

Leveraging Peer-to-Peer Connections to Increase Voter Participation in Local Elections

In local elections outcomes can turn on just hundreds of voters. Small-scale political entrepreneurs find it increasingly difficult to wage effective turnout campaigns relying on traditional, costly methods of outreach. Yet a growing literature on social pressure indicates that recasting voting as a socially motivated act increases the likelihood that voters participate. In this paper I present matching analyses of the impact of a new platform that relies on peer-to-peer voter outreach. I find that the intervention is responsible for significant increases in the likelihood that a voter casts a ballot in low information elections across different sorts of voters.

Keywords: Elections, Networks, Voting, Email, Local Politics

Acknowledgments: This research did not receive any specific grant from funding agencies in the public, commercial, or not-for-profit sectors.

Introduction

Political campaigns are costly and time intensive. For local office-seekers or ballot activists, an additional hurdle is the low-salience environment in which many races occur. Electorates are small, issues are oftentimes not of universal interest, and voting does not always occur on general election day. Enterprising political actors continue to develop and seek low-cost and easy to implement strategies to win elections. Campaign workers and scholars know that peer-to-peer canvassing is the most effective way to encourage voter participation, but it is also extraordinarily costly. This study assesses the combination of peer-to-peer connections and a new technology platform, VoterCircle, in the context of an all-mail ballot tax election. Despite some limitations, the results offer a glimpse into a new style of campaigning and mobilizing that will influence future local election outcomes.

The paper proceeds as follows. First, I discuss research on the effectiveness of various outreach strategies. Second, I introduce a set of hypotheses about a new platform that allows voter directed, rather than campaign directed outreach. Third, I perform an analysis of the efficacy of this voter outreach platform by applying a new estimation method appropriate for the non-experimental nature of the original deployment of the intervention. I conclude with a discussion of the results and implications.

Strategies of Voter Outreach and Mobilization

The mobilization of voters has been studied in a variety of election level, timing, geographic, and strategy contexts. As the discipline moves towards using administrative data and field experiments to study voter mobilization, methods for making more tenable causal statements about strategies that influence turnout are more robust than ever. Some techniques are more effective than others in getting citizens to participate in the electoral process, yet across all contexts there appear to be two places of consensus in the turn out literature.

First, mobilization and persuasion are hard. There are generally small differences in participation rates for those who received a mobilization intervention compared to those who do not, and there is little systemic evidence that campaigns change minds (Kalla and Broockman 2017). Second, more personal methods of contact are more effective than other methods such as phone outreach or mass mailings for both motivating voter participation and persuasion (Green, McGrath and Aronow 2013).

In the past 10 years, there has been an explosion of interest in measuring the effects of various non person-to-person voter turnout efforts at all levels of politics with an understanding that person-to-person contact is the most effective, but with the acknowledgment that many campaigns cannot rely on large canvassing efforts. Yet, the results from these studies come to less well-defined conclusions. Text, social media, and email are each relative newcomers to voter mobilization. Research provides different assessments on the effectiveness of strategies deployed across the newer mediums.¹

One of the earliest studies on email outreach involved 13 randomized control trials and over 230,000 subjects. This work reported that registration and turnout are generally unaffected by email encouragement (Nickerson 2007). Other large-scale research finds that emails from the local, elected, County Registrar of Voters is related to small positive effects on turnout (Malhotra, Michelson and Valenzuela 2012). But yet other research fails to find any effect and slight evidence that - sometimes - email reminders are related to slightly lower levels of turnout (Bennion and Nickerson 2011). The emails under study in this paper are unlike those that have been assessed before, as they do not originate with a political campaign or an elected official as most of the large-scale previous experiments do. This paper analyzes the effects of receiving an

¹ Given the ubiquity of accessing email, texts, and social media through applications on smart phones, the boundaries that used to characterize these different media may be receding.

email from an already known email contact. There is reason to believe there may be different findings considering the different message source.

Social Connectedness as a Way to Encourage Political Action

There is a great deal of literature relating the habits of voters to other potential voters in their familial and social circles (Campbell, et al. 1960) (Beck, et al. 2002) (Andolina, et al. 2003) (Fieldhouse and Cutts 2012). Connectedness matters; individual efforts have larger social impacts depending on the sorts of social networks and social circles of individuals exist in (Christakis and Fowler 2009). The main finding in this line of research indicates that voting can be *contagious*; the more people in your network who vote, the more likely you are to vote. There are multiple mechanisms by which that contagion operates. Sometimes it is just that like-minded people are more likely to interact with each other and also hold similar political behavior patterns. On Facebook, there is evidence that people take cues about information seeking and voting from other individuals who are close to them (Bond, et al. 2012) (Jones, et al. 2017). There are also feedback loops between on the ground mobilization efforts and social network activity. Those who participate in more political activities post more online about political activities and those who post more are prone to do more real-life activism in the future (Vissers and Stoller 2014). A meta-data analysis of the state of social media in elections confirms the nature by which any individual event happens a part of a greater social media ecosystem and thus may have impacts that radiate out to influence others (Boulianne 2015). There is also evidence of more purposeful style of encouragement contagion. Emails using social pressure from a closely linked individual are effective at increasing the likelihood that someone joins a professional group (Druckman and Green 2013).

