the presumptive nominee on April 25th. During this period, there was little general election campaign activity, so any differences between voters in battleground and non-battleground states are unlikely to be attributable to campaigning.[10]

The CCAP asked two questions each week that are useful for detecting any underlying differences that might plague the results of our study. First, they asked each respondent about their level of interest in politics/current events, allowing respondents to say that they are "very much interested," "somewhat interest," "not much interested," or "not sure." Just over half of the respondents report that they are very much interested, so we code a binary variable indicating whether each respondent chose that particular option. The CCAP also asked each respondent who they would support in the presidential election. Respondents could choose among the following choices: "The Democratic Party candidate," "The Republican Party candidate," "Other," "Not

_____

[10] Obama may have already engaged in general election campaigning at this point, but this activity would create a bias toward a failure of our placebo test.

sure," or "I would not vote." We code a binary variable indicating whether a respondent reported that they would not vote.

Tables 2 and 3 below present a series of placebo regressions testing whether individuals in targeted states reported greater levels of interest in politics or greater intentions to vote before the general election campaigns got underway. If voters in targeted states are more likely to turn out even in the absence of campaigning, we would expect to find large differences in these surveys. However, if there are no meaningful differences before campaigning began, then we can more confidently attribute our subsequent results to campaigning. Because we cannot match the CCAP respondents to low levels of geography, we do not include media market fixed effects in these placebo analyses. Also, because we pool survey responses from different weeks, we include week fixed effects to account for any changes over time in these attitudinal variables. In all other respects, the regressions in Tables 2 and 3 mimic those in Table 1. For each outcome and for each measure of campaign effort across states, we present three separate regressions. The first has no controls. The second includes individual controls for age, race, and gender. The third also includes a control variable indicating whether the respondent's state held an early primary (before Super Tuesday). In total, across both tables, we obtain 30 placebo estimates. Only 2 of these estimates are statistically significant (we would expect 1.5 by chance), and they go in the opposite direction that we would expect if battleground voters were more motivated to vote independent of campaigns.[11]

[Tables 2 and 3]

In short, we find no evidence that individuals in targeted states had different turnout behaviors, intentions to vote, or levels of political interest before presidential campaigning began. These results suggest that our identifying assumptions are sound and that our subsequent results can

---

[11] Individuals living in Tier 2 states were slightly more likely to report that they will not vote compared to other states (Columns 5 and 6 in Table 3)—a result that is most likely explained by chance.

be attributed to campaigning as opposed to spurious factors and independent effects of battleground environments. We further discuss these issues at the end of the paper. In the next section, we present our main results where we repeat the same specifications using voter turnout in 2012 as the outcome of interest.

Effects of the Ground Campaign in 2012

Table 4 presents our main results using 2012 turnout as the dependent variable. These results measure the aggregate effect of the ground campaign in 2012. Columns 1-3 show results for Equation 1, using battleground states as the key independent variable, Columns 4-6 show results for Equation 2, using state tiers as the key independent variable, and Columns 7-9 show results for Equation 3, using *RomneyEffort* as the key independent variable. Columns 1, 4, and 7 include no controls other than media market fixed effects. Columns 2, 5, and 8 also include individual controls for gender, race, and income. Columns 3, 6, and 9 include an additional control for turnout in 2010. The inclusion of control variables has virtually no impact on our estimates, indicating that these covariates are well-balanced across state lines. These results are also unchanged if we include a control variable indicating which states had a competitive gubernatorial or senatorial race in 2012.[12] Additional regressions controlling for different election laws and voting methods across states also produce nearly identical results.[13]

Conditional on the media market and other factors, voters in battleground states were 4 percentage points more likely to turn out, suggesting that ground campaigning increased

---

[12] We code competitive Senate or gubernatorial races as those classified beforehand as "toss up" races by Real Clear Politics. These races took place in Indiana, Massachusetts, Montana, Nevada, North Dakota, Virginia, Wisconsin, and Washington.

[13] Data on voting methods was collected from the National Council on State Legislatures. We code dummy variables which correspond to their categorizations (no-excuse absentee voting, early voting, early voting and no-excuse absentee voting, all-mail voting, and no early voting and excuse required for absentee).

participation in battleground states by 4 percentage points, on average. Comparing the top tier of campaign priority to non-targeted states, we see that turnout was 6-7 percentage points higher in the more targeted regions. This difference is 2-4 percentage points for the second tier of states and 1-4 percentage points for the third tier. As expected, for each specification, the estimated effect of Tier 1 residence is greater than the estimated effect of Tier 2 residence, which is greater than the estimated effect of Tier 3 residence, which is greater than zero. Similarly, in the last three columns, going from states with no campaign effort to those with maximal campaign effort corresponds to a 7-8 percentage point increase in turnout. With the exception of the Tier 3 estimates, all point estimates are substantively meaningful and statistically significant ($p < .01$ with standard errors clustered by state). Our aggregate results are consistent with large, powerful mobilizing effects of the presidential campaigns that increase with the extent to which a state was targeted. Voter mobilization appears to have significantly increased turnout in the most targeted states.

[Table 4]

Our data and results allow us to say something about the marginal effects of multiple campaign contacts. The Romney campaign placed 35.8 million phone calls, sent 3.4 million pieces of mail, and knocked on 2.8 million doors in the state of Ohio alone. The voting-eligible population of Ohio was approximately 8.6 million, so this means that the Romney campaign placed 4.2 phone calls, sent 0.4 pieces of mail, and knocked on 0.3 doors per person. These figures are similar for other heavily targeted states. Nevada—the state with the greatest level of effort—saw 5 phone calls, 0.7 pieces of mail, and 0.5 door knocks from the Romney campaign for every eligible voter. Qualitative reports suggest that the Obama campaign exerted similar if not greater mobilization efforts and focused more on door knocks and less on phone calls.[14] Combining the average numbers from the Obama and Romney campaigns, we approximate that the typical eligible voter in a heavily

---

[14] Personal communication with high-level Obama campaign staff, May to July 2013.

targeted state received 7-10 phone calls, 1-2 pieces of mail, and 1-2 door knocks from the presidential campaigns. Of course, some individuals received little contact and other received much more contact, but these numbers represent our best approximation of the average. Therefore, we conclude that a treatment including approximately 7-10 phone calls, 1-2 pieces of mail, and 1-2 door knocks increases turnout by 7-8 percentage points.

The effect of multiple voter contacts was significant in 2012. In fact, our estimated effects are consistent with a hypothesis that treatment effects are additive. Using the average effects of interventions from Green, McGrath, and Aronow (2013),[15] our point estimates are close to the sum of the average intent-to-treat effects of 10 phone calls, 2 pieces of mail, and 2 door knocks (10*0.4 + 2*0.7 + 2*1.0 = 7.4).

