

Youth Voter Mobilization Through Online Advertising: Evidence From Two GOTV Field Experiments*

Kevin Collins

Analyst Institute

kcollins@analystinstitute.org

Joshua L. Kalla

(corresponding author)

Department of Political Science

University of California, Berkeley

kalla@berkeley.edu

Lauren Keane

Analyst Institute

lkeane@analystinstitute.org

Abstract

An increasingly fragmented media environment poses a challenge to campaigns and political organizations trying to persuade young voters, as young people increasingly eschew television for online video streaming and represent a growing cell-phone-only population, which is more costly to reach. As a result, each cycle campaigns are spending more on online advertising. However, the effectiveness of this outreach is an open question with disparate findings in the literature (Bond et al. 2012, Broockman and Green 2012). To test the effectiveness of this outreach for voter mobilization, we partnered with a national organization, Rock the Vote, to conduct a seven-state, 731,568-person GOTV experiment using advertising on Facebook during the 2012 presidential election, which involved nearly four million advertising impressions, and a 14-state, 93,053-person replication during the 2013 general election. Across both experiments, we find no evidence that voter turnout in the treatment group was greater than the control group. While exploratory analysis from the 2012 experiment suggested the ads were effective among those who had a demonstrated a prior history of clicking on online advertising, the follow-up 2013 experiment failed to confirm this finding. These results suggest that while online advertising can reach young voters, it may not be a panacea for encouraging them to vote. This research also underscores the importance of replications of even large-scale experimental findings.

* Acknowledgements: Thanks are owed to the team at Rock the Vote who made this experiment possible, including Heather Smith, Amanda Brown, and David Pruter. Thanks also to the work of Analyst Institute staff in conducting and analyzing this experiment, most especially Josh Rosmarin. We thank David Broockman for helpful feedback. This research was reviewed by the Yale University Human Subjects Committee, Protocol #1404013677. Authors are listed in alphabetical order.

Introduction

How do campaigns reach young voters and shape their political behavior in the digital media age? A decade and a half of experimental research on campaign practices have transformed political scientists' and practitioners' understanding of the effectiveness of campaigns, and along the way it has reshaped the way that campaigns operate. This research has shown that personalized voter contact is a powerful tool for mobilizing (Gerber and Green 2008) and persuading (Arceneaux 2007, Arceneaux and Kolodny 2009), but the large field operations necessary to implement these insights require extensive management of volunteers and paid staff, and as a result do not scale easily.

The limitations campaigns face when trying to reach people through personalized modes is particularly acute for young people. While they are equally susceptible to canvasses when they are reached, one experiment suggests that young voters are approximately three times as hard to reach at the door than older voters (Nickerson 2006). Similarly, a greater proportion of young voters do not have landlines, and cell phones are much more expensive to reach due to current laws about the use of predictive dialers (Link et al. 2007).

Despite the demonstrated effectiveness of personalized voter contact, television advertising remains the main tool for campaign communication in terms of campaign expenditures, due in no small part to the fact that it is a scalable method. Unsurprisingly, scholars have also sought to understand the mobilization effects of these ads. In a 2004 study of cable television ads produced by Rock the Vote, Donald Green and Lynn Vavreck (2008) found turnout 3.0 percentage points (S.E.=1.4) higher among 18- and 19-year old voters in treatment group cable zones than in control group cable zones.

However, voters are less reliably reachable through television ads than they once were. As Markus Prior (2007) has documented, the rise of cable advertising has eroded news audiences, a traditional focus for television advertising buys. Of course, it is possible to advertise political

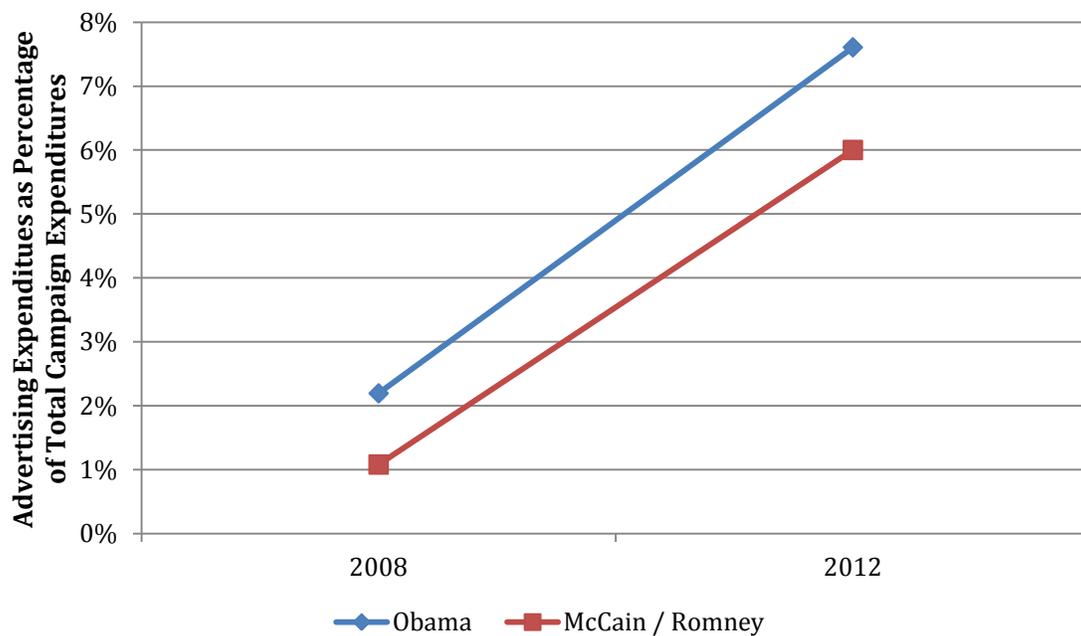
messages on cable (the means of delivery in the Green and Vavreck study, for example). But the decline of broadcast news audiences has also coincided with a proliferation of cable channels, making it harder to reach target audiences in any one place. Moreover, there exists a recent trend toward “cord cutting,” as audiences give up cable and satellite television for online content (Ramachandran 2012).

At the same time that television audiences are fragmenting, voters are increasingly engaging with political information online. This is particularly true for young voters; 67% of all adults aged 18-24 were politically active on social networking sites compared to just 13% of adults aged 65 and older (Smith 2013). But conditional on being a social network user, the differences between young and old adults diminishes. Among social network site users, 57% of those over 65 years old were politically active on them compared to 72% among those aged 18-24. Thus, much of this variation is due to younger adults being more likely to use social network sites rather than how different age cohorts behave once they are on a site. Survey evidence from 2008 also shows that young voters are more likely to use the Internet to engage in political behavior than are older voters (Schlozman, Verba, and Brady 2010). Social networking sites like Facebook are particularly appealing avenues for contacting younger voters because they are places where they can successfully communicate with these voters. This is especially appealing if the contact to voters through social networking sites is a more personalized form of voter contact, carrying with it social cues about the importance of participation (Bond et al. (2012).

And indeed, campaigns are increasingly devoting resources to online outreach, in an attempt to address media fragmentation. Campaigns also recognize that online organizing dramatically reduces the marginal cost of reaching additional targets (Garcia-Castanon, Rank, and Barretto 2011). Online voter outreach has taken a variety of forms, from social networking sites to video games (Yenigun 2012). The most important emerging area for reaching voters is online advertising.

