The Marginal Effects of Direct Mail on Vote Choice[*]

William Cubbison
George Washington University

March 25, 2015

Abstract

This paper presents a field experiment directly testing the marginal effects of direct mail on individual vote choice. Although most direct mail sent by campaigns is intended to change voter's candidate choice, the majority of prior research into the effects of direct mail has focused on turnout effects. This experiment is unique in two respects. Using three State Legislative districts in North Carolina during the 2014 election I survey individuals on vote choice after the treatment directly measuring individual level effects, and I vary the amount of mail received by voters. This experiment found no effect on vote choice from increased amounts of direct mail.

Prepared for presentation at the 2015 Annual Conference of the Midwest Political Science Association.

1 Introduction

The majority of direct mail sent by campaigns to voters is intended to influence vote choice. In spite of this, political science has largely focused on measuring get out the vote operations (GOTV). This focus has continued even with increased interest from partisans in using the most up to date scientific findings in their campaigns (Issenberg 2012, 2013). This paper presents a field experiment directly testing the marginal effects of direct mail on individual vote choice. Field experiments have become increasingly popular in measuring the effects of political campaigns, in part because they allow us to better isolate causal effects through the use of randomized treatments. This paper presents the results of an experiment measuring the effects on vote choice, and the decision to vote, from different amounts of direct mail from an independent interest group on unaffiliated voters in three North Carolina legislative elections.

Candidates and consultants often face a difficult strategic choice, whether to focus more resources on a small number of individuals or to broaden their campaign efforts and expend fewer resources per person. In other words, is it better to send six pieces of mail to one person or one piece of mail to six people? Although this might seem like a question that is only relevant to practitioners, it is relevant to political scientists interested in the effects of campaigns on vote choice, and understanding the process that goes into voter decision making. The debate over whether campaigns matter for electoral outcomes has produced a wide and diverse field of research, but the question of marginal effects has been largely unexplored. This experiment is an initial attempt to investigate marginal effects of campaign effects, but the topic deserves much more research to understand the effects of candidates and campaigns on public opinion and voter decision making.

To investigate this question I approached a number of former colleagues and associates starting in late 2013 to explore interest in conducting this experiment. The North Carolina Chamber of Commerce was willing to participate, which began a long process of creating the formal experiment. Funds for the post election survey were raised above and beyond the budget that was already dedicated by the Chamber to independent communications in these districts. These campaigns would have been included in the Chamber's political program without the experiment, and the experiment was conducted within the budget set before the decision to participate was made.

2 Campaign Effects Research

Prior research into the persuasive effects of campaign activity have been limited, and many of these studies have measured effects at the aggregate, rather than the individual, level. Studies of individual

level persuasion efforts have included experimental and observational studies of the content and delivery method used by campaigns (Adams and Smith 1980, Ansolabehere and Iyengar 1995, Arcenaux 2007). These studies show certain methods to be more effective at convincing voters, and indicate message content or communication methods that are more likely to be used by campaigns in certain electoral contexts.

Partisan mail does seem to influence vote choice, but how much is not a settled question. Gerber et al (2011a) attempted to determine the aggregate effect of partisan mail on vote choice through the use of precinct level data. Their study involved six mailers sent by an advocacy group supporting the Democratic Nominee for Attorney General in Kansas in 2006. The organization restricted mail to voters living in precincts below a set income threshold, allowing for a regression discontinuity design. The results indicated that for every 10% increase in the amount of mail sent to a specific precinct, the supported candidate saw an increase in vote share of 2.5 to 3.5%. Their results did not seem to indicate any GOTV effect from the persuasion mail, "At least for this specific mailing, the effect on vote share appears to come through persuasion of individuals who were already going to vote (Gerber et al p. 153)."

These results indicate a larger effect on vote choice than Gerber (2004), which may be due to different effects for incumbents and challengers on vote choice. In three experiments with Democratic incumbents in Connecticut and New Jersey, mail increased the vote share in one race and decreased it in two others. These results were all small and within the margin of error. Alternatively, in a local election, mail sent by a challenger had a larger, but still statistically insignificant increase on vote share.

During the 2012 election Doherty and Adler (2014) conducted an experiment to measure the effect of campaign mail on name recognition and the intention to vote. They found small increases in the amount of voters who recognized the name of the incumbent following an early mailing, and a very small increase in the self-reported intention to vote after a negative mailing. However, they did not find any GOTV effect among either groups, both those that received a mailing early in the election and those that received a communication close to the election.

The majority of campaign field experiments investigate one treatment or one level of treatments, which is beneficial for scientific purposes, but fails to replicate the choices faced by campaign professionals and fails to directly address the question of marginal returns. With limited resources campaign actors often face a tradeoff, whether to communicate using less resources with a broader range of the electorate or to overwhelm a smaller group of voters. Research showing diminishing marginal returns argues for fewer mailings to a broader section of the electorate. However, an overwhelming advantage is much more likely to affect vote choice than evenly matched communications from both sides, arguing for many more mailings to a smaller section of the electorate. This experiment is designed to directly address these questions.

My experimental design will also allow me to test for effects of varying amounts of direct mail on the decision to vote or sit out the election. A summary of prior GOTV experiments (Gerber et al 2008) does not show any clear pattern in past studies for turnout differences with varying levels of mailers. Some studies showed turnout increasing in direct proportion with the number of mailers, others showed unexpected patterns such as lower turnout with one mailer compared to none, higher with two mailers and lower again (but still higher than zero) with three mailers. Although these results are not a direct test of the research question here, they indicate that there may not be a significant benefit to sending multiple pieces of mail to voters. Alternatively, an attempt at this question (Green and Gerber 2008) indicates that six pieces of mail might be a pivot point. Their results showed increases in voter turnout from two to four to six pieces, but indicated six pieces caused higher turnout increases than eight pieces. There do not appear to have been any attempts to replicate this experiment, and no one else has attempted to apply a similar technique to persuasion mail.

3   Campaign Context

The overwhelming majority of studies that have found substantive effects of campaign activities have been conducted in non-competitive or non-salient elections. A non-competitive election leads to a number of benefits for researchers, but it likely presents an ideal scenario for finding campaign effects. Multiple researchers have found evidence for diminishing marginal returns of campaign spending, indicating that campaign activities have less of an effect on vote choice as overall spending increases (Green 1998, Jacobson 2006, Bonneau and Cann 2011). The effects found in prior studies (observational and experimental) for campaign effects should therefore likely be thought of as maximum possible effects. However, these effects could still be essential in determining the winners of electoral campaigns. For example, Gerber Kessler and Meredith (2011) estimate that six pieces of mail increased vote share by 1.5% to 3.5%. Even the lower estimate is larger than the gap between candidates in many competitive elections every year, but this effect is of much smaller interest when a candidate wins by 15 or more percentage points.