Social media ushered in an era of citizen-initiated campaigning, where people are free to recruit and persuade other voters via online media without the specific direction of a political

actor. The turning point for this sort of citizen engagement is widely understood to have begun in the 2008 U.S. Presidential election (Gibson 2015). This reality muddies the ability for scholars to isolate the impact of any given campaign effort, but it is a reality not likely to change and thus must be adopted into the background understanding of how voters interact within an election environment. Despite this never-ending endogenous system, there is still merit in trying to isolate and understand the impacts of different sorts of campaign strategies.

Practitioners of emerging voter outreach strategies would be well suited to run studies that uphold the standards of field experiments, but limitations on time, money, know-how, and willingness often prevail. There are alternate, non-random styles of intervention deployment that can increase our understanding of mobilization mechanisms but require special caution in estimation techniques. In this study I analyze a novel peer-to-peer email technology that leverages peer-to-peer connections to encourage people to vote for certain candidates or ballot initiatives. The deployment of this technology was not done as an experiment, yet there is information to be gleaned that follows from assessing theoretical expectations with appropriate strategies.

The platform studied here uses the existing contacts of a potential voter to link people for an encouraging email exchange. Both academic research and political operative folk knowledge indicate that more direct and personal attempts tend to do better to increase political participation than anonymous or distant attempts. While election specific research is lighter on the efficacy of personal email, related research supports the notion that sender-receiver connectedness enhances the signal of a message. Thus, I hypothesize that peer-to-peer email results in greater voter participation for those who receive such a communication than similarly situated individuals who do not receive such a communication.

Peer-to-Peer Connections via VoterCircle: How it Works

For local, and generally low-salience elections, hiring or recruiting volunteer door-to-door canvassers is prohibitively costly – yet seasoned campaign staff know the more personal the appeal, the more effective it is. VoterCircle is packaged as a compromise solution. To use VoterCircle a campaign first obtains the voter file of all eligible voters for a given election.² Second, from the list of eligible voters, a campaign identifies known supporters and recruits them to be “seeds” for further network outreach.³ Once seeds agree to use VoterCircle the platform combs through seed’s email contact list, which is populated by friends, neighbors, coworkers, and anyone whose email address has been contacted previously by an individual. Some contacts may be eligible to vote in the election of interest, and some may not. In order to assess eligibility the platform cross-references the voter file to identify eligible voters – based on a matching algorithm of first and last names, as well as email, phone numbers, and addresses.⁴ There is no element of selectivity or microtargeting in this stage, any voter identified as eligible is flagged as a match and provided to the seed as someone eligible to vote in the election of interest. After the platform matches, seeds are given a list of eligible contacts and on their own schedule can send emails to their eligible contacts urging support for a candidate or initiative. The VoterCircle platform supplies contact seeds with a template to craft outreach emails. The template simply indicates the day and subject of the election and includes an appeal from their

2 The amount and quality of data available in a voter file vary from state to state – in this instance the voter file contains information on a voter’s name, the precinct they are registered to vote in, gender, age, year first registered, and the voting history on whether they participated or not in the past two elections.

3 It is important to note that the author has no access to the identity of the seeds in the deployment under study here. Seeds are typically those who are personally or professionally close to a candidate or initiative organizers.

4 A video summarizing the process and platform compatibility is available at: <https://votercircle.com/>.

peer to vote in a certain way. Seeds are free to modify messages.⁵ VoterCircle emails look more personal than mass campaign emails because it is individualized from each peer sender to each peer recipient.

Here is the basic template for the initiative used in this study:

Subject: Asking for your YES vote on A

Message:

Friend,

By now you should have received a special mail-in election ballot in your mail regarding an important measure for our school district - Measures A. I am sending this email to ask you to join me and the many other community and school leaders listed below in voting YES on this very important. Every vote is important, so please return your ballot today!

For more information, you can visit <http://www.supportmenloparkschools.org> or feel free to reach out to me directly anytime.

Thanks for doing this!