Expenditure on political advertising has increased dramatically over the past few election cycles. In 2008, the Obama campaign spent \$16 million online (just over 2% of total campaign expenditures) and the McCain campaign spent \$3.6 million online (just over 1% of total campaign expenditures) (Kaye 2009). However, in 2012 the Obama campaign spent \$52 million in online advertising (almost 8% of total campaign expenditures) and the Romney campaign spent \$26 million (almost 6% of total campaign expenditures) (Stampler 2012).

Figure 1: Growth in Online Political Advertising



These ads serve a variety of purposes, from trying to persuade voters, soliciting campaign contributions, recruiting volunteers, registering voters, and turning supporters out to vote. However, despite the growing centrality of online advertising to campaign practices in the United States, relatively little is known about the effectiveness of this advertising. According to Barnard and Kreiss's survey of extant research on online political advertising, "We know little about ... what effects online advertising has on voters' political participation, exposure to political communication, and attitudes about candidates" (Barnard and Kreiss 2013 p.2047).

So how can we effectively evaluate the impact of online advertising? Barnard and Kreiss write “Online advertising is a “closed loop”; staffers can judge an ad’s effectiveness by viewing data on click-throughs. Tracking user behavior in real time enabled staffers to continuously measure these outcomes and calculate returns on investment (ROI) for all the campaign’s online advertising” (Barnard and Kreiss 2014 p.2053). However, this analysis hinges on a strong counterfactual assumption that people exposed to advertising would not have taken the action in question if they were not exposed to that advertising. If the action is a signature on a petition, this assumption may be realistic. If the action is voting or voting for a specific candidate, it clearly does not hold.

A far superior strategy for estimating the effectiveness of online advertising is to conduct field experiments, which is to say a randomized experiment that meets several criteria summarized by Gerber and Green (2012):

“[W]hether the treatment used in the study resembles the intervention of interest in the world, whether the participants resemble the actors who ordinarily encounter these interventions, whether the context within which the subjects receive the treatment resembles the context of interest, and whether the outcome measures resemble the actual outcomes of theoretical or practical interest” (Gerber and Green 2012 p.10-11).

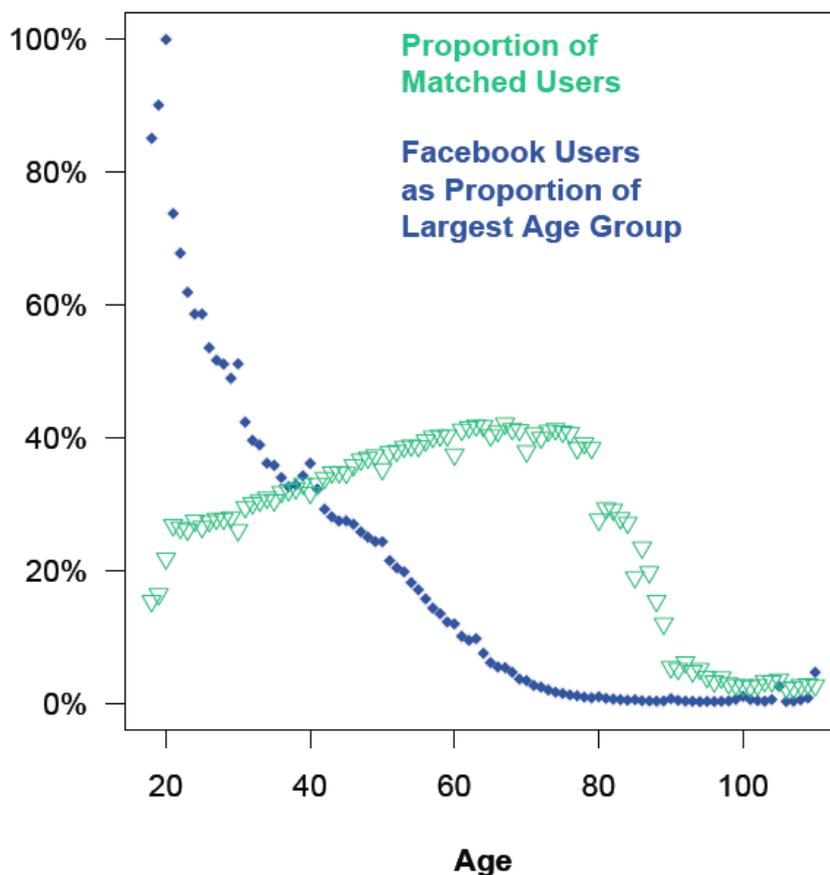
In the case of online political advertising, this means randomly assigning actual voters to receive actual ads while actually viewing content on websites in the midst of an actual election campaign; then measuring their actual voting behavior. To date, little such research has been conducted. Perhaps this gap in the literature is unsurprising; as Lewis and Rao (2013) point out, the low cost of online advertising necessitates very large samples in order to reliably estimate return-on-investment in field experiments of online advertising. However, given the growing importance of online advertising to political campaigns, the importance of closing this gap in our understanding is growing as well.

The limited academic research that evaluates the effectiveness of online political advertising through field experiments has shown online advertisements to have small average effects, though as of yet we have little guidance on the effect of online advertising campaigns on voter turnout. The most prominent study of turnout in relationship to postings on Facebook is Bond et al. (2012). This experiment, conducted in coordination with Facebook's data science team, posted one of two messages at the top of a Facebook user's "News Feed" on Election Day. Messages that contained information about how many of the user's friends had already reported voting had a turnout rate of 0.39 percent higher than a control group. With 6,338,882 observations matched to the voter file, even this small effect is statistically significant.

However, this experiment is of somewhat limited value when trying to learn about the effectiveness of online ads for voter mobilization for two reasons. First, the experimental treatment here is not an ad unit that Facebook actually makes available to advertisers. It appeared as a banner ad at the top of a user's News Feed all day, unless the target clicked to remove the ad or interacted with it in another fashion. While News Feed advertising is a product Facebook sells, these ads are not nearly as prominent as the treatment in this experiment.

Second, rather than starting with a list of Facebook users already matched to the voter file, these researchers randomized anonymous Facebook users to experimental conditions, and then after the election attempted to match the users to voter files in a subset of states with inexpensive voter files. Only one of three users in the experiment matched positively to the voter file. In addition to leaving open the possibility of differential attrition, this matching procedure left a sample with verified voter records that was much older than the young-skewing set of Facebook users, as is evident from the researchers' graph S2 from their supplemental materials, reprinted below as Figure 2. If the effects of online ads are greater among younger users, then this matching procedure could downwardly bias the estimated average treatment effect.

Figure 2: Distribution of Facebook Users and Matched Users by Age



Note: Reprinted Figure S2 from Bond et al. (2012)

Other research on Facebook advertising also shows modest effects at best. Broockman and Green (2013) conducted two experiments with Facebook advertising in conjunction with two campaigns, one for state legislature and one for Congress. They found non-significant treatment effects, though the relatively small size of the studies (2,701 Facebook users who responded to the phone surveys, pooled across both experiments) meant that these experiments are relatively underpowered to reliably detect as statistically significant effects of the observed magnitude.