However, to answer the question of whether campaigns matter requires political scientists to investigate campaign activities in a competitive electoral context. In measuring the effects of the 2008 Obama campaign in Wisconsin, Bailey, Hopkins, and Rogers (2013) found potential evidence of a backlash, with lower support from voters who received appeals from the campaign. During a special local election that followed a tied general, Enos and Fowler (2014) found GOTV effects of about 7%, even though purely additive effects would have predicted more than a 20% increase in voting. A key purpose of this experiment is to expand research in this field, and test for campaign effects on persuasion in a competitive

electoral context.

Based upon this prior research I expect small, but still meaningful effects of persuasion. A competitive election with a large amount of spending is the worst possible scenario to find effects for campaign communications, and therefore prior estimates of effect sizes up to 3.5% (Gerber et al 2011a) are unlikely. However, even studies showing effects of spending that diminish to almost zero (Bonneau and Cann 2011) show some effect of a campaign, an effect that should be measurable through this experimental design. Experimental and observational studies of GOTV effects have found small or insignificant effects on the decision of whether to vote. Therefore, I expect that in this competitive environment my experiment will find small but meaningful effects on vote choice, but not find GOTV effects.

This experiment was conducted in three state legislative districts in North Carolina that were considered competitive at the beginning of the election season.[1] The districts were North Carolina Senate District 18, and House Districts 41 and 116. The nonpartisan research organization North Carolina Free Enterprise Foundation rated Senate 18 and House 41 as competitive and House 116 as Lean Republican.[2] Based upon work by Daily Kos Elections in 2012 Obama received 47.2% of the 2 party vote in Senate District 18, and 50.7% and 44.2% in House Districts 41 and 116 respectively.[3] Polling conducted by the Chamber of Commerce over the summer of 2014 indicated that all three of these races were close. Additionally, each district saw spending by outside organizations beyond this experiment, discussed below.

As anticipated, all of these elections were highly competitive, with the final outcomes decided by just a few percentage points. In House 41, incumbent Tom Murry faced a town council member, Gale Adcock. In House 116, incumbent Tim Moffitt faced Brian Turner, a first time candidate. In Senate 18, incumbent Chad Barefoot faced Sarah Crawford, a first time candidate with deep political ties. Table 1 reports district level results and spending by the candidates and outside groups. These races were among the most expensive in the state in 2014.[4] This highly competitive partisan environment makes this experiment unique in both design and electoral context.

[Insert table 1 about here.]

---

[1]The Chamber initially chose a mixture of two Republicans and two Democrats for the experiment, but their ideal messages would have introduced too much variation. Hence the decision to choose these three specific races.

[2]Conventional Voting Behavior District Ratings are based on every statewide election result since 2008. More details are available at http://ncfef.org/research-analysis/election-analysis/cvb-district-rating/

[3]Calculation of 2012 Presidential Election Results by Congressional District Statement of Methodology available at http://images.dailykos.com/i/user/303419/Statement_of_Methodology.pdf

[4]Data on direct candidate spending from the North Carolina State Board of Elections, candidate reports filed January 2015, independent spending reports best possible number from October 2014 and January 2015 organizational reports.

4   Experimental Design

This experiment was conducted on three state legislative races in North Carolina in cooperation with the North Carolina Chamber of Commerce's PAC. The North Carolina Chamber PAC is non-partisan, although generally conservative and pro-business. In 2014 they endorsed legislative candidates from both parties, and followed that endorsement with independent expenditures on behalf of a small bipartisan group in both the primary and general. The Chamber identified a pool of likely voters that they believed were persuadable, which was then randomly assigned to control and treatment groups. Voters in the treatment groups received 2, 6, 8 or 12 pieces of direct mail asking them to vote for the Republican incumbent in their district. Immediately following the election voters were surveyed on their vote choice to measure the effect of the mail on individual level vote choice.

The individuals included in this experiment were identified by the Chamber as individuals expected to vote and expected to be persuadable. This could potentially exaggerate the treatment effects compared to persuasion mail sent to a random selection of voters. However, it would not be rational for a campaign to send large amounts of mail to randomly chosen voters in a non-competitive election environment, and the purpose of this experiment is to test the effect campaigns have on election outcomes in a competitive context.

The difficulty in conducting a study in the effects of partisan mail is partially illuminated by the methods used within GOTV studies. Whether or not someone voted is independently verifiable, subject to some unknown (and likely random) level of error by the local or state election authority (Ansolabehere and Hersh 2012). However, individual vote choice can only be established through the use of surveys, which are subject to many sources of potential error. The remainder of this section will lay out the methods used to mitigate those biases.

As stated above, the voters in this experiment were chosen because they were expected to vote and were expected to be more persuadable than other voters. To be included in the experiment voters had to fall into one of three categories: registered to vote as an unaffiliated[5] voter and voted in either the 2006 or 2010 election and at least two of the most recent four general elections (2006, 2008, 2010 and 2012), registered to vote as an unaffiliated voter after 2010 and voted in the 2012 general election, or registered to vote as an unaffiliated voter after the 2012 election. This created a universe of roughly 25,000 individuals. The large amount of individuals was purposefully chosen in anticipation of low response rates to the survey instrument after the election.

The treatment groups included a control group, which received no direct mail, and groups receiving

---

[5]North Carolina allows voters to choose the Democratic, Libertarian or Republican Party when registering to vote. If they choose not to register with any of those parties they are considered to be "Unaffiliated" and are eligible to vote in either the Democratic or Republican primaries.

2, 6, 8 or 12 pieces of mail. These groups were chosen in part to replicate prior experiments, the majority of which have sent 1, 4, 6 or 9 pieces of direct mail. Further, Gerber and Green (2008) provide some evidence that effects on voter turnout went away after six mailings, setting a good base line for measuring diminishing marginal effects against. Keeping the message consistent also required even numbered groups, as discussed further below. The group receiving two mailings was split into a group receiving the mail early and a group receiving it late. This was done to attempt to measure the anticipated decay in effects, which has been found in studies of campaign communications in a variety of methods (Gerber et al 2011, Mitchell 2012, Hill et al 2013, Bartels 2014, Doherty and Adler 2014). The two early mailings were sent at the same time as the first two mailers for the group receiving twelve mailings, while the two late mailers were sent at the same time as the last two mailers for the six, eight and twelve mailing groups. The exact mail schedule is included in the appendix.