Fellow Endorsements for YES on Measures A

[List of already pledged community leaders removed here for space]

2016 Test Case Specifics

⁵ Like information on seed names, the author does not have access to the individual modifications some seeds may choose to use. The vendor, VoterCircle knows the senders and recipients of each email sent using the platform – but the individual modifications to the text of the message are not accessible to the platform or the individual campaign. As is the case with other directed outreach, the expectation is that senders will add in personalizing information to make a more effective appeal. Examples would be using first names in the address, adding a life update or inquiry to the recipient, signing off in a personalized way, etc. This information can be communicated to contact seeds when they are originally recruited.

The Board of Education for the Menlo Park City Elementary School District in California conducted an election by mail ballot on a single ballot measure (Measure A) regarding a parcel tax referred to as the “Excellence in Education Tax” in 2016. The text of the ballot initiative proposed to voters was:

To maintain our community commitment to existing small class sizes, high quality teachers and comprehensive academic programs, and renew the expiring education special tax, shall the Menlo Park City Elementary School District continue to levy at the current annual rate of \$201.38 per parcel, spent only on teachers and education programs, beginning July 1, 2016, adjusted annually for inflation, with an exemption available to individuals 65 and over?

The initiative received typical media coverage and had typical turnout for elections of this sort. There were 16,372 eligible, registered voters. The 2010 Census indicates that the total population of Menlo Park is 32,026, 70% white, 10% Asian, 9% Pacific Islander, 5% African American, and the remaining identify as other races; 18% identify as Hispanic or Latino. Of the 12,347 households in the city, 33% have children under 18. The top employer is Facebook, 70% of the population has at least a bachelor’s degree and the median family yearly income is \$121,000. This community is particularly attuned to email and electronic communication, a point I return to in the conclusion.

Data and Analyses

The data consist of (1) the voter file of all registered voters eligible to vote on Measure A⁶, (2) a database of every voter who received a VoterCircle email from a peer urging them to vote for the measure⁷, and (3) administrative data recording whether or not an eligible voter voted on

6 This data was graciously provided to the author by Political Data <http://politicaldata.com/>.

7 This data was graciously provided to the author by VoterCircle <http://votercircle.com/>.

Measure A.⁸ Four hundred twenty-three eligible voters received VoterCircle emails sent from seeds during the campaign.⁹

Explanatory and Control Variables

In each of the following analyses, the variable of interest is whether someone received a VoterCircle email. I use additional information to model the likelihood that someone receives a VoterCircle email in the first place, and then to create an appropriate “control” group by matching those who received an email to the most feasible set of counterfactual individuals who did not.

Generally speaking, women tend to turnout more than men (Levine and Lopez 2002) and women tend to be more active in supportive social networks than men (Antonucci and Akiyama 1987) (Cornwell 2011). Therefore, voter gender is important to include in the analyses of who is more likely to vote and who is more likely to receive a VoterCircle email initially. Age is also related to the likelihood that someone votes as well as the size and types of their social circle (Morgan 1988) (Strate, et al. 1989). Thus, controls of voter age are included.

In this election, there were 10 parties in the voter file, but for the analysis I use only members of the top 5 largest parties.¹⁰ As parties serve as organizing forces for campaigns, I

⁸ There is not individual level data on whether or not a person voted for or against Measure A.

⁹ For the following analyses, the total population is not always listed at 16,372 as there is some data missingness in some fields that reduces the overall number of individual data points used.

¹⁰ The parties eliminated either have too few registered voters to make any inferences from, or they perfectly predict turnout and the reception of a VoterCircle email and are therefore dropped in the regression analysis. The eliminated parties are Libertarian (101 registered voters) Natural Law (1 registered voter), Peace and Freedom (13 registered voters), Reform (10 registered voters) and YY (6 registered voters).

control for the party identification of voters despite the fact that this ballot contest was not defined with great partisan intensity. The voter file has a field indicating the date a person registered to vote as well as if an individual voted in the previous election. Based on research that indicates voting is habitual (Gerber, Green and Shachar 2003) (Aldrich, Montgomery and Wood 2011) (Denny and Doyle 2009) and the idea that people who have a history of voting are more likely to be in contact with campaigns than newer or less seasoned voters, I use a continuous measure of years registered and an indicator of whether or not someone cast a ballot in the 2014 election in each of the subsequent analyses.

Analysis strategy

The roll out of VoterCircle in this campaign was not a true field experiment as VoterCircle emails were not sent at random. Despite this limitation, there are methods to analyze the efficacy of the platform using approaches typically employed in field experimental research. Like a field experiment, the intervention took place unobtrusively in the context of an actual campaign, the participants were actual voters, and the results are individualized, publicly recorded, verifiable, and of real-world importance.