Therefore, there is no prior research on the effect of a Facebook (or other online) advertising campaign that can reliably estimate the effectiveness of online advertising for mobilizing

voters, despite the growing importance of this mode of contact for political campaigns. This prior research leaves us with our first research question.

- *Research Question 1: What is the average treatment effect (ATE) of online GOTV advertising at mobilizing voters?*

However, not all voters are equally likely to see these ads, even if we intend to treat them with advertising. Due to privacy concerns, Facebook is not able to provide individual-level data on how many times an individual was served the advertisements. Instead, Facebook reports the percentage of individuals who were served at least one advertisements and the average number of advertisements served by person.

As the above discussion noted, young voters are particularly likely to engage with political content online. Furthermore, younger people have been found to spend more time on Facebook than older adults (McAndrew and Jeong 2012). By spending more time on Facebook, younger people in the treatment condition are likely to be exposed to more of a campaign's advertisements. Thus, we believe that age can serve as a partial proxy for the number of advertisement views. Thus, we would expect younger people to be more treatment responsive than older adults. This means we not only want to examine the overall ATE, but also the conditional average treatment effects (CATEs) among younger voters:

- *Research Question 2: What is the conditional average treatment effect (CATE) of online GOTV advertising among young voters (who are heavier users of Facebook)?*

After conducting the first experiment in 2012, we received additional data from the organization documenting what method individuals in the experimental universe had used to register to vote (e.g., at a campus booth, from an email, through an online advertisement). We conducted exploratory analysis to see if individuals who had registered through an online advertisement or specifically through an advertisement on Facebook, from the same organization sponsoring the

GOTV ads, were more likely to be treatment responsive. We hypothesize that previous experience with an online advertisement demonstrates that these types of individuals are more likely to pay attention to these advertisements and thus be more treatment responsive. This leads to our final research question:

- *Research Question 3: What is the CATE of online GOTV advertising among voters who have a demonstrated history of interacting with online ads from the sponsoring organization?*

While this research question grew out of exploratory data analysis in the 2012 experiment, we use a second experiment conducted during the 2013 general election as an out-of-sample replication of this exploratory finding.

Study 1: Online Advertising Experiment in 2012

Research Design

We use a large-scale field experiment to measure the effectiveness of online advertising on voting participation. The field experiment was conducted across seven battleground states prior to the November 6, 2012 general election using the social network site Facebook.

The experimental population consisted of 731,568 individuals who had previously registered to vote with Rock the Vote. These voters lived in Colorado, Florida, North Carolina, Ohio, Pennsylvania, Virginia, and Wisconsin. All of these individuals had previously provided an email address to Rock the Vote, mostly while in the process of registering to vote in person or online. Randomization was conducted at the individual level.

The experimental population was randomly assigned to one of two experimental conditions: 1) an uncontacted control group ($n = 365,668$), or 2) a group that was assigned to receive Rock the Vote's get-out-the-vote Facebook advertisements ($n = 365,900$). Random assignment occurred within strata, which were defined by the individual's voter registration state and terciles of age.

The GOTV advertisements were delivered on Facebook using their “Custom Audiences” tool. This feature allows a client to upload a list of email addresses, which is then matched to Facebook’s user database. Rock the Vote’s email list is particularly well suited to this type of program, as Rock the Vote obtains email addresses from users when they register, and in the case of the online registration tool in states without fully electronic registration, uses those email addresses to deliver the voter registration form to the voter. This should provide a higher match rate than, for example, commercially available email lists. Unfortunately, Facebook does not disaggregate the data into which email addresses were successfully matched to a Facebook account, so we were not able to remove non-Facebook users from the experimental sample. Instead, Facebook informed us that 54.2% of the treatment group email addresses successfully matched to a Facebook user and that among those who matched 76% were served with at least one Rock the Vote advertisement. This treatment group noncompliance does not bias our results. Because email addresses were randomly assigned to treatment or control, in expectation, the proportion of Facebook users and thus those individuals who are ineligible to receive Facebook advertisements, is balanced across experimental conditions. We limit our results to the intent-to-treat estimate because we do not know the compliance rate among the subgroups.

Beginning October 31st and continuing through Election Day, ads were served to treatment targets with email addresses matched to a Facebook user. These ads were placed in both the user’s sidebar and “News Feed.” News Feed advertisements were restricted to those Facebook users who had “liked” Rock the Vote on Facebook or whose friends had liked Rock the Vote, a policy that has since been relaxed. 150,598 individuals were served at least one advertisement (76% of those who matched to Facebook and 41% of the overall treatment group) and those who were served

advertisements viewed 25 advertisements on average. Sample advertisements are presented below in Figure 3:¹

Figure 3: Sample Facebook Advertisements

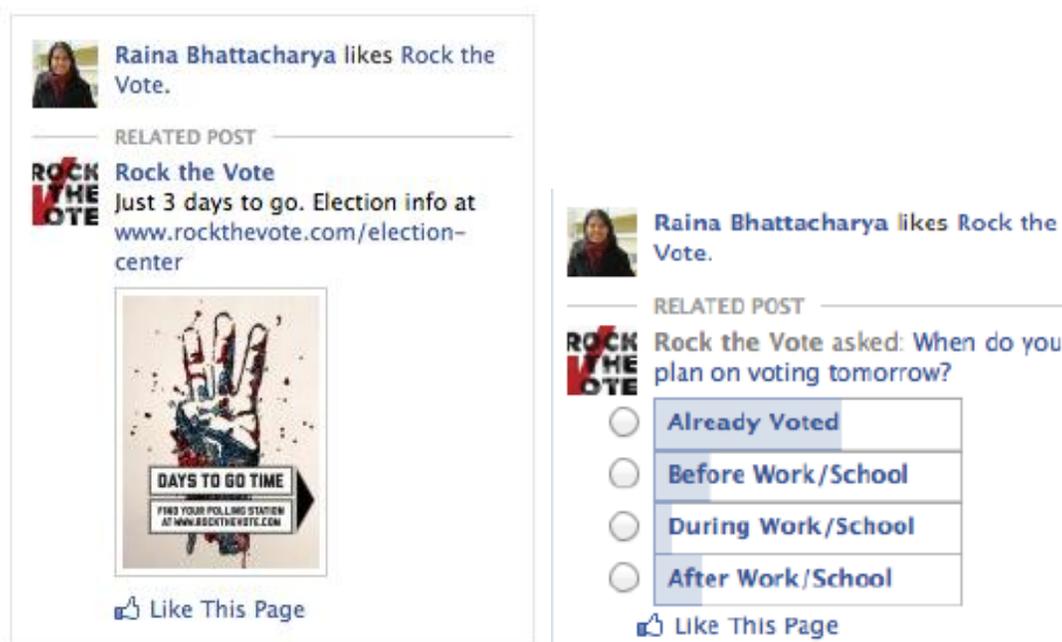


Table 1 provides descriptive statistics for each condition following random assignment. As expected, it shows no relationship between experimental condition and observable covariates. A logistic regression to predict treatment assignment as a function of the covariates confirms the proper implementation of the random assignment: a likelihood ratio test with 11 degrees of freedom (11 covariates x 1 treatments) finds $\chi^2=4.22, p=0.96$.