Using the voter file the Chamber of Commerce provided a list of targeted individuals in each district matching the characteristics above. The random assignment of treatment was stratified by district. In houses where the initial list identified more than one voter matching the experimental universe the mail vendor removed the second voter in the house according to the order they appear in the voter file, which is roughly alphabetically. This so-called "householding" process occurred before randomization. More individuals were placed in the groups receiving eight and twelve mailings to satisfy the request of the Chamber of Commerce. The control group for this experiment is 13.1% of the overall sample. Each group receiving two mailings was 8.6%, the group receiving six mailings was also 13.1% and the groups receiving eight and twelve mailings were 28.3% of the sample.[6]

A pre-experiment power calculation concluded that if 50% of the control group supports the targeted candidate and each mailer increases the vote share by 1%, except the group receiving two mailers early which was expected to decay, then the necessary sample size to find an effect is 1650.[7] All individuals were eligible for the follow up survey, however some individuals fall outside the sampling frame of the live caller telephone survey, even with a very high match rate with the survey firm able to identify phone numbers for a large majority of the targeted voters. Households without a land line were included if the survey firm was able to match a cell phone to that registered voter based on consumer data. The survey was conducted from the Wednesday immediately following the election until the next Tuesday. The majority of responses came in the first 3 days. Overall 1700 individuals responded, although fewer than 1500 said they voted for the Democratic or Republican candidate, falling below the goal for a properly powered experiment. Individuals who were not surveyed were still included in the analysis of GOTV

---

[6]Prior experiments show little consensus on the relative sizes of treatment groups. Bond et al (2012) placed almost 80% of their sample in one group, with roughly 10% each in a separate treatment group and a control group. Alternatively, many GOTV studies include a control group containing anywhere from 27% of the total sample to over 95% (Enos et al 2014, p. 284).

[7]Calculation was done with STATA 13's power command.

results presented below.

The mail was created by the Chamber based upon their messaging and a "benchmark poll"[8] conducted of a sample of all likely voters. Since the Chamber was unwilling to send the exact same mailer every time, the treatments were written and designed to control for different messaging effects. Half of the pieces contained messaging on education, and half contained messaging on the economy. The message was consistent across all pieces of mail, such that voters receiving twelve pieces of mail saw the same language and information as those receiving two pieces[9]. There is some variation in the design of each piece, but the language was kept as consistent as possible, such that voters receiving twelve pieces of mail received the same information as voters receiving only two pieces of mail.

As with any field experiment, there are a number of threats to the validity of the experiment. There is a potential compliance problem since there is no certainty that the targeted individuals received the mail sent to them, or if they received the mail they may not have read the mail. Additionally, the Chamber of Commerce was continuing to communicate its endorsements and message through other channels across the state, including paid television ads and through the media. The topics included in the direct mail, including teacher pay, tax cuts, and job growth were included in communications from other groups and organizations including candidates in these races and other campaigns.

Any study of partisan vote choice, and the effects of partisan campaign messaging, is subject to a much larger level of bias than GOTV studies, because vote choice is not independently verifiable. A source of error in conducting a survey on vote choice after the election is over reporting support for the winner (Wright 1993). A potential fix for this is asking earlier (Atkeson 1999) or even on Election Day (Jackson and Rainey 2012). Ideally the survey would have been conducted entirely on the day of the election, but that is not possible. However, the daily results for each district do not indicate that this was a problem, and do not provide evidence of a systematic bias against the losers. The effect of survey day on vote choice was not significant in any of the three elections nor in the pooled results. Therefore, this bias of over reporting support for the winner either didn't materialize in these elections, or the effect occurred prior to my survey taking place.

A potential problem for determining treatment effects is highlighted by Bailey et al (2013) and Enos et al (2014). In their study of the effects of the field program of the Obama campaign in Wisconsin in 2008 Bailey et al found a significant and non-random bias in survey non-response. Survey non-response is a significant problem if the likelihood of response is correlated with our variables of interest (Berinsky 2011). Bailey et al and Enos et al both indicate that prior differences between high turnout propensity

---

[8]A "benchmark poll" is a term of art used by campaigns and refers to a large survey, often testing 20-40 messages to determine campaign message and strategy. This can also be referred to as "message testing" and is often confused with push polls by the public. An explanation of the difference can be found at: `http://www.aapor.org/AAPORKentico/Education-Resources/Resources/AAPOR-Statements-on-Push-Polls.aspx`

[9]A limited set of examples are included in the appendix, more can be made available if requested.

and low turnout propensity voters are exacerbated by campaign treatments and response rates are reduced for lower turnout voters following campaign activity. However, Baiely et al "find no similar pattern of heterogeneous treatment effects on survey response for those who received campaign mailings" (p 16). Alternatively, there is some evidence (Enos et al 2014, Zaller 1992) that voters with extremely high or low political awareness or activity are the least likely to be affected by campaign activity. This could potentially affect the survey response rates for this experiment, since the treatment population is chosen under the expectation that campaign activity can change their vote choice.

The timing of political communication in relation to both vote choice and GOTV is the subject of some research, but this study faces some limitations in this area. It is obviously impossible for all treatment groups to receive the direct mail at the exact same time. Gerber et al (2011b) and Hill et al (2013) show a significant decay in the effects of campaign activity. A failure to control for this timing problem could pose a problem for the internal validity of this study, as results could be measuring timing instead of dosage. In an attempt to test for this effect the group receiving two mailers was split into an early and late group. This allows for measurement of the varying timing of the treatments. The results from Hill et al (p. 522) are that large effects have a half-life of approximately four days at the Presidential level, but the effects do not disappear entirely after six weeks. Alternatively, Mitchell (2011) showed that new information displaced old information rather quickly, such that the effects of information lasted less than a week. This can be controlled for by keeping the information constant, but there is still potential for decay in treatment effects.

Following the election the list of voters was independently matched to vote history and other information by Catalist. Catalist, LLC is a data vendor maintains an augmented voter file available for purchase by left-of-center organizations and candidates, and for academic research. Their data has been used by Ansolabehere and Hersh 2012, Hersch and Schaffner 2013 and other researchers. Data on partisanship and education is modeled by Catalist based upon consumer data and surveys. Although the effects on GOTV are not a core purpose of this study, it does not make sense to not take advantage of the data and to explore these effects.

5 Results

5.1 Persuasion

Given the unsurprisingly low survey response rates, responses were pooled across districts[10]. The post-election survey did not find substantive or significant effects of mail on vote choice. There is some variation in the overall support for the Republican incumbents among group means, as can be seen in Figure 1, but these differences are not significantly different from the control group. To say the effect is not significant wouldn't necessarily preclude some effect that might have been found with a larger sample, but the variation lacks any systematic pattern, and it is unlikely that a larger sample would have found substantively or statistically significant effects of direct mail in this electoral context.