The largest concern from an inferential standpoint is that the “treatment” of getting a VoterCircle email may be systematically related to citizen attributes that are also related to the likelihood that someone votes in the first place. In this election just 3% of the total electorate were exposed to the VoterCircle intervention. Comparing this slice of the electorate to other potential voters reveals that those who received a VoterCircle email are particular in some ways and unremarkable in others.

With the recognition of these limitations, I attempt to approximate a randomized field experiment by matching the treatment sample to the most reasonable set of counterfactual individuals who did not receive a VoterCircle email. I use the latest strategy in matching

developed and described in (King, Lucas and Nielson 2016). While the VoterCircle deployment lacks randomization, this estimation strategy provides the soundest basis for making inferences as to the effectiveness of VoterCircle in a way that best approximates experimental standards given the limitations of the reality of previous deployment.

Results

I consider first the results describing the differences between people who received and VoterCircle email to those who did not. The following analyses assesses the efficacy of the VoterCircle platform in encouraging voter participation. Table 1 provides summary statistics comparing eligible voters who received a VoterCircle email to those who did not.

[TABLE 1 HERE]

Title: Average Characteristics of Eligible Measure A voters by VoterCircle Email Reception

As Table 1 makes clear, the deployment of VoterCircle was non-random. Women, slightly older voters, and voters who had participated in an election recently were more likely to receive a VoterCircle email. The gender split in the population of non-VoterCircle exposed voters is 52% female, 48% male. Of those who received a VoterCircle email, the split is 60% female, 40% male. The average age of a VoterCircle recipient is 48, while the non-VoterCircle population is 50 on average. The distribution of age is more telling than the average age, there is a dense, age-based clustering of those who received a VoterCircle email centered around age 48 versus the broader distribution in the general voting population. This finding is likely because people tend to group with and be contacts with others who are closer to their own age.

The largest difference among recipients and non-recipients is past voting. Those who received VoterCircle emails are significantly more likely to have voted in the 2014 general election than those who did not; 74% of VoterCircle recipients voted in 2014, while just 55% of

non-VoterCircle recipients did. The mechanism of VoterCircle that relies on politically supportive and active seed peers to seek out other peers at least partially explains why we observe more politically active people in the list of those who get a VoterCircle email.

Of the similarities across recipient status, there is not a strong relationship between the party an eligible voter is registered with and the likelihood that he or she receives a VoterCircle email. Overall, the distributions of partisanship within those who did and did not receive the treatment are very similar and a Fisher's exact test of association retains ($p = 0.60$) the null hypothesis that they are drawn from similar populations. There is not a significant relationship between the time a citizen has been registered to vote and the likelihood he or she was exposed to a VoterCircle email.

VoterCircle Turnout Effects: Matching

A simple probit regression that considers the relative impacts of receiving a VoterCircle email as well as a host of other controls on the outcome of voting, unsurprisingly results in a large and positive coefficient relating the fact that someone received a VoterCircle email and cast a ballot.¹¹ But getting a VoterCircle email in the first place is highly correlated with other explanatory variables related to voting so the common and necessary assumptions underlying unbiased regressions estimation and inference are not met and do not provide an adequate assessment of causal effects.

To obtain better effectiveness estimates with less bias I use a technique that relies on a matching *frontier* metric to optimize balance between the groups of eligible voters who received an email and those who did not, as well as the sample size from which inferences are to be made

¹¹ Specifically voting is a 1 if a person cast a ballot and a 0 otherwise; independent variables include an indicator on whether or not someone got a VoterCircle email, an indicator denoting if someone voted in 2014, a sex indicator, voter age, the number of years registered, party indicators, and precinct indicators.

(King, Lucas and Nielson 2016). Every treated voter (received a VoterCircle email) is matched with the closest non-treated voter (did not receive a VoterCircle email) by minimizing the Mahalanobis distance between a treated and non-treated voter on all explanatory variables. The set of eligible voters on which inferences are to be made is smaller than the universe of all eligible voters. This reduces the number of observations, but means that the impact of a VoterCircle email is estimated in the cleanest, least biased way by comparing only the most similar potential voters.

A potentially straightforward way to find counterfactual individuals for the 441 recipients of VoterCircle emails might be to find 441 other potential voters (1:1 matching) among the non-recipients or to match the original 441 to some weighted other number (1:N matching) of potential voters. Yet, each of these strategies assume and force a match from the non-treated population, which may sometimes assign an ill-matched counterfactual to a treated observation for the sake of keeping the number of observations “high”. That strategy may in turn lead to biased estimation. The technique I use here allows both treated and non-treated observations to be discarded and weighted. That is, there is an effort to optimize balance between both the sample size and likeness among the matches used for inferential analyses.