Figure 2 offers an approximate visualization of the effects of campaigning, based on the regression result in Column 7 of Table 4. Each data point represents a pair of states that share a media market. For example, the intersection of Massachusetts and New Hampshire in the Boston media market represents one data point in the figure. The difference in voter turnout between the two regions is plotted on the vertical axis, with higher values indicating higher turnout among the region that received more campaign contact. The difference in campaign effort between these two regions, as indicated by the *RomneyEffort* variable, is plotted on the horizontal axis. Consistent with the regression result, we see that, on average, the turnout differential increases significantly as the campaign effort differential increases. The solid line represents a linear fit, and the dashed curve represents a nonparametric kernel regression. In both cases, each observation is weighted by the population of the smaller region (represented visually by the size of the circle). This weighting

---

[15] We believe these figures, obtained from over 200 field experiments, represent the best available estimates of these average effects.

approach appropriately mitigates the influence of the data points with very few observations.[16] The slope of the line, .071, is nearly identical to the regression results in Columns 7-9 of Table 4. Moreover, the relationship between the effort differentials and the turnout differentials are approximately linear, further supporting our claim that the effects of GOTV are largely additive. If there were diminishing or increasing marginal returns to additional campaigning, then we would see the turnout differentials level-off or accelerate at a certain point. However, the effects of campaigning appear to increase linearly in proportion to its intensity, helping to explain the large cumulative effects that we detect.

[Figure 2]

Large-scale campaigns using modern GOTV tactics appear to significantly increase electoral participation. Our results suggest that voter turnout would have been 7-8 percentage points lower in heavily targeted states like Nevada, New Hampshire, and Ohio had the presidential campaigns and interest groups not deployed their ground campaigns. Even though a single campaign contact typically has a negligible effect on overall rates of participation, the widespread deployment of many such interventions appears to add up to a significant change in participation.

The regression results in Table 4 allow us to approximate the total mobilizing effect of the ground campaign in 2012. If we multiply the coefficient in Column 7 by the level of *RomneyEffort* in a state and the voting-eligible population in a state, we obtain an approximation of the number of people mobilized. Repeating this for every state and computing the sum, we estimate that approximately 2.6 million individuals turned out to vote in 2012 who would have otherwise abstained in the absence of the presidential ground campaigns. Because the campaigns attempted more than 400 million

---

[16] For example, the extreme observation with a turnout differential of −20 percentage points represents the intersection of Michigan and Minnesota in the Duluth-Superior media market, where the Michigan region has only 17,000 residents. Some regions are tiny—e.g., Nebraska's overlap with the Rapid City market has only 780 residents, while others are huge—e.g., the Philadelphia market covers 4.7 million people in Pennsylvania and 2.1 million people in New Jersey.

contacts, each attempt increased an individual's chances of turning out by about 0.5 percentage points, on average—a number that is strikingly consistent with previous experimental evidence.

Variation across Partisanship

In this section, we test whether the aggregate results from the previous section vary predictably across partisanship. By isolating the effects of a campaign on the most partisan voters, who are also were the voters most targeted by the campaigns, this analysis allows us to observe how strong campaign mobilization can be when used most heavily. This analysis also allows us to compare the relative effectiveness of the Obama and Romney campaigns. The Obama campaign of 2012 has been championed as the most technologically-sophisticated, evidence-based campaign in history while the Romney campaign was more traditional (e.g., Issenberg 2013). When we began this project, we surveyed 46 academics, and they predicted that Obama's campaign was almost 3 times as effective as Romney's in mobilizing supporters.[17] Do these perceptions manifest themselves in the data? Did the technological sophistication of the Obama campaign lead their GOTV efforts to be significantly more effective than Romney's? In this analysis, we also offer additional tests on the validity of our empirical design. If our estimated effects are greatest among the subgroups that were most heavily targeted, then we can be more confident that our results are attributable to campaigning rather than other spurious factors that might differ between targeted and non-targeted states.

Campaigns do not target all citizens equally within a targeted state. Logic suggests that partisan campaigns will focus their efforts in such a way to give them the largest return on their investment. Because strong partisans are most likely to support their candidate if they turn out, campaigns heavily target these individuals with GOTV messages. Our interviews with operatives in

---

[17] The median respondent predicted a 10.5 percentage point effect for Obama but just a 3.8 percentage point effect for Romney.

both campaigns revealed that targeting partisans was, in fact, their strategy. The most targeted individuals were strong partisans or those who look that way to the campaign based on their public records and demographics. For example, the Obama campaign generated an "Obama Support Score" using demographic and consumer data (similar to our data) in addition to their own contacts and focused their GOTV efforts on individuals with high scores (Nickerson and Rogers 2014).

We do not know precisely which citizens were targeted by the presidential campaigns. Our interviews with campaign operatives and other accounts reveal that targeting decisions were made based on support and turnout scores based on a small number of variables and internal polling. Our interviews with Obama campaign operatives revealed that decisions on targeting were made by directors at the state level, sometimes with discretion by operatives at local levels, so there was variation across states. We do not know the specific scores that were assigned to specific voters in our data. However, even if campaigns were willing to make this data available, it would not be helpful to us. Even if we knew which individuals in battleground states were contacted, we wouldn't know which individuals in non-battleground states would have been contacted had they lived in a battleground states. Moreover, it would not be possible to impute the type of voter that would be contacted in non-battleground states since the campaigns used data from their own communications with voters to determine who would be contacted, and these communications did not occur in non-battleground states.

Instead, we generate our own measure of partisanship, using our data from public and consumer records. Because vote choices are not public, the only direct indication of partisan attitudes is party registration, which is only available in some states. Moreover, we do not want to use party registration directly in our analysis, because campaigns work to register supporters in targeted states, making party registrants no longer comparable across state boundaries. And, of course, party registration is not present in all states. Therefore, we predict party registration using

race, gender, and income, using data from the non-battleground states where party registration is present. We code a party registration variable which takes a value of 1 for registered Democrats, 0 for registered Republicans, and .5 for all other individuals. Then, we calculate the average level of this variable for every possible combination of race, gender, and income in non-battleground states with party registration. This score represents our predicted levels of partisan support (with higher values indicating more Democratic and less Republican support), which we impute for the entire sample. The variable ranges from .09 to .84, with a mean of .51 and a standard deviation of .13.