[Insert figure 1 here.]

After comparing overall treatment group means I conducted a probit analysis of candidate choice among the 1413 respondents who gave a valid answer on vote choice[11]. After controlling for self-reported partisanship and gender, the experimental treatments still did not appear to be statistically or substantively significant. There is reason to expect differential effects for in-partisans and out-partisans from campaign communications(Bartels 2014, Taber and Lodge 2006). To control for this I interacted partisanship and the direct mail treatment and again found no effect for the direct mail. The survey measured partisanship on a seven point scale, but the differences between Strong Republicans and Independents who lean Republican were statistically indistinguishable, therefore the results for this presented below combine these answers into a three point scale for partisanship.

[Insert figure 2 here.]

The results above are the pooled responses on candidate support across all three races. These results indicate that controlling for the observed gender and partisanship there is no evidence that the amount of mail received by respondents affected their vote choice. Although there is some noise in the individual level, the group differences are not statistically significant. The results of a contrast test for differences

---

[10]There is reason to test for systematic differences between the survey takers and those who did not respond (Bailey and Rogers 2012, Enos, Fowler and Vavreck 2014), but this has not been done yet.

[11]Of 1716 respondents, 19 claimed to have voted for a third party candidate although none of the three races had a third party candidate, 116 said they did not vote, and 168 either said they didn't know who they voted for or refused to answer the question. A future version of this paper should test for whether those individuals differ from the rest of the survey takers in important ways.

across treatment groups within different levels of partisanship do not show statistically significant differences comparing the treatment effects with the control group.

Unsurprisingly, individuals who identified as a Republican largely voted for the Republican candidate, while individuals who identified as a Democrat mostly supported the Democratic candidate. The survey asked for partisanship on a seven point scale, but the differences in vote choice between strong Republicans and Independents who lean towards the Republican Party were insignificant, similarly for Democrats. The association between party choice and candidate choice was weaker among independents who only chose a party when pushed, but was still strong. As an example of the specific predicted probability and the lack of effect of direct mail, the predicted probability that a voter who identifies as Republican and received no mail supported the Republican candidate is 93% with a confidence interval of 3.3%, while the predicted probability for a Republican who received twelve pieces of mail is 88% with a confidence interval of 2.7%.

The necessary sample size calculated under the assumption that the control groups would have 50% support, the 2 early mailers would have no effect, and each mailer would increase support by 1% was 1650. This experiment is therefore under powered, and the lack of statistically significant effects could be based on that. However, if that was the problem it should be expected to see a small and correctly signed effect, which was not found. There are also potential threats to this experiment, including spending on both sides of the aisle with similar talking points or the Democratic response. Additionally, there is no guarantee that targeted individuals received the treatment.

5.2   GOTV

Similar to the persuasion results, the data also showed no substantive or significant effects when looking at the overall treatment group means for whether or not individuals voted. These results should be considered as intent to treat, since there is no way to know whether individuals actually received the direct mail or whether they read the mail if it was received by the specific targeted individual.

[Insert figure 3 here.]

In addition to measuring the basic group means I conducted a probit analysis of whether or not individuals voted, controlling for observed values of gender, past voting behavior and whether they are college educated. Past voting behavior is treated as a continuous variable, with a sum of the number of general elections from 2006-2012 that the individual voted in (ranging from 0-4). Partisanship is

modeled by Catalist on a 0-100 scale and as a categorical variable with Strong and Weak Partisans and Independents. For this analysis I used the categorical variable version.

[Insert figure 4 here.]

The above results show no effect of direct mail on the predicted probability that an individual voted based on their partisanship and treatment group, controlling for observed gender, past voting behavior and college education. Consistent with prior research on voting behavior, stronger partisans were more likely to vote controlling for other characteristics than weak partisans and independents. Although it visually appears that the mail might have had some effect on the decision to vote for weak partisans and independents, the differences are not statistically significant. With over 25,000 individuals it is unlikely that this result is due to a lack of power.

6 Discussion and Conclusion

This experiment failed to provide evidence for an effect of direct mail on the decision of individuals to vote, and the decision of which candidate to support. Spending on direct mail is large, representing almost \$231 million in the 2012 presidential election (Washington Post 2012, Eggen 2012). Prior experiments have estimated persuasion effects as large as a 3.5% increase in vote share for the supported candidate (Gerber 2004, Gerber et al 2011a). In light of this high spending and prior research, this result is surprising, and raises the question of whether and why candidates should invest precious resources in direct mail. However, these results should be read with caution, and may be highly contextual.

The 2014 election in North Carolina featured a large amount of political spending, including over \$100 million in disclosed spending on the United States Senate race, and as stated above, each of these races included a large amount of spending. In an era of polarization it should not be surprising that there is a correlation between support for the Republican U.S. Senate nominee Thom Tillis and support for these three North Carolina General Assembly candidates of 0.82, a number driven artificially lower by the high levels of 3rd party support in the US Senate race, and refusals from these survey respondents. It is entirely plausible that this high level of spending drove results in these down ballot races, and along with the high level of spending in these races may have caused a high level of mail recipients to discard or ignore the direct mail sent as part of this experiment.

There is also a possibility that the null result is driven by the message included in the direct mail. The mail was split evenly between communications about education and the economy, which consistently

polled as the most important issues in North Carolina and in these districts, and were expected to be highly salient. However, the communication was positive, and on behalf of an incumbent. It is possible that negative mail or mail on behalf of the challenger in the same environment would have had an effect on vote choice. Furthermore, these issues were front and center during the election campaign, and it is plausible that members of the control group were exposed to information similar to the experimental treatment by the Republican incumbent's campaign, by the Democratic challenger's campaign, or by other outside groups.

The difficulty with measuring individual level vote choice is the high cost of obtaining a large enough sample to make strong conclusions. Despite starting with a very large sample, this experiment falls short of the desired level of power to measure persuasion effects, but still provides compelling evidence for a lack of an effect of direct mail on both vote choice. The large sample does provide compelling evidence for no effect on the decision to vote, but future experiments utilizing this or a similar design will have similar difficulties in measuring persuasion effects at the individual level. This is especially true if there is an effort to replicate this design across multiple elections.

This paper presents a unique and important addition to the literature on campaign effects. The competitive elections present a unique and important context for addressing the importance of campaign activities on electoral outcomes. Further, this is one of a very small number of experiments to test the marginal effect of campaign communications on political outcomes, and the first to test dosage effects of direct mail on both persuasion and GOTV effects. The failure to find effects of any amount of mail in this context should raise serious questions for campaign professionals and for researchers interested in whether campaigns matter.