Table 1: Descriptive Statistics by Experimental Condition

	Control	Treatment
Age	33.53	33.57
Voted 2010	23.15%	23.07%
Voted 2008	52.81%	52.77%
Voted 2006	11.92%	11.93%
Voted 2004	25.57%	25.58%

¹ The exact advertisement that an individual in the treatment group received was not randomly assigned. Instead, treatment group individuals were eligible to receive any of these advertisements. The advertisement that a given individual in the treatment group viewed was determined by Facebook's algorithm.

% Democrat	30.65%	30.68%
% Republican	15.07%	15.02%
% Black	15.90%	15.87%
% Caucasian	69.87%	69.83%
% Hispanic	6.99%	6.93%
% Female	54.47%	54.42%
<i>N</i>	365,668	365,900

Note: Cells report the mean.

Results

After the election, an updated voter file was obtained from Catalist. Table 2 reports the voter turnout rate by experimental condition. We found that 56.485% voted in the treatment group compared to 56.491% in the control group. The difference between the two (0.003 percentage points) is statistically insignificant ($p = 0.96$), and so we cannot say that the GOTV advertisements boosted turnout among the treatment group.

Table 2: 2012 Voter Turnout by Experimental Condition

	Control	Treatment
2012 Voter Turnout	56.49%	56.485%
<i>N</i>	365,668	365,900
T-test: $p=0.96$		

CATE by Age

Our second research question was whether younger voters were more treatment responsive than older voters. Following the advice of Feller and Holmes (2009), we estimated a treatment-age interaction using a generalized additive model (GAM) to account for any non-linear heterogeneous treatment effects. The conditional average treatment effect for age is shown in Figure 4. The CATE plot demonstrates that the treatment had no effect on voter turnout across all ages. Statistical testing from the GAM confirms a treatment-by-age interaction effect with a p-value of 0.77. This suggests that the solution to the effectiveness of online advertising is not as simple as targeting younger audiences.

Figure 4: CATE by Age

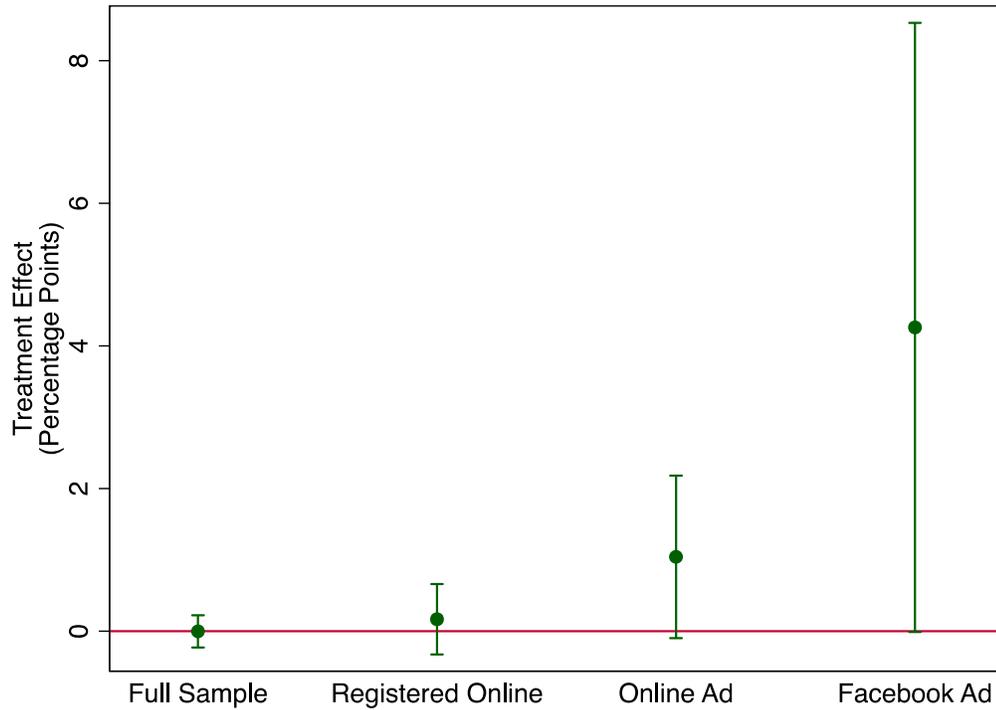
[SEE SEPARATE FILE FOR FIGURE 4]

*Note: Bands represent the 95% confidence interval of the treatment effect by age.**Exploratory Analysis*

After conducting an initial analysis, we received additional data from Rock the Vote reporting whether an individual had registered to vote with Rock the Vote online or in person, whether through an online advertisement, and if so, if that online advertisement was placed on Facebook. We hypothesize that because these individuals have revealed previously that they take action based on online advertisements, and the registration advertisement was sponsored by the same organization sponsoring the GOTV advertisement, then they might be more likely to vote after viewing a GOTV advertisement on Facebook. Due to the growing concern within the social sciences over “fishing,” or the post hoc determination of researchers to search for interesting results found in the data, we choose to present these results as “exploratory analysis” (Gerber et al. 2001, Gerber and Malhotra 2008, Ioannidis 2005, Humphreys et al. 2013). In addition, we replicated this experiment during the 2013 general election. These results are presented as Study 2.

The full universe for this experiment included all voters whom Rock the Vote registered during 2012 in targeted states. These registrations occurred both online at Rock the Vote’s website and in person at various events. A reasonable theory is that those who registered online ($n = 152,059$) may be more likely to be exposed to the ads and perhaps be more attentive. Thus, they may be more receptive to the Facebook ads than those who registered in person. We conducted exploratory analysis in which we limited our analysis only to those individuals who had registered online. These results are reported in Figure 5 and Table 3, while controlling for strata. While we find a larger point estimate than the average treatment effect (0.17 percentage points), this effect is nevertheless indistinguishable from zero ($p=0.507$).

Figure 5: Voter Turnout by Online Subgroups



Note: Error bars report the 95% confidence intervals. All estimates include the strata.

To take this hypothesis a step further, we further limited our analysis to the portion of the universe who registered in response to a voter registration advertising campaign Rock the Vote ran online earlier in the 2012 cycle ($n = 26,948$). Our theory is that those who have been found to interact with online ads in the past should be more likely to interact with them in the future. These results are presented in Figure 5 and in Table 3. They suggest that those who registered in response to a Rock the Vote online ad were about 1 percentage point more likely to vote than individuals in the control group who had similarly responded to an online ad ($p = 0.073$).

Finally, Rock the Vote collected data on which online ads these individuals clicked on in order to register to vote. When we examine only those individuals who registered to vote in response to a Rock the Vote Facebook advertisement ($n = 1,839$), we find that the Facebook voter mobilization advertisements increase turnout by a statistically significant 4.3 percentage points ($p = 0.051$). These results are presented in Figure 5 and in Table 3.