In order to assess the variation in our estimated effects across partisanship, we divide our predicted partisanship variable into 6 categories and re-estimate the regression from Column 7 of Table 4 for each category. These estimates represent the effect of the ground campaign on turnout for voters at different levels of partisanship. Figure 3 presents the point estimates and 95% confidence intervals graphically. As expected, we see large estimates of mobilization for those voters predicted to be strongly partisan, and we obtain significantly smaller estimates for less partisan subgroups. Our estimated effects of campaign mobilization are 11.3 and 10.8 percentage points for our most Republican and Democratic subgroups, respectively the left- and right-most categories in Figure 3, while our estimates of mobilization are only 5.5 and 6.8 percentage points for the two least-partisan categories in the middle of Figure 3. These results provide additional validation of our empirical assumptions, because we see the largest effects of mobilization among the most heavily targeted subgroups.

[Figure 3]

The estimates of aggregate voter mobilization from the previous section are averages across all eligible voters, and many individuals in targeted states likely received little direct contact from campaigns. In Figure 3, when we focus on smaller subsets of individuals that were more likely to have been contacted, we obtain larger estimates. Among predicted partisans, the large-scale

21

interventions of the presidential campaigns appear to have increased turnout by more than 10 percentage points. Moreover, recall that our predictions of partisanship are imperfect. Presumably, if we had better information about the types of individuals targeted, we would estimate even greater mobilization effects among the most targeted groups, and we would expect starker differences in mobilization effects across subgroups.

A notable implication of these results is that campaigns can significantly change the composition of the voting population. By targeting partisan supporters, campaigns appear to produce a more partisan, polarized set of voters. Not only does this potentially distort the incentives of Presidential candidates to appeal to more partisan voters, but this also changes the set of individuals participating in down-ballot races, potentially affecting Congressional and legislative polarization. Consistent with previous research (Enos, Fowler, and Vavreck 2014), these results suggest that voter mobilization does not necessarily lead to a more representative set of voters. Furthermore, as mobilization techniques are more widely adopted,[18] we might expect the polarization of the voting population to increase further.

As discussed above, this analysis allows us to roughly compare the effectiveness of the Obama and Romney campaigns in mobilizing their respective supporters. Despite the purported technological sophistication of the Obama campaign and its devotion to a data-driven, evidence-based campaign, we see similar mobilization effects on both sides of Figure 3. The two campaigns were roughly comparable in their ability to turn out supporters. Moreover, this result is not an artifact of using the *RomneyEffort* variable; we obtain similar results with the other measures of campaign effort. Of course, this analysis does not preclude the possibility that the Obama campaign exceeded Romney in other areas that are more difficult to measure like persuasion or targeting.

---

[18] See, for example, Steve Friess, "Ahead of the Midterms, GOP Operatives Are Obsessively Studying a Book About the Obama Campaign", *New York Magazine*, September 12, 2013.

However, on the dimension of mobilizing people who look demographically like supporters, both campaigns appear to have been effective at comparable levels.[19]

One interpretation of the Romney campaign's slight advantage with their own partisans is that Democrats are simply harder to mobilize than Republicans. Indeed, previous research suggests that, on average, GOTV interventions are more effective for conservative and high-socioeconomic-status citizens (Enos, Fowler, Vavreck 2014). The amount of effort and resources needed to mobilize Democratic supporters may be greater than that needed to mobilize Republican supporters. With this in mind, even if the Obama campaign was more advanced than the Romney campaign, this difference was not great enough to overcome this structural disadvantage

A Victory Lab Revolution?

With the rebirth of GOTV field experiments in the 21st century and the adoption of these techniques by campaigns (e.g., Green and Gerber 2008), the ground game has become significantly more prevalent in recent elections. Television advertising continues to be the bread and butter of presidential campaigns, but campaigns have significantly shifted resources toward more personal contacts like direct mail, phone calls, and door-to-door canvassing (e.g., Issenberg 2012). In particular, the 2008 presidential campaign saw an unprecedented increase in the use of ground campaigning, and this was sustained in 2012. These tactics are gradually expanding to all levels of campaigns[20]. Given the substantively large effects that we estimate in this paper, this "Victory Lab Revolution", which we name after the best-selling book (Issenberg 2012), may have significantly

---

[19] We take this as one data point suggesting that the accounts of the Obama campaign's effectiveness and the big data revolution in campaigning may be overstated. One potential explanation for this overstatement is that political reporters are guilty of a fundamental attribution error, allowing the result of the election—influenced by many external factors—to shape their assessments of the campaigns.

[20] See, for example, Sasha Issenberg, "Dept. of Experiments," POLITICO, 27 February 2014.

altered the size and composition of the voting population. Perhaps the effectiveness of large-scale GOTV campaigns, combined with their increasing use, has marked a significant change in American campaigns.

Figure 4 provides evidence along these lines. For every presidential election between 1980 and 2012, we determine which 5 states were most pivotal in the following way. We sort the states according to the their two-party vote share, and determine which state would have tipped the electoral college had one candidate won every state below it and had the other candidate won every state above it. We then selected the two additional states on each side of this pivotal state to obtain our 5 pivotal states where much of the campaign effort was likely concentrated.[21] The top panel of Figure 4 plots the average difference in voter turnout between these pivotal states and all other states for each election, along with the corresponding 95% confidence intervals. Consistent with previous findings (e.g., Gerber et al. 2009), the differences in turnout between battleground and non-battleground states were smaller before 2008—hovering between 1.7 and 5.8 percentage points. Then, coinciding with a significant increase in GOTV efforts, this difference shot up to 9.1 percentage points in 2008 and remained at 6.9 percentage points in 2012.

[Figure 4]

The bottom panel of Figure 4 presents a similar analysis that draws inferences within states instead of between states. We run a differences-in-differences regression, taking advantage of states that switch in and out of the pivotal category, and estimate the extent to which within-state changes in battleground status correspond to within-state changes in turnout for each year.[22] These within state estimates allow us to approximate the effect of being a battleground state, and by proxy, the

---

[21] Alternate codings, such as identifying the states categorized as "swing states" by news organizations, produce similar results.

[22] Specifically, we regress turnout on year dummy variables, interactions between the year dummies and the indicator for a battleground state, and state fixed effects: $\text{Turnout}_{it} = \beta_1 * \text{Pivotal} * 1980 + \beta_2 * \text{Pivotal} * 1984 + \ldots + \delta_t \text{ (year fixed effects)} + \gamma_i \text{ (state fixed effects)} + \varepsilon_{it}$. Figure 4 shows the interactive coefficients.

approximate effect of campaigning on voter turnout in each election. Here, these effects are indistinguishable from zero between 1980 and 2000 but increase monotonically in each of the next three elections, with an estimated effect of 6.1 percentage points in 2012. The results are slightly different between the two approaches (between-state vs. within-state comparisons), but they both suggest that the growth of GOTV in the 21st century has significantly altered the level and distribution of participation in battleground states.