Because these inferences are based observational data the estimator of interest is the feasible sample average treatment effect on the treated individuals (FSATT). Figure 1 presents the estimates of FSATT over various levels of treated voters to successively smaller numbers, but more closely matched non-treated voters. At the left most part of the figure, the difference attributed to getting a VoterCircle and voting on Measure A is .38. This is the uncontrolled difference in vote rates for those who got a VoterCircle email (72% cast a Measure A ballot) and those who did not get an email (34% cast a Measure A ballot). As the matching further reduces the distances between treated and non-treated voters by pruning the sample of non-treated voters

to be more similar to treated voters this effect is reduced, but still robust. I end up with an optimal comparison group of 272 treated VoterCircle recipients and 574 non-treated potential voters.

[FIGURE 1 HERE]

Figure 2 shows how the sample changes as controls are pruned to better approximate the treated population. As is expected based on the breakdowns in table 1, partisanship does not appreciably change, but whether someone voted in 2014 and is a woman or not, changes such that those with greater past electoral participation and more women are in the smaller, but more evenly balanced samples.

[FIGURE 2 HERE]

Given that the estimates of the treatment effect are most conservative as the matching sample is pruned, I use a pruned result of an optimally matched sample of 846 eligible voters to perform a further set of analyses. Table 2 presents the results of OLS regression and a probit regression relating the explanatory variable as well as the control variables to the outcome of casting a ballot on Measure A on the matched population.

[TABLE 2 HERE]

As table 2 shows, the intervention of receiving a VoterCircle email is significantly and positively related to voting on Measure A. The only other variable consistently related to a higher likelihood of voting is if someone voted in 2014 or not. To better understand VoterCircle influence across voter history differences, I split the matched sample into voters who voted in 2014 and those who did not. Then I perform simulations holding all other variables at their

respective means and varying only the reception of a VoterCircle email using CLARIFY (Tomz, Wittenberg and King 2003).¹²

While those who did not vote in 2014 have a lower baseline probability of voting in 2016 than those who did – both populations are more likely to vote upon reception of a VoterCircle email. For those who did not vote in 2014, the receipt of a VoterCircle email is associated with a 25% (s.e. 7%) increased likelihood in voting, and for those who did vote in 2014, the associated increased likelihood in voting is 30% (s.e. 3%). For all sorts of voters, receiving a VoterCircle email is related to an increase in the likelihood that someone casts a ballot, even under successively conservative estimation techniques.

Conclusions and Limitations

The results here offer a somewhat novel finding and a support for of a variety of other conclusions for different assessments of voter mobilization. These results suggest that is it not so much the medium, but the messenger that matters in driving turn out. In an increasingly technologically powered and impersonal world, the power of human connection still prevails.

In this case study, a real deployment of a new email voter mobilization tool was used in a low-salience, all-mail election in a well-educated, high earning, tech savvy population. To be sure, there are limitations and cautions to be taken when considering the strictly non-experimental deployment of this tool. However, given the assessment mechanism using only the most closely similar set of individuals for analyses, the results point to a positive increase in voter mobilization. The peculiarity about the population should actually bias against finding any sort of impact as these are the sorts of people who are more keenly aware about spam email and

¹² The simulation specification assumes a woman, democrat, of age 46, and registered for 8 years.

prone to dismiss electronic communications deemed unimportant. The fact that I find a positive, and significant result from receiving an email from a known contact is quite a powerful insight.

Campaigns are not in a position to ensure that voters have cultivated a habit of voting, but campaigns do have the ability to leverage social connections of supporters in order to encourage citizens to vote. For potential voters of all sorts, the intervention of a VoterCircle email is associated with an increased likelihood of participating in the electoral process. For voters who do not have an established voting habit, a VoterCircle email from a friend may make the key difference in taking someone from a non-voter to an active political participant.

By employing a cutting edge, yet straightforward matching procedure, the non-random deployment of VoterCircle in the 2016 Menlo Park election can be assessed in a way that fairly considers other influences on voting, while still offering a robust test of my hypothesis.

There are some important limitations of this study. The type of election used for Measure A was conducted as an all-mail ballot. This is standard for some local races in California and elsewhere but is still not typical. Oregon, Washington and Colorado have provisions to conduct all elections by mail, and 19 other states have some type of mail-in ballot option. With that in mind, the lessons on the efficacy of VoterCircle emails may not be the same for elections that require voters to go to the polls. This sort of local election ballot issue is unlike an election featuring partisan candidates. In elections for political candidates, voter mobilization is not just about candidates, but the greater party forces also serve to influence voter turnout. Yet for many ballot initiatives that lack clear support or opposition based on partisanship a VoterCircle email from a known contact may be uniquely effective as people have fewer political cues to rely on. Another reservation to keep in mind is the scale of the election and this sort of intervention. The power of this type of outreach strategy is probably most uniquely situated to local, low-

information, low-salience elections where the push from a friend to vote in a certain way might be the most useful input any individual voter receives for that sort of contest.