Works Cited

Adams, William C., and Dennis J. Smith. 1980. "Effects of Telephone Canvassing on Turnout and Preferences: A Field Experiment." *The Public Opinion Quarterly* 44 (3): 389-95.

Ansolabehere, Stephen, and Eitan D. Hersh. 2012. "Validation: What Big Data Reveal about Survey Misreporting and the Real Electorate." *Political Analysis* 20 (4): 437-59.

Ansolabehere, Stephen, and Shanto Iyengar. 1995. *Going negative : how attack ads shrink and polarize the electorate.* New York: Free Press.

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 (1): 43-65.

Arceneaux, Kevin, and David W. Nickerson. 2009. "Who is Mobilized to Vote? A Re-Analysis of 11 Field Experiments." *American Journal of Political Science* 53 (1): 1-16.

Atkeson, Lonna Rae. 1999. ""Sure, I Voted for the Winner!" Overreport of the Primary Vote for the Party Nominee in the National Election Studies." *Political Behavior* 21 (3): 197-215.

Bailey, Michael, Daniel J. Hopkins, and Todd Rogers. 2013. "Unresponsive and Unpersuaded: The Unintended Consequences of Voter Persuasion Efforts." Working Paper `http://scholar.harvard.edu/files/todd_rogers/files/unresponsive.pdf`.

Bartels, Larry M. 2014. "Remembering to Forget: A Note on the Duration of Campaign Advertising Effects." *Political Communication* 31 (4): 532-44.

Berinsky, Adam J. 2011. "Representative Sampling and Survey Non-Response." In *The Oxford Handbook of American Public Opinion and the Media*, ed. Robert Y. Shapiro and Lawrence R. Jacobs. New York: Oxford University Press.

Blumenthal, Mark. 2009. "The Big Rock Candy Mountain of Data." `http://www.pollster.com/blogs/the_big_rock_candy_mountain_of.html` .

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 (7415): 295-8.

Bonneau, Chris W., and Damon M. Cann. 2011. "Campaign Spending, Diminishing Marginal Returns, and Campaign Finance Restrictions in Judicial Elections." *The Journal of Politics* 73 (4): 1267-80.

Doherty, David, and E. Scott Adler. 2014. "The Persuasive Effects of Partisan Campaign Mailers." *Political Research Quarterly* 67 (3): 562-73.

Dan Eggen, "Direct Mail Still a Force in Campaigns," *The Washington Post.* `http://www.washingtonpost.com/politics/decision2012/direct-mail-still-a-force-in-campaigns/2012/10/12/24f6f830-0bf9-11e2-bb5e-492c0d30bff6_story.html`.

Enos, Ryan D., and Anthony Fowler. 2014. Aggregate Effects of Large-Scale GOTV Campaigns: Evidence from 400 Million Voter Contacts. Working Paper `http://people.hmdc.harvard.edu/~renos/papers/EnosFowlerGOTV/EnosFowler_AggregateGOTV.pdf`.

Enos, Ryan D., Anthony Fowler, and Lynn Vavreck. 2014. "Increasing Inequality: The Effect of GOTV Mobilization on the Composition of the Electorate." *The Journal of Politics* 76 (1): 273-88.

Gelman, Andrew, and John Carlin. 2013. "Beyond Power Calculations to a Broader Design Analysis, Prospective Or Retrospective, using External Information." Working Paper `http://www.stat.columbia.edu/~gelman/research/unpublished/retropower.pdf`.

Gerber, Alan. 1998. "Estimating the Effect of Campaign Spending on Senate Election Outcomes using Instrumental Variables." *The American Political Science Review* 92 (2): 401-11.

Gerber, Alan S. 2004. "Does Campaign Spending Work?: Field Experiments Provide Evidence and Suggest New Theory." *American Behavioral Scientist* 47 (5): 541-74.

Gerber, Alan S., James G. Gimpel, Donald P. Green, and Daron R. Shaw. 2011b. "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): 135-50.

Gerber, Alan S., and Donald P. Green. 2000. "The Effect of a Nonpartisan Get-Out-the-Vote Drive: An Experimental Study of Leafletting." *The Journal of Politics* 62 (3): 846-57.

Gerber, Alan S., and Donald P. Green. 2012. *Field Experiments: Design, Analysis, and Interpretation.* New York: W. W. Norton & Company.

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 (1): 33-48.

Gerber, Alan S., Daniel P. Kessler, and Marc Meredith. 2011a. "The Persuasive Effects of Direct Mail: A Regression Discontinuity Based Approach." *The Journal of Politics* 73 (1): 140-55.

Green, Donald P., Brian R. Calfano, and Peter M. Aronow. 2014. "Field Experimental Designs for the Study of Media Effects." *Political Communication* 31 (1): 168-80.

Green, Donald P., and Alan S. Gerber. 2008. *Get Out the Vote: How to Increase Voter Turnout.* 2nd ed. Washington, D.C.: Brookings Institution Press.

Green, Donald P., Alan S. Gerber, and David W. Nickerson. 2003. "Getting Out the Vote in Local Elections: Results from Six Door-to-Door Canvassing Experiments." *The Journal of Politics* 65 (4): 1083-96.

Groves, Robert M. 2009. *Survey methodology.* Hoboken, N.J: Wiley.

Hersh, Eitan D., and Brian F. Schaffner. 2013. "Targeted Campaign Appeals and the Value of Ambiguity." *The Journal of Politics* 75 (2): 520-34.

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

Issenberg, Sasha. 2012. *The victory lab: the secret science of winning campaigns.* New York: Crown.

Issenberg, Sasha. 2013. A more perfect union: how President Obama's campaign used big data to rally individual voters. *Technology Review, Inc.*

Jacobson, Gary C. 2006. "Campaign Spending Effects in U.S. Senate Elections: Evidence from the National Annenberg Election Survey." *Electoral Studies* 25 (2): 195-226.

Mitchell, Dona-Gene. 2012. "It's about Time: The Lifespan of Information Effects in a Multiweek Campaign." *American Journal of Political Science* 56 (2): 298-311.

Nickerson, David W. 2008. "Is Voting Contagious? Evidence from Two Field Experiments." *American Political Science Review* 102 (1): 49-57.

Rainey, Carlisle, and Robert A. Jackson. 2012. "Modeling Misreports in Self-Reported Vote Choice Data." Working Paper `http://www.carlislerainey.com/papers/misreports.pdf`.