Separating the Effects of Internet Advertising and the Ground Campaign

Our data also allow us to partially separate the effects of two aspects of modern campaigning: internet advertising and the ground campaign. Internet advertisements can be finely targeted, so voters in battleground states were more likely to see them than those in non-battleground states and a portion of our estimated effects may be attributable to this new form of advertising. However, our own analysis indicates that the effect of internet advertising is minimal and virtually all of our previously estimated effects can be attributed to ground campaigning.

We expect the relative effectiveness of digital campaigning vs. traditional ground campaigning to be small for several reasons. First, the sheer scale of internet advertising was relatively small compared to other campaign efforts. We communicated with the digital teams of both campaigns and each spent less than 10 million dollars on registration and GOTV.[23] This amount pales in comparison to more traditional forms of GOTV, where the efforts of each campaign exceeded 100 million dollars. Second, previous evidence suggests that television advertisements do not significantly increase turnout (Krasno and Green 2008) and our own analysis in the Appendix confirms this result for 2012. Since internet advertising is similar to television

---

[23] Personal communications with Romney and Obama staff, January-February 2014.

advertising but smaller in scale, we might expect similarly small effects.[24] Third, the few experiments related to internet campaigning show small effect sizes. E-mails have zero (Nickerson 2007a, 2007b) or negligible (Malhotra, Michelson, and Valenzuela 2012) effects on registration and turnout. An extremely salient treatment utilizing social pressure conducted through Facebook on Election Day in 2010 increased turnout by only 0.3 percentage points (Bond et al. 2012). The only experiment, to our knowledge, directly assessing political internet advertisements, finds no effect on candidate evaluations (Broockman and Green 2013).

We conduct our own test of the relative importance of digital vs. ground campaigning by utilizing data from the Obama campaign on their digital campaigning activities in each state. Specifically, we acquired information on dollars spent on digital campaigning and advertisements run in each state.[25] As with ground campaign effort, we divided the dollars spent by the voting-eligible population in each state and rescale this variable to range from 0 to 1. As expected, digital advertising varied significantly across states with virtually no such activity in 41 states or DC. In Iowa, the Obama campaign spent almost one dollar per eligible voter, and they also made significant digital efforts in Colorado, Nevada, Ohio, and New Hampshire. Also as expected, digital and GOTV campaign effort are highly correlated across states, but enough variation exists that we can include both in a regression to estimate the relative importance of each strategy. Specifically, we repeat the regressions from Columns 7-9 of Table 4 but also include the continuous variable indicating digital effort which also ranges from 0 to 1. In other words, we regress voter turnout on continuous measures of ground effort and digital effort along with media market fixed effects and several control variables. Table 5 presents the results.

[Table 5]

---

[24] Both campaigns relied on video advertisements for their internet advertising, which is a medium very similar to that of television.

[25] Multiple files obtained from Obama staff, March 18, 2014. Data aggregated by the authors.

Consistent with our previous claim that our results are largely explained by traditional ground campaigning, we see that the coefficients for ground effort are similar to those in Table 4, while the coefficients for digital campaign effort are actually negative.[26] Of course, the strong correlations between ground and digital effort mean that these results should not be interpreted too strongly. Nonetheless, to the extent that we detect effects of campaigning on turnout, it appears to be attributable to the traditional forms of GOTV campaigning that have consistently been shown to boost turnout in experimental studies.

Discussion of Challenges

Due to the nature of our research question, we must make stronger assumptions than would be necessary with a randomized experiment, yet the discontinuous nature of campaigning allows us to make weaker assumptions than those typically necessary for observational studies. Given the low chances that a researcher could conduct an experiment on such a large scale or that a presidential campaign would randomize their ground campaign across a large-scale or long time-period, we believe this strategy presents the best opportunity to estimate these effects. In this section, we discuss and address potential challenges to our empirical design.

Possibly the most salient challenge to our empirical design has to do with the differential incentives and psychology of individuals in battleground and non-battleground states. We assume that, on average, the differences in turnout between battleground and non-battleground voters in the same media market are attributable to campaigning, and we designed placebo tests to validate this assumption. Here we return to this issue and draw on literature and additional tests to address this

---

[26] One explanation for this negative coefficient is that the Obama campaign attempted to strategically target these advertisements in areas where the turnout of supporters, based on early and absentee voting, was below their targets (personal communication with an Obama staffer on March 19, 2014). Digital advertising is more nimble than traditional GOTV and can be strategically targeted at local geographies in this manner. For these reasons, the negative coefficient should not be interpreted as evidence of a negative effect of digital campaigning.

challenge. Voters in battleground states may have greater incentives to vote independent of campaigning, for the same reasons that campaigns target these states. An individual's probability of casting a pivotal vote is higher in swing states, which could increase the direct returns to voting and make one more likely to turn out (Riker and Ordeshook 1968). On the other hand, the probability of casting a pivotal vote is miniscule in both battleground and non-battleground states (Gelman, King, and Boscardin 1998). As Schwartz (1987) aptly points out, "Saying that closeness increases the probability of being pivotal . . . is like saying that tall men are more likely than short men to bump their heads on the moon."

Despite the tiny probability of a pivotal vote in a presidential election, there are other reasons to expect higher turnout in battleground states even in the absence of campaigns. A combination of uncertainty and altruism could provide a rational basis for turnout and could explain higher participation in battleground states (Myatt 2012). Perhaps individuals in battleground states overestimate their probability of casting a pivotal vote, are more interested in the election for psychological reasons, or exert more social pressure on others for reasons independent of the campaign. Remember that all of our inferences are drawn within media markets, so differences in television advertising and news coverage are unlikely to explain our results, but these other psychological or social explanations could be relevant. We must acknowledge these possibilities, but we believe them to be minimal if not negligible for several reasons. First, in the absence of campaigning, we find no detectable differences between battleground and non-battleground voters: our placebo tests in Tables 2 and 3 show that individuals in targeted states were no more likely to be interested in politics or express an intention to vote before the presidential campaign began. Next, if other factors, as opposed to campaign mobilization explain our results, then we would not expect to find larger estimates among subgroups that were more likely to be targeted, as we see in Figure 3. Similarly, we would not expect to see the differences in turnout between battleground and non-

battleground states rise in accordance with the rise of ground campaign efforts in Figure 4, as these other factors would have presumably been present before the Victory Lab Revolution. Furthermore, experimental evidence from Enos and Fowler (2014) and Hoffman, Morgan, and Raymond (2013) suggests that considerations of pivotality have minimal effects on turnout. Lastly, using data collected by Enos and Hersh (2013) we find that individuals in battleground states are no more likely than those in non-battleground states to predict a close result of the 2012 Presidential Election in their state.[27] Collectively, these independent sources of evidence suggest that the extent to which our estimates are biased by the independent effects of battleground states is minimal.