If greater voter turnout can be achieved with VoterCircle interventions, this tool may be of benefit for those wishing to see more citizen participation in the political process. But a word of caution is in order. Knowing the power of social networks and the effects of peer-to-peer information sharing means that those choosing to use this sort of tool must do everything possible to ensure the security of the platform use. It is not difficult to imagine a scenario in which a seed's contacts are hacked, then contacted with an email address that is nearly identical to email address of the hacked seed. This sort of debacle could lead to misleading or misattributed messages in order to undermine one side in a campaign. These sorts of implications are potentially far reaching and merit very close scrutiny and preparedness on behalf of campaigns.

For practitioners and campaigns the take away is quite clear – the reception of an email from a peer facilitated by the VoterCircle platform is associated with a significant increase in the likelihood that the recipient casts a ballot. For no demographic did a VoterCircle email lead to fewer votes, and instead it is quite impactful across many subgroups. In a reality where local elections turn on fewer voters, campaigns continue to rise in cost, and levels of voter trust erode, platforms such as VoterCircle may be a best bet. This method leverages peer-to-peer connections to inform voters, reduce costs for those running campaigns, and effectively increases voter participation.

Bibliography

Aldrich, John H., Jacob M. Montgomery, and Wendy Wood. 2011. "Turnout as a Habit."

Political Behavior 33 (4): 535-563.

- Andolina, Molly, Krista Jenkins, Cliff Zukin, and Scott Keeter. 2003. "Habits from home, lessons from school: influences on youth civic engagement." *Political Science and Politics* 36 (2): 275-280.
- Antonucci, Toni C., and Hiroko I. Akiyama. 1987. "Social Networks in Adult Life and a Preliminary Examination of the Convoy Model." *Journal of Gerontology* 42 (5): 519-527.
- Beck, Paul Allen, Russell J. Dalton, Steven Green, and Robert Huckfeldt. 2002. "The social calculus of voting." *American Political Science Review* 91 (1): 57-73.
- Bennion, Elizabeth A., and David W. Nickerson. 2011. "The cost of convenience: An experiment showing email outreach decreases voter registration." *Political Research Quarterly* 64: 858-869.
- Bond, Robert M., Christopher J. Fariss, Jason J. Jones, Adam D. Kramer, Cameron Marlow, Jaime E. Settle, and James H. Fowler. 2012. "A 61-million-person experiment in social influence and political mobilization." *Nature* 295-298.
- Boulianne, Shelley. 2015. "Social media use and participation: a meta-analysis of current research." *Communication and Information Technologies Section (ASA) Special Issue* 18 (5): 524-538.
- Campbell, Angus, Philip E. Converse, Warren E. Miller, and Donald E. Stokes. 1960. *The American Voter*. New York: Wiley.
- Christakis, Nicholas A., and James H. Fowler. 2009. *Connected: The Surprising Power of Our Social Networks and How They Shape Our Lives*. New York: Back Bay Books.
- Cornwell, Benjamin. 2011. "Independence Through Social Networks: Bridging Potential Among Older Women and Men." *J Gerontol B Psychol Sci Soc Sci* 66B (6): 782-794.
- Denny, Kevin, and Orla Doyle. 2009. "Does voting history matter? Analysing persistence in turnout." *American Journal of Political Science* 53 (1): 17-35.

- Druckman, James N., and Donald P. Green. 2013. "Mobilizing Group Membership : The Impact of Personalization and Social Pressure e-Mails." *SAGE Open* 1-6.
- Fieldhouse, Edward, and David Cutts. 2012. "The companion effect: household and local context and the turnout of young people." *The Journal of Politics* 74 (3): 856-869.
- Gerber, Alan S., Donald P. Green, and Ron Shachar. 2003. "Voting May Be Habit-Forming: Evidence from a Randomized Field Experiment." *American Journal of Political Science* 47 (3): 540-550.
- Gibson, Rachel K. 2015. "Party change, social media and the rise of 'citizen-initiated' campaigning." *Party Politics* 21 (2): 183-197.
- Green, Donald P., Mary C. McGrath, and Peter M. Aronow. 2013. "Field Experiments and the Study of Voter Turnout ." *Journal of Elections, Public Opinion and Parties* 23 (1): 27-48.
- Jones, Jason J., Robert M. Bond, Eytan Bakshy, Dean Eckles, and James H. Fowler. 2017. "Social influence and political mobilization: Further evidence from a randomized experiment in the 2012 U.S. presidential election." *PLoS ONE* 12 (4).
- Kalla, Joshua L., and David E. Broockman. 2017. "The Minimal Persuasice Effects of Campaign Contact on Candidate Choice in General Elections: Evidence from 45 Experiments." *17th Annual Northeast Political Methodology Meeting*.
- King, Gary, Christopher Lucas, and Richard A. Nielson. 2016. "The Balance-Sample Size Frontier in Matching Methods for Causal Inference." *American Journal of Political Science*.
- Levine, Peter, and Mark Hugo Lopez. 2002. "Youth Voter Turnout has Declined, by Any Measure." *Center for Information and Research on Civic Learning and Engagement* 1-11.