Rogers, Todd, and Joel A. Middleton. 2012. "Are Ballot Initiative Outcomes Influenced by the Campaigns of Independent Groups? A Precinct-Randomized Field Experiment." Working Paper `http://dash.harvard.edu/bitstream/handle/1/9830357/RWP12-049_Rogers.pdf?sequence=1`.

Taber, Charles S., and Milton Lodge. 2006. "Motivated Skepticism in the Evaluation of Political Beliefs." *American Journal of Political Science* 50 (3): 755.

Wright, Gerald C. 1993. "Errors in Measuring Vote Choice in the National Election Studies, 1952-88." *American Journal of Political Science* 37 (1): 291-316.

Zaller, John, 1949. 1992. *The nature and origins of mass opinion.* New York: Cambridge University Press.

"2012 Presidential Campaign Finance Explorer," *The Washington Post.* `http://www.washingtonpost.com/wp-srv/special/politics/campaign-finance/`.

Figures

Table 1: District Spending and Election Results

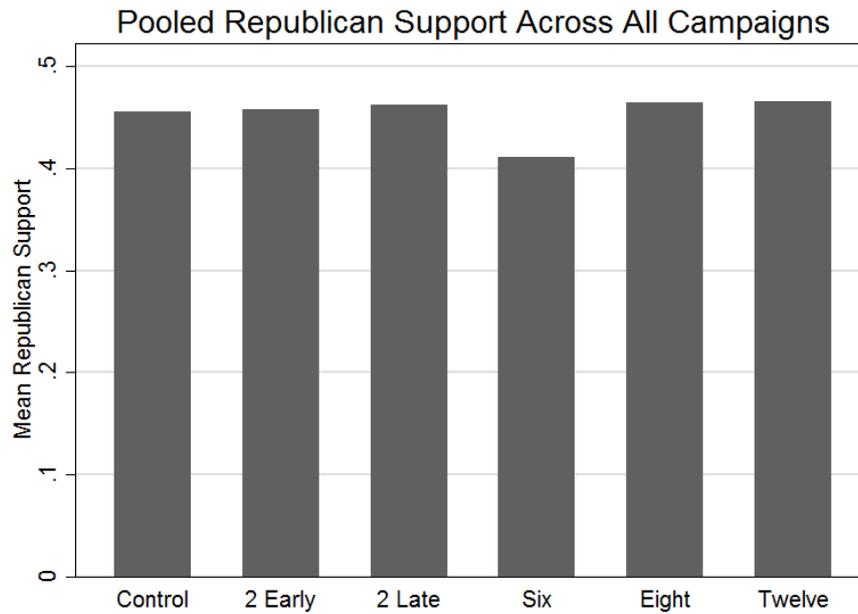| District	House 41	House 116	Senate 18
GOP % of Vote	48.68%	48.09%	52.89%
GOP Spending	$740,781	$698,873	$1,053,549
DEM Spending	$528,754	$559,509	$932,717
Pro-GOP Independent Spending	$53,482	$135,408	$244,750
Pro-DEM Independent Spending	$746,489	$2,180,127	$2,987,796
Total Voters	29,543	25,619	65,507