By making comparisons within media markets, we can be sure that residents of targeted and non-targeted states receive similar levels of television advertising and news coverage. Nonetheless, we might worry that the effects of television advertisements and news coverage are somehow greater in battleground states. We acknowledge this possibility, but we believe it holds negligible implications for our estimates for several reasons. First, our best evidence suggests that television advertising has little impact on turnout (see Green and Krasno 2008 and our similar analysis in the Appendix).[28] Second, only a small fraction of Presidential advertisements—about 5 percent—contained state specific content.[29] As such, there is little reason to believe that advertising, based on content, would systematically be more salient to viewers in battleground states. Third, the turnout differential between competitive and uncompetitive states was small before 2008 when television advertising was just as prominent if not more so in presidential campaigns. If differential effects of

---

[27] When asked to predict Obama's vote share in their state in 2012, CCAP and CCES survey respondents in battleground states believed that their states were less than 0.1 percentage points more competitive than those in non-battleground states—a negligible and statistically insignificant difference.

[28] Recent evidence suggests that the effects of advertisements decay rapidly, perhaps within a few days (Gerber et al. 2011; Hill et al. 2013), implying that even if advertisements did affect turnout, the bulk of the effect may be completely gone by the day of the election.

[29] We collected data on all of Obama and Romney's televised ads from Stanford's Political Communication Lab, and we searched for state-specific references. Only 19 of 360 ads contained references to a specific state.

television in battleground states explain our results, one would not expect these effects to be greater in 2008 and 2012.

Conclusion

Campaigns have long been theorized to change the size and shape of the voting population, and the advent of the modern, scientifically-driven campaign further raises these prospects. However, rigorous evidence has been lacking. Experiments show that individual voter contacts increase turnout, but the substantive size of these effects is small. To our knowledge, the 2012 presidential campaign offers the best available opportunity to assess the effects of modern, large-scale campaigns, because these methods were deployed at a larger scale than ever before and because these efforts varied idiosyncratically across states.[30] According to our estimates, the ground efforts of the 2012 presidential campaign increased average levels of turnout by approximately 7 percentage points in the most heavily targeted states, mobilizing 2.6 million individuals who would have otherwise not turned out. In short, large-scale campaigns can significantly increase political participation.

Our calculations, along with our graphical analysis in Figure 2, suggest that the effects of many voter contacts are approximately additive. The effect of multiple door knocks, mailings, and phone calls is comparable to the sum of the average intent-to-treat estimates arising from single-shot experiments. This result may be surprising to those who expect diminishing returns to multiple treatments, but there are compelling reasons to expect additive or even multiplicative effects, at least up to a point. Only a small fraction of citizens will answer the door, pick up the phone, or read their

---

[30] Notably, our analysis of the 2012 campaign may be one of the final opportunities to parse the effects of various campaign tactics—television, internet, and ground—by taking advantage of discontinuities created by the Electoral College. This is because changes in television technology are quickly allowing the micro-targeting of advertisements to particular households, which will render ineffective the strategy of comparing voters across states but in the same media market.

mail, and each new treatment may reach a new subset of citizens. Moreover, some citizens that are unresponsive to single treatments may be eventually mobilized by multiple treatments.

In this paper, we have offered the first systematic assessment of the cumulative mobilization effect of a campaign. Contrary to some expectations, large-scale campaigns can significantly increase the size and composition of the voting population, rather than simply mobilizing a small fraction of voters on the margin. These findings may also lend insights for increasing participation in general, as the benefits of multiple and different treatments may be greater than previously expected. This phenomenon, in conjunction with recent increases in the use of ground campaigning, marks a significant change in the American electoral landscape, with millions of otherwise nonparticipating voters going to the polls. Indeed, as politicians in other countries adopt these new campaign tactics,[31] there is a potential for a world-wide increase in participation. These findings may also lend insights for increasing participation in general, as the benefits of multiple and different mobilization efforts may be greater than previously expected.

---

[31] See, for example, Annie Gowen and Rama Lakshmi, "Indian parties are using Obama-style campaign tactics in crucial election", *Washington Post*, April 7, 2014.

References

Ansolabehere, Stephen and Eitan Hersh. 2012. Validation: What Big Data Reveal about Survey Misreporting and the Real Electorate. *Political Analysis* 20(4):437-459.

Ansolabehere, Stephen and Eitan Hersh. 2013. Gender, Race, Age and Voting: A Research Note. *Politics and Governance* 1(2):132-137.

Bond, Robert M., Christopher J. Fariss, Jason J. Jones, Adam D. I. Kramer, Cameron Marlow, Jaime E. Settle, and James H. Fowler. 2012. A 61-Million-Person Experiment in Social Influence and Political Mobilization. *Nature* 489:295-298.

Broockman, David E. and Donald P. Green. 2013. Do Online Advertisements Increase Political Candidates' Name Recognition or Favorability? Evidence from Randomized Field Experiments. *Political Behavior*, forthcoming.

Enos, Ryan D. and Anthony Fowler. 2014. Pivotality and Turnout: Evidence from a Field Experiment in the Aftermath of a Tied Election. *Political Science Research and Methods*, forthcoming.

Enos, Ryan D., Anthony Fowler, and Lynn Vavreck. 2014. Increasing Inequality: The Effect of GOTV Mobilization on the Composition of the Electorate. *Journal of Politics* 76(1):273-288.

Enos, Ryan D. and Eitan D. Hersh. 2013. Elite Perceptions of Electoral Closeness: Fear in the Face of Uncertainty or Overconfidence of True Believers. Midwest Political Science Association Annual Meeting, Chicago.

Cardy, Emily Arthur. 2005. An Experimental Field Study of the GOTV and Persuasion Effects of Partisan Direct Mail and Phone Calls. *The Annals of the American Academy of Political and Social Science* 601:28-40.

Gelman, Andrew, Gary King, and W. John Boscardin. 1998. Estimating the Probability of Events That Have Never Occurred: When Is Your Vote Decisive? *Journal of the American Statistical Association* 93(441):1-9.

Gerber, Alan S., James G. Gimpel, Donald P. Green, and Daron R. Shaw. 2011. How Large and Long-lasting Are the Persuasive Effects of Televised Campaign Ads? Results from a Randomized Field Experiment. *American Political Science Review* 105(1):135-150.

Gerber, Alan S., Gregory A. Huber, Conor M. Dowling, David Doherty, and Nicole Schwartzberg. 2009. Using Battleground States as a Natural Experiment to Test Theories of Voting. APSA 2009 Toronto Meeting Paper.

Green, Donald P. and Alan S. Gerber. 2008. *Get Out the Vote: How to Increase Voter Turnout*. Brookings Press.

Green, Donald P., Mary C. McGrath, and Peter M. Aronow. 2013. Field Experiments and the Study of Voter Turnout. *Journal of Elections, Public Opinion, and Parties* 23(1):27-48.