- Malhotra, Neil, Melissa R Michelson, and Ali Adam Valenzuela. 2012. "Emails from official sources can increase turnout." *Quarterly Journal of Political Science* 7: 321–332.
- Morgan, David L. 1988. "Age Differences in Social Network Participation." *The Journal of Gerontology* 43 (4): 129-137.
- Nickerson, David W. 2007. "Does email boost turnout?" *Quarterly Journal of Political Science*, 369-379.
- Strate, John M., Charles J. Parrish, Charles D. Elder, and Coit Ford. 1989. "Life span civic development and voting participation." *American Political Science Review* 83 (2): 443-464.
- Tomz, Michael, Jason Wittenberg, and Gary King. 2003. *CLARIFY: Software for Interpreting and Presenting Statistical Results Version 2.1. Cambridge, MA.*
- Vissers, Sara, and Dietlind Stoller. 2014. "Spill-Over Effects Between Facebook and On/Offline Political Participation? Evidence from a Two-Wave Panel Study." *Journal of Information Technology & Politics* 11 (3): 259-275.

Table 1:
Average Characteristics of Eligible Measure A voters by VoterCircle Email Reception

	VC Treatment	No VC	Difference
Average Age	48.5	50.5	2.0*
Women (%)	60%	52%	8.0*
Voted in 2014 (%)	74%	55%	19.0*
Years Registered	12	12	0
Party: Republicans (%)	25.4%	21.6%	3.8
Party: Democrats	46.0%	49.4%	3.4
Party: Democrat Socialists	26.5%	26.0%	0.5
Party: American Independents	1.1%	1.7%	0.6
Party: Green	0.7%	0.6%	0.1
N	441	15656	

Note: * p<0.05

Table 2: Estimates of VoterCircle Email impact on voting, matched sample only

	Dependent Variable: Voted 2016 Measure A		
	OLS		Probit
	(1)	(2)	(3)
VoterCircle Email	0.287*** (0.035)	0.292*** (0.034) 0.001 (0.003)	0.860*** (0.104) 0.003 (0.008)
Voted 2014		0.357*** (0.041)	1.013*** (0.008)
Female		0.028 (0.034)	0.083 (0.098)
Years Registered		-0.0003 (0.002)	-0.001 (0.007)
AI Party		-0.163 (0.327)	-0.465 (0.926)
Dem Party		0.001 (0.043)	0.004 (0.125)
Dem Soc. Party		0.02 (0.051)	0.067 (0.147)
Green Party		-0.687** (0.327)	-5.598 (96.481)
Republican Party		Omitted	Omitted
Constant	0.467*** (0.020)	0.120 (0.142)	1.109*** (0.422)
Observations	846	846	846
R2	0.073	0.158	
Adjusted R2	0.072	0.149	
Log Likelihood			-508.188