Table 3: CATE by History of Interacting with Online Advertisements

	Full Universe	Registered Online	Clicked Online Ad	Clicked Facebook Ad
<i>N</i>	731,568	152,059	26,948	1,839
Treatment <i>n</i>	365,900	78,410	13,400	946
Control <i>n</i>	365,668	73,649	13,548	893
Treatment Effect (SE)	-3.02 x 10 ⁻⁵ (0.0012)	0.0017 (0.0025)	0.0104 (0.0058)	0.0426 (0.0218)

Note: Each column reports the OLS regression results of 2012 voter turnout on treatment, controlling for strata. The Supplementary Information contains robustness checks for models with and without strata and covariates.

Study 2: Replication in 2013

Research Design

To address the fishing concerns noted above, we decided to replicate this experiment during the November 5, 2013 general election. Like the 2012 experiment, we randomly assigned voters who had registered with and provided an email address to Rock the Vote to an uncontacted control group ($n = 46,540$) or to a treatment group that received voter mobilization Facebook ads ($n = 46,513$). Full disclosure of exploratory analysis coupled with replication will help to overcome the fishing problems and a lack of confidence in empirical social science research (Laitin 2013).

The experimental population consisted of individuals who had registered to vote with Rock the Vote. These voters lived in 14 states that held both a November 2013 general election (see Table 4) and were an organizational priority for Rock the Vote.

Table 4: Experimental Universe by State

State	Control	Treatment
Arizona	92	119
Florida	2,065	1,991
Georgia	2,068	2,051
Massachusetts	2,453	2,496
Michigan	524	539
Minnesota	2,050	1,948
New Jersey	42	37
New York	7,484	7,487
North Carolina	4,108	4,268
Ohio	3,549	3,599
Pennsylvania	2,569	2,603
Texas	4,407	4,361
Virginia	13,353	13,291
Washington	1,776	1,723
Total	46,540	46,513

Like in Study 1, all of these individuals had previously provided an email address to Rock the Vote, mostly while in the process of registering to vote in person or online. Because we were primarily interested in studying the turnout effects of online Facebook ads on individuals who have

indicated a previous familiarity with online ads by registering through an online form, 98.31% of the experimental universe included individuals who had done so. Randomization was again conducted at the individual level within strata defined by mode of registration (online or in person) and overlap with two other non-digital experiments.²

Beginning 30 October 2014 and continuing through Election Day (5 November 2014), ads were served to treatment targets with email addresses matched to a Facebook user. Facebook informed us that 29,500 (63.4%) of the treatment group email addresses successfully matched to a Facebook user and that among those who matched 85.4% were served with at least one Rock the Vote advertisement, so 54% of the total treatment group was served advertising. These ads were placed in both the user’s sidebar and “News Feed.” In this replication study, Facebook allowed the placement of News Feed advertisements to everyone in the experimental universe and no longer imposed a restriction to those Facebook users who had “liked” Rock the Vote on Facebook or whose friends had liked Rock the Vote. Those who were served advertisements viewed them 48.33 times on average.

Table 5 provides descriptive statistics for each condition following random assignment. A logistic regression to predict treatment assignment as a function of the covariates confirms the proper implementation of the random assignment: a likelihood ratio test with 10 degrees of freedom (10 covariates x 1 treatments) finds $\chi^2=12.10$, $p=0.2786$.

Table 5: Descriptive Statistics by Experimental Condition

	Control	Treatment
Age	24.63	24.62
Voted 2012	68.06%	68.17%
Voted 2010	16.08%	15.95%
Voted 2008	43.18%	43.22%
% Democrat	22.05%	22.42%

² The overlap with one experiment, a volunteer GOTV phone call, is reported below. The remainder of these additional experiments are not reported here.

% Republican	5.54%	5.61%
% Black	21.56%	21.92%
% Caucasian	62.20%	62.48%
% Hispanic	10.00%	9.75%
% Female	60.62%	60.43%
N	46,540	46,513

Note: Cells report the mean.

Results

Table 6 reports the voter turnout rate by experimental condition. We found that 14.03% voted in the treatment group compared to 14.56% in the control group. The difference between the two (0.53 percentage points) is statistically significant ($p = 0.02$). Given our initial expectation that Facebook advertisements increase voter turnout, we understand this result as the null hypothesis of no positive effect cannot be rejected using a one-tailed test (Gerber and Green 2000).

Table 6: 2013 Voter Turnout by Experimental Condition

	Control	Treatment
2013 Voter Turnout	14.56%	14.03%
N	46,540	46,513
T-test: $p=0.021$		

CATE by History of Interacting with Online Advertisements

The primary aim of this 2013 experiment was to replicate the earlier exploratory analysis showing that individuals who had previously interacted with online advertisements are more responsive to an online GOTV advertisement than those with no prior history. Due to the growing concern within the social sciences over “fishing,” or the post hoc determination of researchers to search for interesting results found in the data, we choose to present the 2012 results as “exploratory analysis” and then validate these findings on a new, out-of-sample dataset generated through this separate experiment from 2013 (Gerber et al. 2001, Gerber and Malhotra 2008, Ioannidis 2005, Humphreys et al. 2013).

Table 7 presents the results by history of interacting with online advertisements in 2013. Overall, we find no evidence that Facebook GOTV advertisements are particularly effective among those who have previously clicked an online ad in general or a Facebook ad in particular in order to register to vote. Among the subgroup of voters who had clicked an online advertisements, this experiment had sufficient power to detect a minimum treatment effect of 1.8 percentage points, and among those who had clicked a Facebook ad, the experiment could reliably detect a minimum effect of 5.3 percentage points.³ While our 2013 replication was moderately underpowered compared to the 2012 exploratory analysis, which found a 1.0 percentage point effect among the online ad clickers and a 4.3 percentage point effect among the Facebook ad clickers, given that 2013 was an off-year election, we expected any GOTV effect to have a larger impact in 2013 than in 2012 (Arceneaux and Nickerson 2009). In addition, the Facebook match rate and average number of Facebook ad impressions was greater in 2013 than in 2012. Though an experimental universe can always be larger, these results nevertheless fail to confirm the 2012 exploratory analysis.

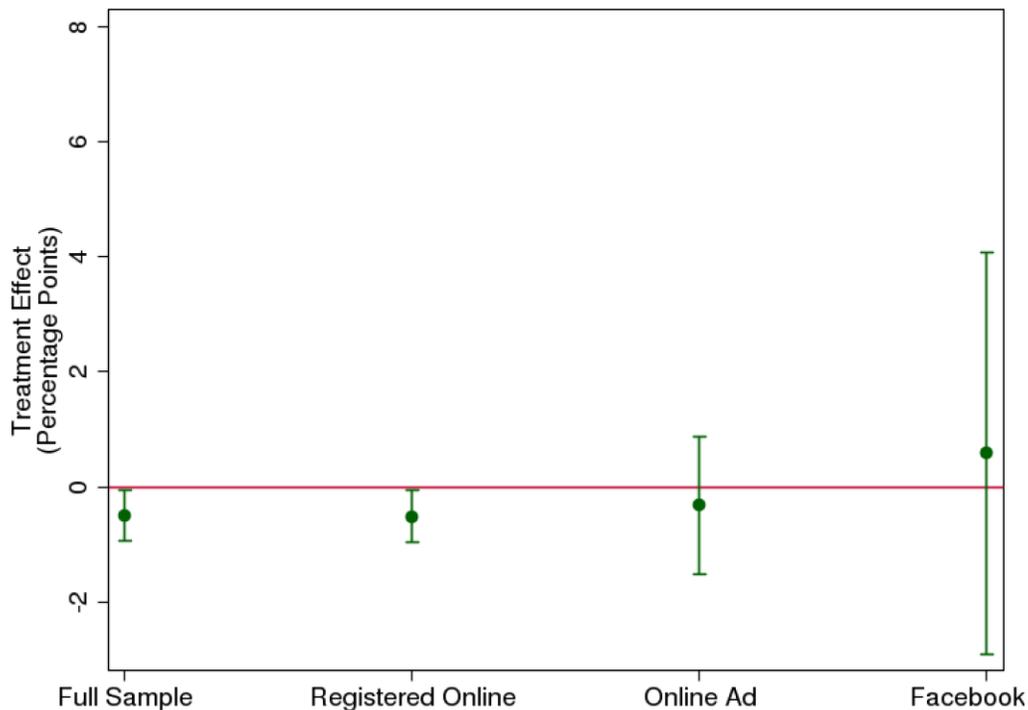
Table 7: CATE by History of Interacting with Online Advertisements