Hersh, Eitan. 2011. At the Mercy of Data: Campaign' Reliance on Available Information in Mobilizing Supporters. Unpublished manuscript.

Hersh, Eitan. 2013. Long-term Effect of September 11 on the Political Behavior of Victims' Families and Neighbors. *Proceedings of the National Academy of Sciences*, forthcoming.

Hill, David and Seth McKee. 2005. The Electoral College, Mobilization, and Turnout in the 2000 Presidential Election. *American Politics Research 33*(5):700-725.

Hill, Seth J., James Lo, Lynn Vavreck, and John Zaller. 2013. How Quickly We Forget: The Duration of Persuasion Effects From Mass Communication. *Political Communication 30* (4):521-547.

Hoffman, Mitchell, John Morgan, and Collin Raymond. 2013. One in a Million: A Field Experiment on Belief Formation and Pivotal Voting. Working paper.

Huber, Gregory A. and Kevin Arceneaux. 2007. Identifying the Persuasive Effects of Presidential Advertising. *American Journal of Political Science* 51(4):957-977.

Issenberg, Sasha. 2012. *The Victory Lab: The Secret Science of Winning Campaigns*. Crown Publishers.

Issenberg, Sasha. 2013. A More Perfect Union: How President Obama's Campaign Used Big Data to Rally Individual Voters. *MIT Technology Review* 116(1):41-51.

Kim, Jae-On, John Petrocik, and Stephen Enokson. 1975. Voter Turnout among the American States: Systematic and Individual Components. *American Political Science Review 69*(1):107-123.

Krasno, Jonathan S. and Donald P. Green. 2008. Do Televised Presidential Ads Increase Voter Turnout? Evidence from a Natural Experiment. *Journal of Politics* 70(1):245-261.

Lipsitz, Keena. 2008. The Consequences of Battleground and "Spectator" State Residency for Political Participation. *Political Behavior 31*(2):1-23.

Malhotra, Neil, Melissa R. Michelson, and Ali Adam Valenzuela. 2012. Emails from Official Sources Can Increase Turnout. *Quarterly Journal of Political Science* 7(3):321-332.

Michelson, Melissa R., Lisa Garcia Bedolla, and Donald P. Green. 2007. New Experiments in Minority Voter Mobilization. Report for The James Irvine Foundation.

Myatt, David P. 2012. A Rational Choice Theory of Voter Turnout. Working paper.

Nickerson, David W. 2007a. Does Email Boost Turnout? *Quarterly Journal of Political Science* 2(4):369-379.

Nickerson, David W. 2007b. The Ineffectiveness of E-Vites to Democracy: Field Experiments Testing the Role of E-mail on Voter Turnout. *Social Science Computer Review* 25(4):494-503.

Nickerson, David W. and Todd Rogers. 2014. Political Campaigns and Big Data. *Journal of Economic Perspectives* 28(2):51-74.

Ramirez, Ricardo. 2005. Giving Voice to Latino Voters: A Field Experiment on the Effectiveness of a National Nonpartisan Mobilization Effort. *The Annals of the American Academy of Political and Social Science* 601:66-84.

Riker, William H. and Peter C. Ordeshook. 1968. A Theory of the Calculus of Voting. *American Political Science Review* 62(1):25-42.

Schwartz, Thomas. 1987. Your Vote Counts on Account of the Way It Is Counted: An Institutional Solution to the Paradox of Voting. *Public Choice* 54(2):101-121.

Sides, John and Lynn Vavreck. 2013. *The Gamble: Choice and Chance in the 2012 Presidential Election.* Princeton University Press.

Wolak, Jennifer. 2006. The Consequences of Presidential Battleground Strategies for Citizen Engagement. *Political Research Quarterly 59*(3):353-361.

Table 1. Placebo Tests using 2010 Voter Turnout

	DV = 2010 Voter Turnout					
	(1)	(2)	(3)	(4)	(5)	(6)
Battleground	.003	.001				
	(.012)	(.011)				
Tier 1			.014	.010		
			(.016)	(.015)		
Tier 2			.024	.027		
			(.013)	(.012)*		
Tier 3			.029	.018		
			(.017)	(.019)		
Romney Effort					.014	.013
					(.020)	(.019)
Media Market Fixed Effects	X	X	X	X	X	X
Individual Controls		X		X		X
N individuals	49,549,516		87,650,509		118,645,707	
N media markets	38		49		74	
N states	34		44		48	
R-squared	.013	.108	.012	.099	.014	.099

*Battleground is an indicator variable for battleground territory where the campaigns made significant efforts to mobilize voters. Tier 1, Tier 2, and Tier 3 are indicator variables for the three categories of states that received different levels of campaign priority from the Obama campaign. Romney Effort is a continuous variable ranging from 0 to 1 indicating the level of mobilization efforts employed by the Romney campaign in each state. Media Market Fixed effects allow for all of the inferences to be drawn within media markets but across state or district boundaries. Specifications with individual controls include dummies for gender, race, and predicted income category. The first two regressions include all individuals living in media markets that span both battleground and non-battleground territory. The next two regressions include all individuals living in media markets that span territories at different tiers. The final two regressions include all individuals living in media markets that span territories receiving differential mobilization effort from the Romney campaign. Because the dependent variable is voter turnout in 2010, a year where there was no presidential campaigning, these regressions constitute placebo tests, indicating whether there are underlying differences between targeted and non-targeted states. All of the point estimates are substantively small, and all but one are statistically indistinguishable from zero. The results are consistent with the underlying assumptions of our empirical strategies. Standard errors, clustered by state, are in parentheses; ** p < .01, * p < .05.*

Table 2. Placebo Tests Using Political Interest before the General Election Campaign

	DV = "very much interested" in politics and current events								
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Battleground	−.005	−.004	−.012						
	(.014)	(.013)	(.013)						
Tier 1				−.003	.000	−.006			
				(.018)	(.020)	(.019)			
Tier 2				−.005	−.011	−.019			
				(.017)	(.014)	(.012)			
Tier 3				.046	.031	.020			
				(.030)	(.024)	(.025)			
Romney Effort							−.008	−.008	−.021
							(.023)	(.022)	(.022)
Early Primary State			.025			.021			.025
			(.015)			(.014)			(.014)
Individual Controls		X	X		X	X		X	X
Survey Week Fixed Effects	X	X	X	X	X	X	X	X	X
Observations					15,998				
R-squared	.001	.106	.106	.002	.106	.106	.001	.106	.106

*All variables are defined as in Table 1. Data is from the 2012 CCAP. We analyze responses for the first 16 weeks of 2012, before Romney was the presumptive nominee and before general election campaigning got underway. We find no detectable differences in political interest between battleground and non-battleground voters before general election campaigning, providing empirical support for our assumptions that these individuals are comparable in terms of their underlying interests and intentions to vote. Standard errors, clustered by state, are in parentheses; ** p < .01, * p < .05.*