Figure 1: Republican Support

Pooled Vote Choice
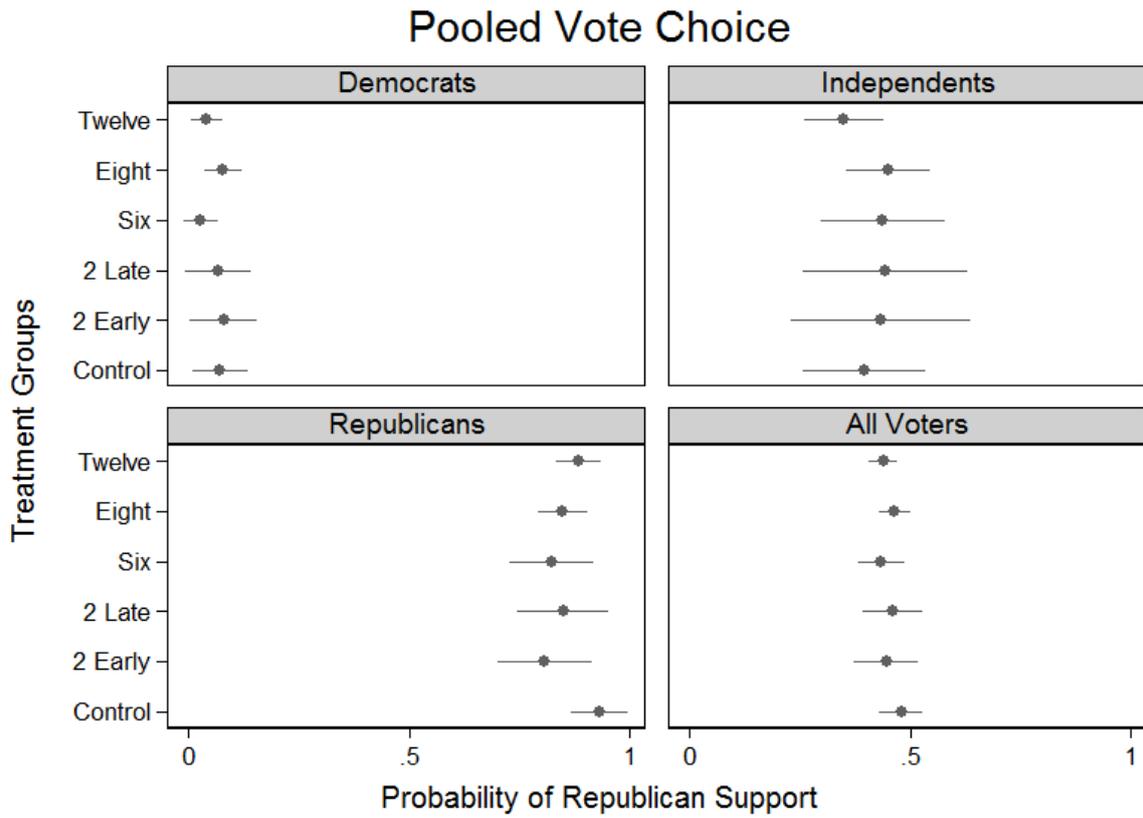

Figure 2: Marginal Effects on Vote Choice by Partisanship
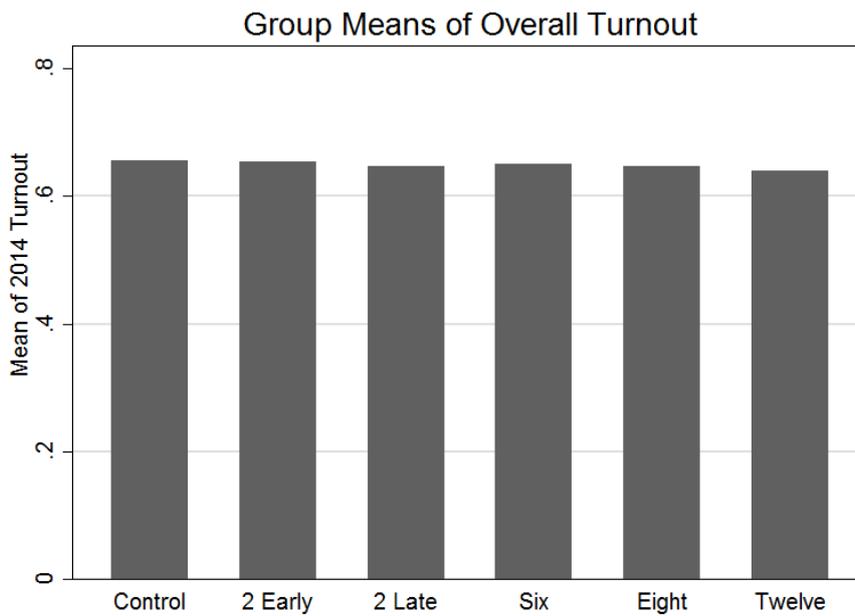

Figure 3: Voter Turnout

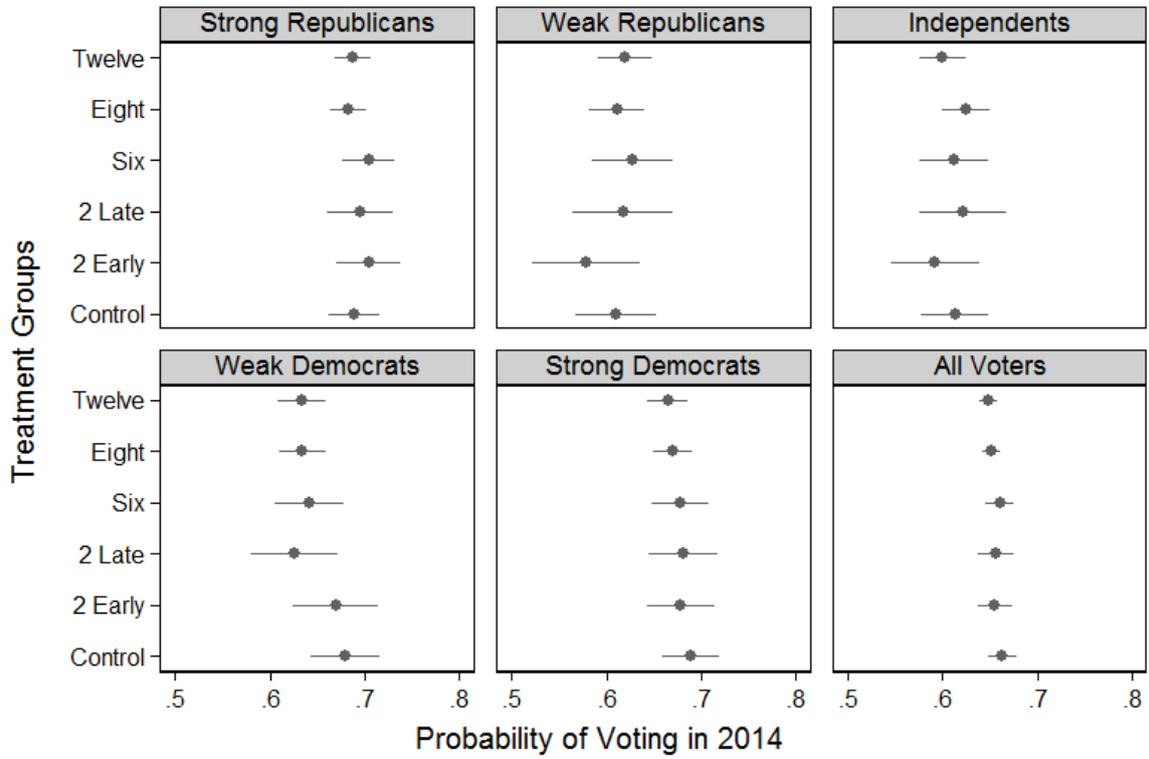Figure 4: Marginal Effects on Voter Turnout by Partisanship

Appendix

Direct Mail Schedule:

Date	Content	Treatment Groups
9-24	Economy 1	2E, 12
9-27	Education 1	2E, 12
10-1	Economy 2	12
10-4	Education 2	12
10-8	Economy 3	8, 12
10-11	Education 3	8, 12
10-15	Economy 2	6, 8, 12
10-18	Economy 3	6, 8, 12
10-22	Education 2	6, 8, 12
10-25	Education 3	6, 8, 12
10-29	Economy 1	2L, 6, 8, 12
11-1	Education 1	2L, 6, 8, 12

Survey Question Wording:

Q0. Did you vote in this year's elections? [1 Yes, voted. 2 Did not vote. 3 Don't Know/Refused]

Q1. Now, some people decide early in a campaign how they will vote. Others make their decisions just before the election. When would you say you made your FINAL decision on which candidate you were going to vote for your State Representative or State Senator? [1 On Election Day. 2 During the last few days before the election. 3 During the last two weeks in October. 4 During the first two weeks in October. 5 Before October. 6 Don't Know/Didn't Vote. 7 Refused]

Q2 In the race for U.S. Senate between Republican Thom Tillis and Democrat Kay Hagan, did you vote for Tillis or Hagan? [1 Hagan. 2 Tillis. 3 Other/Someone Else. 4 Didn't Vote. 5 Don't Know. 6. Refused]

Q3A As you know, we just had an election this week. One of the races in your area was for State Senate. The candidates were Republican Chad Barefoot and Democrat Sarah Crawford. Did you vote for Barefoot or Crawford? [1 Barefoot. 2 Crawford. 3 Other/Someone Else. 4 Didn't Vote. 5 Don't Know. 6 Refused.]

Q3B replaced names with Republican Tom Murry and Democrat Gale Adcock, Q3C replaced names with Republican Tim Moffitt and Democrat Brian Turner. Questions were asked depending on which district the respondent lived in, determined from voter file before the survey began.