	Full Universe	Registered Online	Clicked Online Ad	Clicked Facebook Ad
<i>N</i>	93,053	90,786	12,539	1,410
Treatment <i>n</i>	46,513	45,361	6,305	724
Control <i>n</i>	46,540	45,425	6,234	686
Treatment Effect (SE)	-0.50 (0.23)	-0.51 (0.23)	-0.32 (0.61)	0.59 (1.78)

Note: Each column reports the OLS regression results of 2013 voter turnout on treatment, controlling for strata. The Supplementary Information contains robustness checks for models with and without strata and covariates.

³ Using a 15% baseline voter turnout rate, we calculate the minimum detectable treatment effect given the same size for the subgroup, 80% power and a 95% confidence interval.

Figure 6: Voter Turnout by Online Subgroups



Note: Error bars report the 95% confidence intervals. All estimates include the strata.

Overlap with the volunteer phone call experiment

During the 2013 election, Rock the Vote also conducted an experiment in which voters in Virginia were randomly assigned to receive a volunteer GOTV phone call in the days leading up to Election Day or no contact from Rock the Vote. This random assignment was conducted independent of the Facebook experiment, allowing us to test the comparative effectiveness of a volunteer phone call - a personal form of contact - and an impersonal, digitally-delivered Facebook advertisement. Table 8 reports the overlap in the number of voters receiving no Rock the Vote contact, Rock the Vote Facebook ads, a Rock the Vote volunteer phone call, and both the Facebook ad and the volunteer phone call.

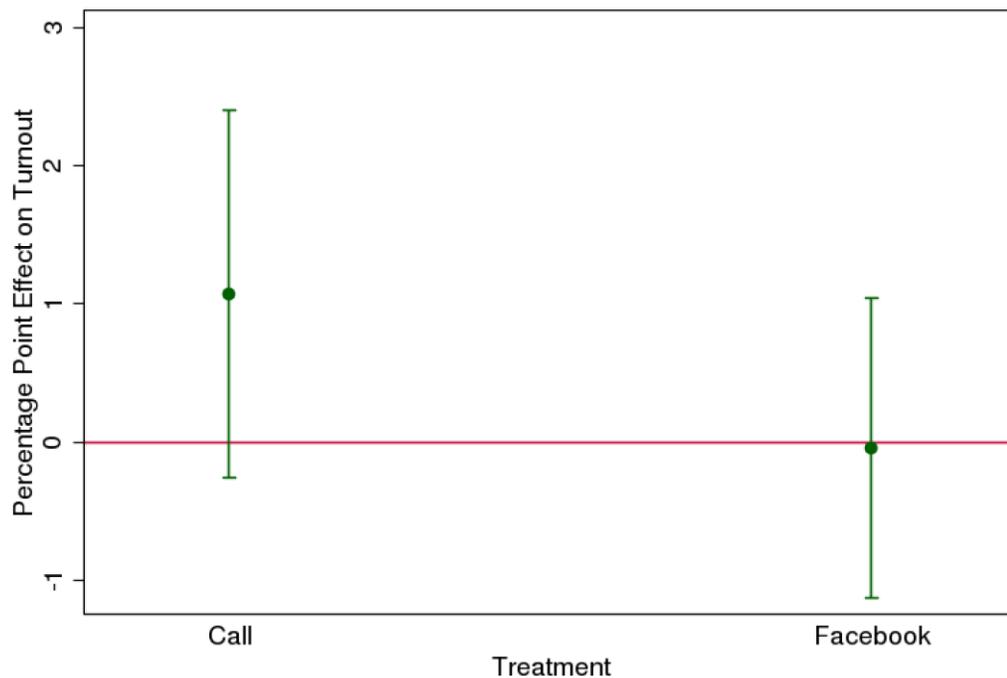
Table 8: Overlap Between Facebook Ad and Volunteer Phone Call Conditions

	No Facebook Ad	Facebook Ad
No Volunteer Call	9,458	9,335
Volunteer Call	2,444	2,537

Table 9 and Figure 7 present the turnout rates and treatment effects of the Facebook ads and volunteer phone calls, relative to the uncontacted control group. We find that while the Facebook ads decreased voter turnout by a statistically insignificant 0.04 percentage points ($p=0.940$), the volunteer phone calls increased turnout by 1.07 percentage points ($p=0.114$).

Table 9: 2013 Voter Turnout by Overlap in Experimental Condition

	Control	Facebook Only	Call Only	Both
2013 Voter Turnout	23.38%	23.63%	25.16%	24.00%
<i>N</i>	9,458	9,335	2,444	2,537

Figure 7: Treatment Effect of Volunteer Calls and Facebook Ads

Note: Error bars report the 95% confidence intervals.

We find is that Facebook advertisements have no effect on voter turnout, while volunteer phone calls do. This finding should be worrying for political campaigns and the civic minded, since the contact rate in the volunteer phone condition was only 23%, while the Facebook advertisements were likely viewed by around 54% of the experimental universe.⁴ As more Americans continue to move to cellphone-only households, the contact rates in future volunteer phone call conditions will likely continue to fall. The benefit of Facebook advertising is that it reaches a substantially larger swathe of young voters, yet, even though these voters view the advertisements, the advertisements have no effect on voter turnout. Thus Facebook advertisements manage to solve the problem of decreasing contact rates, but for now is not a sufficient substitute for the more personal contact of a volunteer phone call.

Discussion

The overall null results from both the 2012 and 2013 experiments are different than those found in the only other experiment conducted to date on Facebook ads' effect on voter mobilization, which found a statistically significant 0.39 percentage point boost from ads with a social message (Bond et al. 2012). With two showing no significant effects, and the other showing small but significant effects, taken together, these three experiments show that online advertising may not be a panacea for low-cost mobilization or mobilization of young voters.

Since the first modern field experiments in voter mobilization, results have consistently found that impersonal appeals fail to increase voter turnout (e.g., Gerber and Green 2000; Nickerson 2007; Gerber, Green, and Larimer 2008) and we confirm that in the 2013 experiment by comparing the effectiveness of Facebook advertisements to volunteer phone calls. It is only through

⁴ Facebook does not provide breakdowns in the contact rate. We only know that in the entire Facebook experimental universe, including states more than just Virginia, Facebook ads were viewed by 54%. It is unlikely that the rate would be much lower in this subgroup.

the active use of personal appeals and social tools that civic-minded organizations and political campaigns can increase voter participation.