Table 3. Placebo Tests Using Vote Intention before the General Election Campaign

	DV = "will not vote" in 2012								
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Battleground	.007	.007	.011						
	(.008)	(.008)	(.009)						
Tier 1				.005	.002	.005			
				(.008)	(.008)	(.007)			
Tier 2				.019	.022	.026			
				(.010)	(.009)*	(.008)**			
Tier 3				−.015	−.018	−.012			
				(.009)	(.010)	(.011)			
Romney Effort							.010	.008	.014
							(.010)	(.011)	(.012)
Early Primary State			−.012			−.011			−.011
			(.008)			(.008)			(.008)
Individual Controls		X	X		X	X		X	X
Survey Week Fixed Effects	X	X	X	X	X	X	X	X	X
Observations					15,998				
R-squared	.002	.033	.033	.003	.034	.034	.002	.033	.033

*All variables are defined as in Table 1. Data is from the 2012 CCAP. We analyze responses for the first 16 weeks of 2012, before Romney was the presumptive nominee and before general election campaigning got underway. We find no detectable differences in vote intentions between battleground and non-battleground voters before general election campaigning, providing empirical support for our assumptions that these individuals are comparable in terms of their underlying interests and intentions to vote. For the two statistically significant coefficients, the point estimate is in the "wrong" direction (tier 2 voters are more likely to say that they will not vote), and the substantive size is small. These differences are most likely attributable to chance and multiple testing. Out of 40 placebo tests in Tables 1, 2, and 3, we obtain 3 statistically significant coefficients. Even if these test were independent (which they are not), we would expect to obtain 3 statistically significant coefficients or more by chance 32 percent of the time. Standard errors, clustered by state, are in parentheses; ** $p < .01$, * $p < .05$.*

Table 4. Average Effects of Ground Campaigning on 2012 Voter Turnout

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
				DV = 2012 Voter Turnout					
Battleground	.042	.041	.041						
	(.011)**	(.010)**	(.010)**						
Tier 1				.067	.063	.056			
				(.006)**	(.006)**	(.012)**			
Tier 2				.039	.043	.023			
				(.006)**	(.005)**	(.007)**			
Tier 3				.035	.025	.011			
				(.016)*	(.017)	(.010)			
Romney Effort							.078	.078	.068
							(.012)**	(.012)**	(.014)**
Media Market Fixed Effects	X	X	X	X	X	X	X	X	X
Individual Controls		X	X		X	X		X	X
Control for 2010 Turnout			X			X			X
N individuals		49,549,516			87,650,509			118,645,707	
N media markets		38			49			74	
N states		34			44			48	
R-squared	.020	.101	.624	.014	.102	.621	.019	.104	.624

*The variables and samples are defined as in Table 1. Conditional on media market and other factors, individuals living in battleground territory were 4 percentage points more likely to turn out to vote. Those living in the first tier of states were 6 to 7 percentage points more likely to turn out, relative to those in non-targeted states. In the second tier, this difference is 2 to 4 percentage points, and in the third tier this difference ranges from 1 to 4 percentage points. Those in the states where the Romney campaign exerted the greatest effort were 7 to 8 percentage points more likely to vote than those where Romney exerted no effort. All of these results are consistent with powerful mobilizing effects of campaign efforts. Standard errors, clustered by state, are in parentheses; ** p < .01, * p < .05.*

Table 5. Separating Effects of the Digital Campaign and Traditional GOTV

	DV = 2012 Voter Turnout		
	(1)	(2)	(3)
Ground Effort	.091	.091	.080
	(.014)**	(.014)**	(.016)**
Digital Effort	−.041	−.044	−.042
	(.014)**	(.012)**	(.015)**
Media Market Fixed Effects	X	X	X
Individual Controls		X	X
Control for 2010 Turnout			X
N individuals	118,645,707		
N media markets	74		
N states	48		
R-squared	.019	.104	.624

*The table mirrors Columns 7-9 of Table 4 with the exception of one new variable. "Ground Effort" is the Romney Effort variable, indicating the level of ground campaigning in each state. "Digital Effort" is a continuous measure, ranging from 0 to 1, indicating the dollars per person that the Obama campaign spent on internet ads in each state. Virtually all of the previous estimates appear to be attributable to ground campaigning as opposed to digital campaigning. Standard errors, clustered by state, are in parentheses; ** p < .01, * p < .05.*
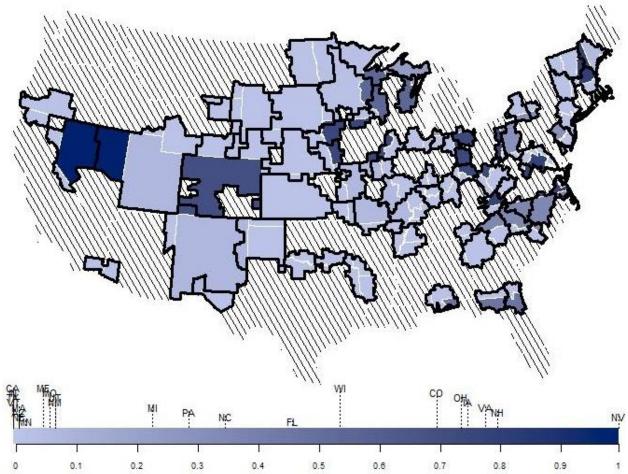
Figure 1: *RomneyEffort* by State and Media Market

Media markets spanning states with different levels of RomneyEffort are surrounded by thick black lines. Portions of each media market are shaded by the level of Romney effort in that state. So, for example, the Reno media market is much darker in Nevada than in California, representing the variation in RomneyEffort across those states. States with no effort or only minimal effort are the lightest shade. These are states, for example Utah, where the campaign made some contact but nowhere near the levels in battleground states. The media markets that do not span states with varying RomneyEffort (not included in our analysis) are indicated by diagonal lines. The states are also aligned horizontally along the bottom of the figure by level of RomneyEffort (vertical alignments are varied for visual ease).
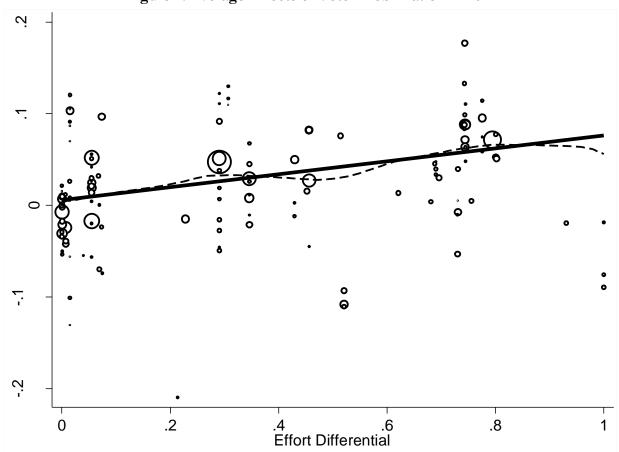
Figure 2. Average Effects of Voter Mobilization in 2012