Q4 I'd like your view on an organization that is in the news from time to time. Please tell me if you are aware or not aware of this organization: the North Carolina Chamber. [1 Aware, Very Favorable. 2 Aware, Somewhat Favorable. 3 Aware, Somewhat Unfavorable. 4 Aware, Very Unfavorable. 5 Aware, No Impression. 6 Aware, Refused Impression. 7 Not Aware.8 Refused.]

Q5 If a candidate for elected office had been endorsed by the North Carolina Chamber, would you be more likely or less likely to support that candidate, or would it make no difference to your vote? [1 Much more likely to support candidate. 2 Somewhat more likely to support candidate. 3 Somewhat less likely to support candidate. 4 Much less likely to support candidate. 5 No difference. 6 Refused.]

Q6 In politics today, do you normally think of yourself as a Republican, an independent, or a Democrat?[1 Strong Democrat. 2 Not-so-strong Democrat. 3 Independent, Leaning Democrat. 4 Independent, No Preference. 5 Independent, Leaning Republican. 6 Not-so-strong Republican. 7 Strong Republican. 8 Refused.]

Q7 When thinking about politics today, do you normally consider yourself to be very conservative, somewhat conservative, moderate, somewhat liberal, or very liberal? [1 Very Conservative. 2 Somewhat Conservative. 3 Moderate. 4 Somewhat Liberal. 5 Very Liberal. 6 Don't Know. 7 Refused.]

Q8 Is your racial or ethnic heritage white, black, Hispanic or something else? [1 White. 2 Black. 3 Hispanic. 4 Asian. 5 American Indian/Native American. 6 Other. 7 Not Ascertained.]

Q9 Gender [1 Male. 2 Female.]

House 41 "Economy 1" Back

House 41 "Economy 1" Front

Senate 18 "Economy 2" Back

23

Lower Taxes
More Jobs

FOR HARD-WORKING NORTH CAROLINA
FAMILIES AND SMALL BUSINESSES

LOWER TAXES
- Chad Barefoot voted for the largest tax cut in North Carolina history.[1]
- The Barefoot Plan is a FLAT TAX of 5.75% – a 25 % rate cut – a fair plan for everybody.[1]
- The plan means more money in your paycheck to save, take care of your family, or spend as you choose.
- The plan means more money for small businesses to invest in North Carolina and create good jobs.

MORE JOBS
- Over 250,000 new jobs created.[2]
- North Carolina had the largest drop in unemployment in the country last year.[3]
- Chad Barefoot voted to reform North Carolina's unemployment system, supporting legislation that paid off debt to the federal government and connected workers with good jobs.[4]

NORTH CAROLINA
CHAD BAREFOOT
SENATE

[1] North Carolina General Assembly – House Bill 998 Information/History (2013-2014 Session). (n.d.). Retrieved September 10, 2014 from http://ncleg.net/gascripts/BillLookUp/BillLookUp.pl?Session=2013&BillID=h998&submitButton=Go
[2] Databases, Tables & Calculators by Subject.* Bureau of Labor Statistics Data. N.p., n.d. Web. 12 Sept. 2014 from http://data.bls.gov/timeseries/SMS37000000000000001?data_tool=XGtable
[3] Giroux, G. (2014, March 18). Bloomberg by the Numbers: 9.2%. Retrieved August 28, 2014 from http://go.bloomberg.com/political-capital/2014-03-18/bloomberg-by-the-numbers-9-2-2/
[4] North Carolina General Assembly – House Bill 4 Information/History (2013-2014 Session). (n.d.). Retrieved September 10, 2014 from http://ncleg.net/gascripts/BillLookUp/BillLookUp.pl?Session=2013&BillID=h4&submitButton=Go

Senate 18 "Economy 2" Front

Probit Analysis of Persuasion Effects

	(1) Republican Support
2 Early	0.0962
	(0.32)
2 Late	0.124
	(0.304)
Six	0.108
	(0.258)
Eight	0.143
	(0.221)
Twelve	-0.118
	(0.22)
Democrat	-1.188***
	(0.292)
Republican	1.755***
	(0.31)
Female	-0.0216
	(0.086)
2 Early X Democrat	-0.0447
	(0.471)
2 Early X Republican	-0.720
	(0.453)
2 Late X Democrat	-0.167
	(0.477)
2 Late X Republican	-0.587
	(0.453)
Six X Democrat	-0.550
	(0.46)
Six X Republican	-0.674
	(0.405)
Eight X Democrat	-0.0965
	(0.351)
Eight X Republican	-0.610
	(0.355)
Twelve X Democrat	-0.157
	(0.377)
Twelve X Republican	-0.181
	(0.36)
Constant	-0.254
	(0.187)
χ^2	776.37***
N	1385

Standard errors in parentheses.

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Probit Analysis of GOTV Effects

	(1) 2014 Voter Turnout
Female	-0.0861***
	(0.017)
College	0.202***
	(0.017)
Past General Voting (0-4)	0.435***
	(0.008)
2 Early	-0.0624
	(0.089)
2 Late	0.0256
	(0.089)
Six	-0.000316
	(0.078)
Eight	0.0364
	(0.067)
Twelve	-0.0380
	(0.066)
Strong Republican	0.236***
	(0.071)
Weak Republican	-0.00823
	(0.085)
Weak Democrat	0.205*
	(0.081)
Strong Democrat	0.237**
	(0.074)
2 Early X Strong Republican	0.112
	(0.115)
2 Early X Weak Republican	-0.0317
	(0.139)
2 Early X Weak Democrat	0.0318
	(0.13)
2 Early X Strong Democrat	0.0297
	(0.117)
2 Late X Strong Republican	-0.00550
	(0.115)
2 Late X Weak Republican	-0.00287
	(0.137)
2 Late X Weak Democrat	-0.193
	(0.129)
2 Late X Strong Democrat	-0.0501
	(0.117)
Six X Strong Republican	0.0521
	(0.102)
Six X Weak Republican	0.0535
	(0.121)
Six X Weak Democrat	-0.120
	(0.114)
Six X Strong Democrat	-0.0368
	(0.104)
Eight X Strong Republican	-0.0550
	(0.087)
Eight X Weak Republican	-0.0323
	(0.104)
Eight X Weak Democrat	-0.177
	(0.098)
Eight X Strong Democrat	-0.0982
	(0.09)
Twelve X Strong Republican	0.0335
	(0.086)
Twelve X Weak Republican	0.0662
	(0.103)
Twelve X Weak Democrat	-0.104
	(0.097)
Twelve X Strong Democrat	-0.0383
	(0.089)
Constant	-0.915***
	(0.059)
χ^2	4370.88***
N	25152

Standard errors in parentheses.
* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Control Variable Effects

Likelihood of Voting Based on Past Behavior