As membership in traditional civic organizations has declined in recent decades, newer professionally managed groups have filled that void, but without the volunteers and activists who would normally provide the social means to mobilize the vote (Skocpol 2003). It was hoped that Internet, with its ability to cheaply reach millions of voters, would stanch the steady decline in civic participation (Hindman 2009). Yet, these experiments demonstrate that citizens cannot be mobilized and elections cannot be won merely on money, management, and the Internet, at least not yet. A vibrant democracy still requires an active *demos* capable of talking to and mobilizing each other. Instead, civic organizations are left with fewer volunteers and are having a more difficult time reaching voters.

However, this research also demonstrates how the technologies of voter contact are rapidly changing. Even between these two studies, the way in which Facebook ads were being delivered. And in 2014, campaigns and organizations wishing to use Facebook ads to reach their target electorate will no longer be limited to targeting abstract demographic groups or voters for whom they already have addresses; Facebook's advertising platform now allows for "multivariate" matching on name, address, and phone number, along with the email addresses that were used to match voters here. This will enable future researchers to run randomized controlled experiments without first needing to collect hundreds of thousands of email addresses. Changes in content of Facebook advertising may also make it a more effective medium in the near future, just as social pressure language turned GOTV postcards from a relatively ineffective voter mobilization tool to a stunningly powerful tactic (Gerber, Green, and Larimer, 2008). So perhaps in the future, we will also see technological and tactical changes to allow digitally-mediated campaign communication to better approximate that ideal of person-to-person contact, a digital *demos* for a digital age.

References

- Arceneaux, Kevin. 2007. "I'm Asking for Your Support: The Effects of Personally Delivered Campaign Messages on Voting Decisions and Opinion Formation." *Quarterly Journal of Political Science* 2:43-65.
- Arceneaux, Kevin, and Robin Kolodny. 2009. "Educating the Least Informed: Group Endorsements in a Grassroots Campaign." *American Journal of Political Science* 53:755-70.
- Arceneaux, Kevin and David Nickerson. 2009. "Who is Mobilized to Vote? A Re-Analysis of 11 Field Experiments." *American Journal of Political Science* 53:1-16.
- Barnard, Lisa and Daniel Kreiss. 2014. "A Research Agenda for Online Political Advertising: Surveying Campaign Practices, 2000–2012." *International Journal of Communication* 7:2046–2066
- 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
- Feller, Avi and Chris C. Holmes. 2009. *Beyond topline: Heterogeneous treatment effects in randomized experiments. Technical report*. Oxford, UK: University of Oxford.
- Garcia-Castañon, Marcela, Alison D. Rank, and Matt A. Barreto. 2011. "Plugged In or Tuned Out? Youth, Race, and Internet Usage in the 2008 Election." *Journal of Political Marketing* 10:115-38.
- Gerber, Alan S. and Donald P. Green. 2000. "The Effects of Canvassing, Telephone Calls, and Direct Mail on Voter Turnout: A Field Experiment." *American Political Science Review* 94:653-663.
- Gerber, Alan S., Donald P. Green, and David Nickerson. 2001. "Testing for publication bias in political science." *Political Analysis* 9:385–92.
- Gerber, Alan S. and Donald P. Green. 2008. *Get out the Vote: How to Increase Voter Turnout*, 2nd Ed. Washington, DC: Brookings Institution Press.
- Gerber, Alan S., Donald P. Green, and Christopher W. Larimer. 2008. "Social Pressure and Voter Turnout: Evidence from a Large-Scale Field Experiment." *American Political Science Review* 102:33-48.
- Gerber, Alan, and Neil Malhotra. 2008. "Do statistical reporting standards affect what is published? Publication bias in two leading political science journals." *Quarterly Journal of Political Science* 3:313–26.
- Green, Donald P. and Lynn Vavreck. 2008. "Analysis of Cluster-Randomized Experiments: A Comparison of Alternative Estimation Approaches." *Political Analysis* 16:138-152.
- Green, Donald P., Peter M. Aronow, and Mary C. McGrath. 2010. "Making Sense of 200+ Field Experiments on Voter Mobilization, Part I: Direct Mail." Presented at the Annual Meeting of the American Political Science Association.
- Hindman, Matthew. 2009. *The Myth of Digital Democracy*. Princeton: Princeton University Press.

- Humphreys, Macartan, Raul Sanchez de la Sierra, and Peter van der Windt. 2013. "Fishing, Commitment, and Communication: A Proposal for Comprehensive Nonbinding Research Registration." *Political Analysis* 21:1-20.
- Ioannidis, John P. A. 2005. "Why most published research findings are false." *PLoS Medicine* 2:696–701.
- Kaye, Kate. 2009. "Google grabbed most of Obama's \$16 million in 2008." *ClickZ*, January 6. <http://www.clickz.com/clickz/news/1703163/google-grabbed-most-obamas-usd16-million-2008> (March 28, 2014).
- King, Gary, James Honaker, Anne Joseph, and Kenneth Scheve. 2001. "Analyzing Incomplete Political Science Data: An Alternative Algorithm for Multiple Imputation." *American Political Science Review* 95:49-69.
- Laitin, David D. 2013. "Fisheries Management." *Political Analysis* 21:42-47.
- Lewis, Randall A. and Justin M. Rao 2013. "On the Near Impossibility of Measuring the Returns to Advertising." Working Paper.
- Link, Michael W., Michael P. Battaglia, Martin R. Finkel, Larry Osborn, and Ali H. Mokdad. 2007. "Reaching the U.S. Cell Phone Generation: Comparison of Cell Phone Survey Results with an Ongoing Landline Telephone Survey." *Public Opinion Quarterly* 71:814-839.
- McAndrew, Francis T. and Hye Sun Jeong. 2012. "Who Does What on Facebook? Age, Sex, and Relationship Status as Predictors of Facebook Use." *Computers in Human Behavior* 28:2359-65.
- Nickerson, David W. 2006. "Hunting the Elusive Young Voter," *Journal of Political Marketing* 5(3): 47-69.
- Nickerson, David W. 2007. "Quality is Job One: Professional and Volunteer Voter Mobilization Calls." *American Journal of Political Science* 51:269-82.
- Prior, Markus. 2007. *Post-Broadcast Democracy: How Media Choice Increases Inequality in Political Involvement and Polarizes Elections*. New York: Cambridge University Press.
- Ramachandran, Shalini. 2012. "Evidence Grows on TV Cord-Cutting." *Wall Street Journal*, August 7. <http://online.wsj.com/news/articles/SB10000872396390443792604577574901875760374> (March 28, 2014).
- Schlozman, Kay Lehman, Sidney Verba and Henry E. Brady. 2010. "Weapon of the Strong? Participatory Inequality and the Internet." *Perspectives on Politics*. 8:487-509.
- Skocpol, Theda. 2003. *Diminished Democracy: From Membership to Management in American Civic Life*. Norman, OK: University of Oklahoma Press.
- Smith, Aaron. 2013. "Civic Engagement in the Digital Age." *Pew Research Center* http://www.pewinternet.org/files/old-media//Files/Reports/2013/PIP_CivicEngagementintheDigitalAge.pdf (March 28, 2014).
- Smith, Aaron, Laurie Segall and Stacy Cowley. 2012. "Facebook reaches one billion users." *CNN Money*. <http://money.cnn.com/2012/10/04/technology/facebook-billion-users/> (4 October 2012.)