Each circle represents the intersection of two states within a single media market (e.g., MA and NH residents living in the Boston media market). The size of each circle is proportional to the population (that of the smaller region). The vertical axis plots the difference in 2012 turnout between the two regions (with higher values representing higher turnout in the more heavily targeted state) and the horizontal axis plots the difference in campaign effort (as measured by the RomneyEffort variable) between the two states. The solid line represents a linear fit and the dashed curve represents a kernel regression. In both cases, each observation is weighted by population (that of the smaller region). Consistent with the regression results in Table 4, the turnout differential increases strongly as the effort differential increases. Moreover, the relationship is approximately linear suggesting approximately additive effects of multiple GOTV contacts.
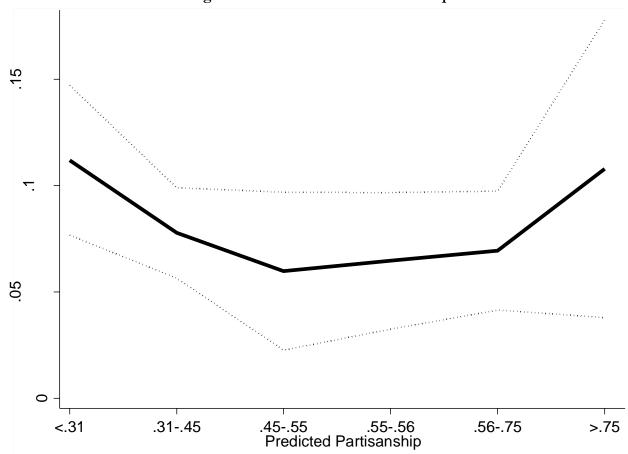
Figure 3. Variation across Partisanship

The figure plots the results for six different regressions, replicating the analysis from Column 7 in Table 4 for different subsets of individuals with different levels of predicted partisanship. Going from left to right, we analyze the 10 percent of individuals with the lowest levels of predicted partisanship (most Republican), then the 10-25th percentiles, 25-50th percentiles, 50-75th percentiles, 75-90th percentiles, and the top 10 percent (most Democratic). The solid line represents the coefficients and the dotted lines indicate 95% confidence intervals (generated from state-clustered standard errors). We see the greatest estimates (11 percentage points) for the most partisan subgroups, and much smaller effects for the least partisan groups which were less likely to be contacted by the campaigns.
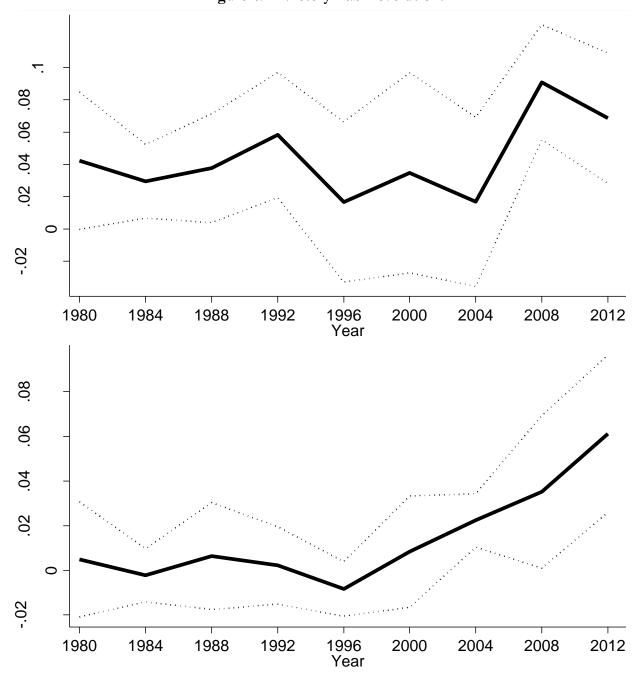
Figure 4. A Victory Lab Revolution?

The figure presents two sets of analysis consistent with the notion that GOTV campaigns significantly increase turnout and that GOTV has become more widespread in the 21st century. We define the 5 most pivotal states in every presidential election from 1980 to 2012 and compare the turnout rates between pivotal and non-pivotal states in each year. The top panel shows the average turnout differential between pivotal and non-pivotal states in each year. The bottom panel shows regression coefficients which explore within-state changes in turnout as a state switches from non-pivotal to pivotal in a particular year. In both cases, the turnout differential between the pivotal and non-pivotal categories has increased dramatically in recent elections, consistent with a "Victory Lab Revolution." Dotted lines indicate 95% confidence intervals (derived from robust/state-clustered standard errors).

Appendix

Effects of TV Advertising on Turnout in 2012

In order to assess the effects of television advertising, we largely replicate the analysis of Krasno and Green (2008) using our individual-level data from 2012. Because television advertising intended for battleground states spills over into the non-battleground states that share a media market with a battleground state, we can exploit this spillover to obtain exogenous variation in television advertising, and estimate the effect of these ads on turnout. We restrict our attention to voters in non-battleground states where some media markets overlap with a battleground state and others do not—there are 24 such states that fit these criteria. We compare the turnout behavior of voters in the media markets that share territory with a battleground state to that of voters in the non-battleground media markets in the same state. To do this, we regression 2012 turnout on an indicator for a "battleground media market" (a media market that shares territory with a battleground state) and state fixed effects. In column 2 we add controls for race, income, age, and gender, and in column 3 we include 2010 turnout as an additional control variable.

Table A1 presents the results of this analysis. Consistent with Krasno and Green (2008) we find no evidence that television advertising increases turnout. The point estimate for each specification is substantively tiny—less than 1 percentage point—and precisely estimated. Therefore, the air campaign appears to have little effect on participation. This result, in conjunction with our main analyses, suggests that the ground campaign and personal contacts are much more effective in mobilizing voters than television.

Table A1. Effects of TV Advertising in 2012

	DV = 2012 Voter Turnout		
	(1)	(2)	(3)
Battleground Media Market	.007	.002	.005
	(.006)	(.005)	(.003)
State Fixed Effects	X	X	X
Individual Controls		X	X
Control for 2010 Turnout			X
N individuals		141,886,753	
N media markets		139	
N states		24	
R-squared	.012	.301	.709

*Standard errors, clustered by media market, are in parentheses; ** p < .01, * p < .05.*