Note: *p<0.1; **p<0.05; ***p<0.01

Figure 1: FSATT estimates over various matched sample sizes

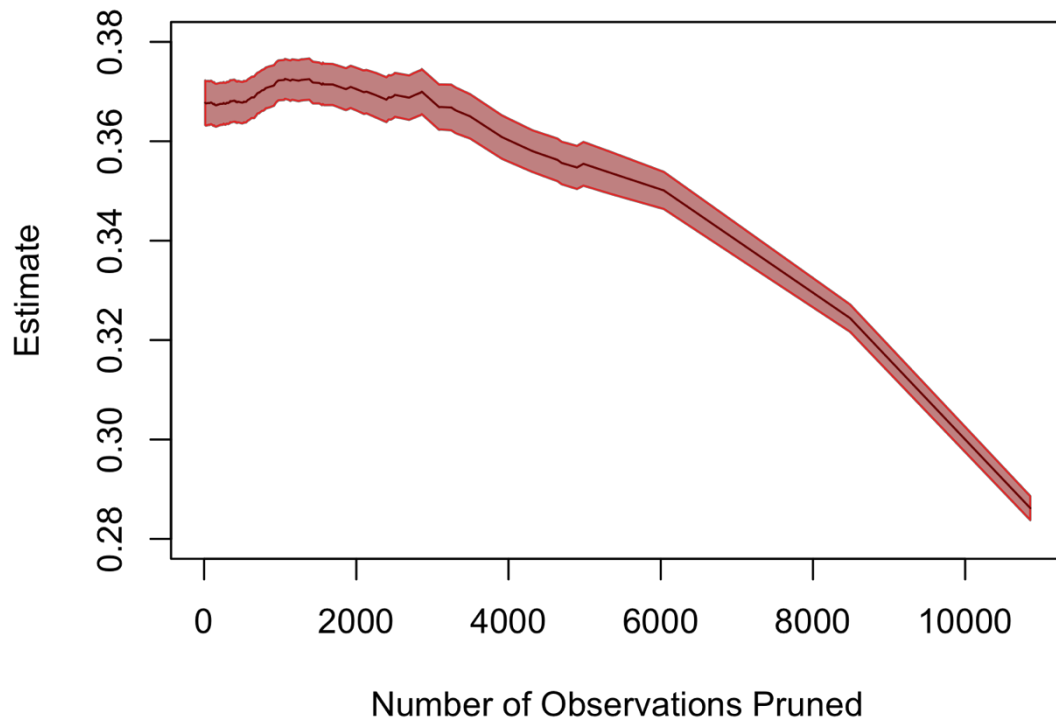



Figure 2: Scaled Means across Matched Sample Sizes



Additional Figures

Figure 1

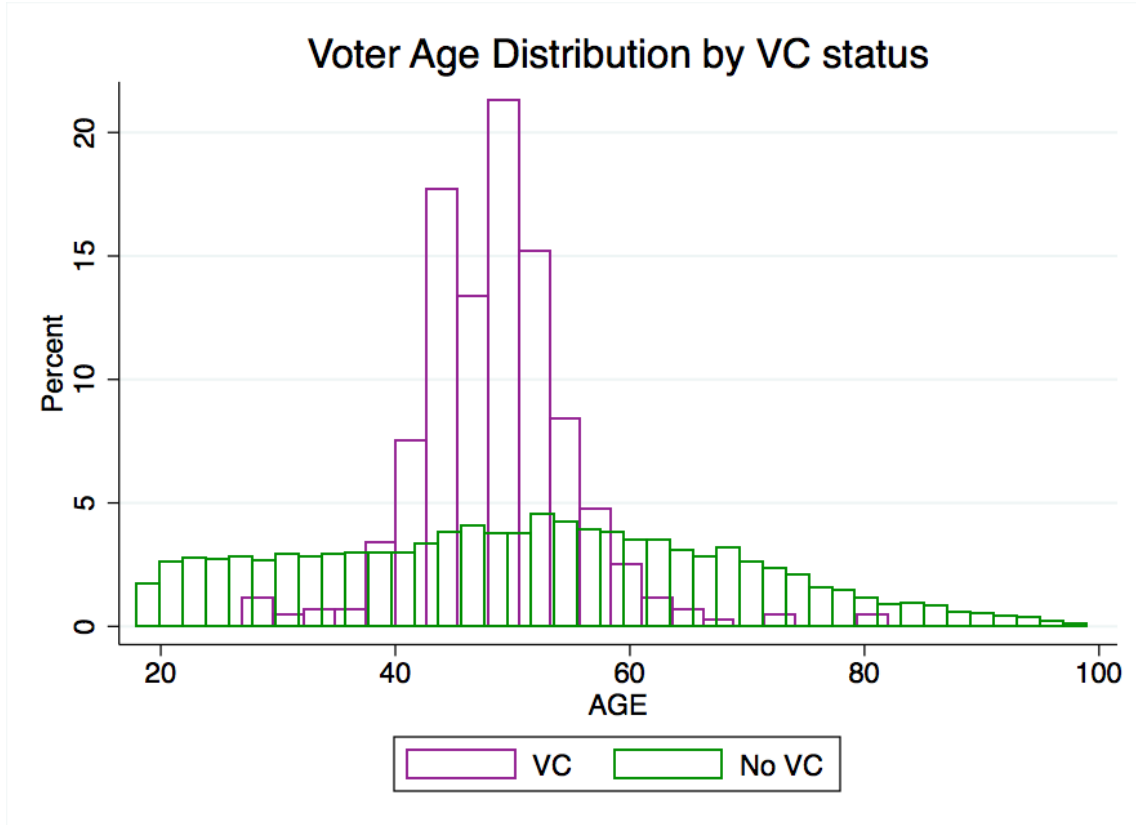


Figure 2: Turn Out Rates by VoterCircle Recipient Status