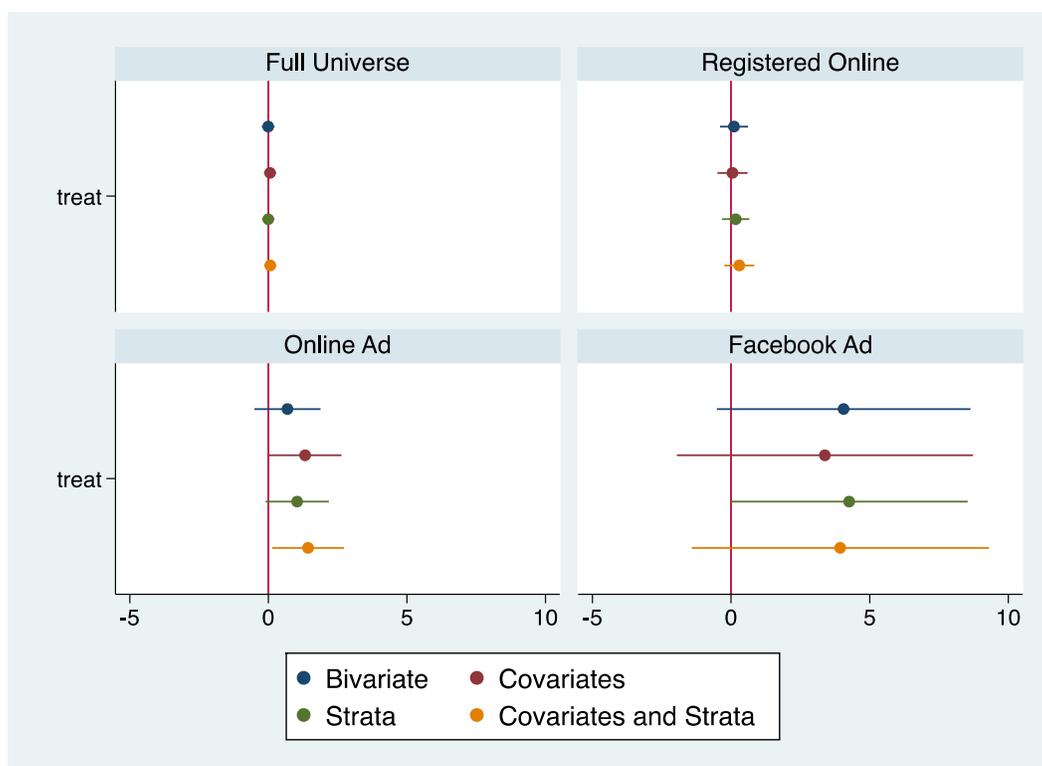
- Stampler, Laura. 2012. "Obama spent more on online ads than it cost to build the Lincoln Memorial." *Business Insider*, November 5. <http://www.businessinsider.com/infographic-obama-romney-final-ad-spend-2012-11> (March 28, 2014).
- Yenigun, Sami. 2012. "Presidential Campaigns Rock The Gamer Vote." *NPR*, October 1. <http://www.npr.org/2012/10/01/162103528/presidential-campaigns-rock-the-gamer-vote> (March 27, 2014).

Supplementary Information

Robustness Checks

SI Figure 1 presents the effect of the Facebook advertisements with and without strata and covariates for all subgroups. The covariates are age, vote history for 2004, 2006, 2008, and 2010, and binary variables for whether the voter is a registered Democrat, a registered Republican, black, white, Hispanic, and female.

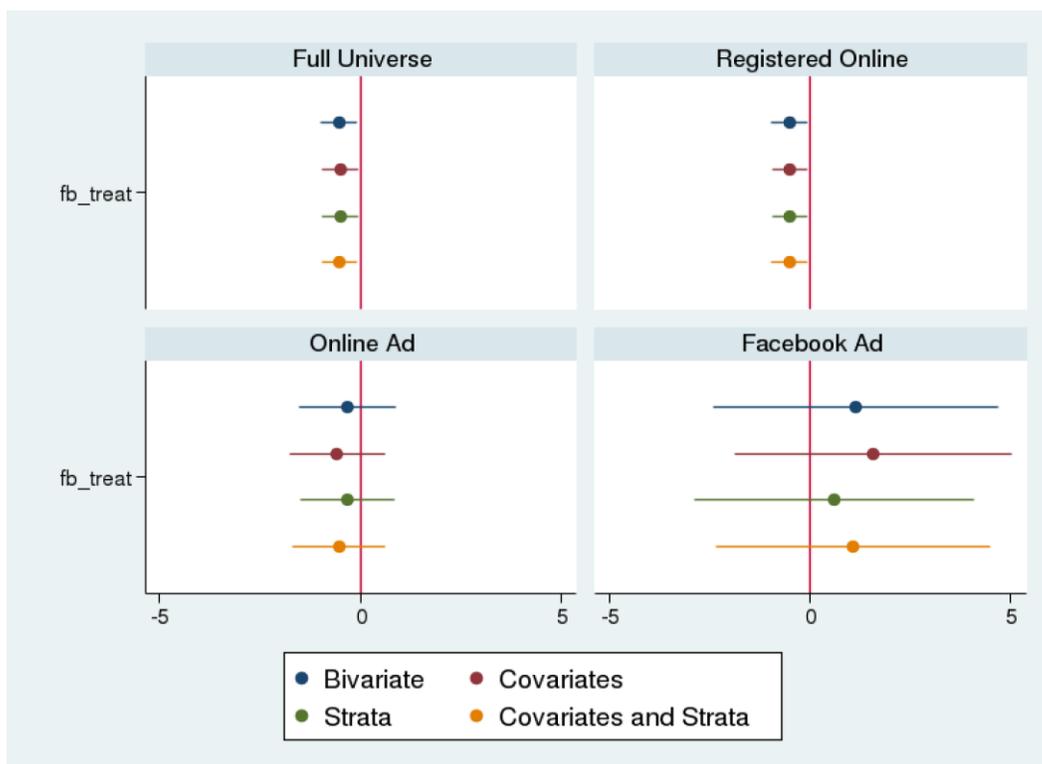
SI Figure 1: Robustness checks for Study 1



Note: Plotted dots denote the estimated treatment effect in percentage points of the Facebook advertisements on 2012 voter turnout. The error bars report the 95% confidence intervals.

SI Figure 2 presents the effect of the Facebook advertisements with and without strata and covariates for all subgroups. The covariates are age, vote history for 2008, 2010, and 2012, and binary variables for whether the voter is a registered Democrat, a registered Republican, black, white, Hispanic, and female.

SI Figure 2: Robustness checks for Study 2



Note: Plotted dots denote the estimated treatment effect in percentage points of the Facebook advertisements on 2013 voter turnout. The error bars report the 95% confidence intervals.