


Does digital advertising affect vote choice? Evidence from a randomized field experiment

Research and Politics
 January-March 2022: 1–7
 © The Author(s) 2022
 Article reuse guidelines:
sagepub.com/journals-permissions
 DOI: 10.1177/20531680221076901
journals.sagepub.com/home/rap


Alexander Coppock¹ , Donald P. Green² and Ethan Porter³ 

Abstract

Despite the increasing sums devoted to online political advertising, our understanding of the persuasive effects of such advertising is limited. We report the results of a ZIP code level randomized field experiment conducted over Facebook and Instagram during the 2018 U.S. midterm elections in Florida. The ads, produced by a Democratic-leaning political action committee, were designed to spur Democratic vote share and were seen more than 1.1 million times with over 100,000 full views. This wide saturation notwithstanding, we find that these advertisements had very small estimated effects on Democratic vote share at the precinct level (−0.04 percentage points, SE: 0.85 points). Our results underline the challenges of political persuasion via digital advertisements, particularly in competitive electoral contexts.

Keywords

elections, advertising, social media, persuasion

While the study of political advertising has been largely focused on televised advertising, television is not the only means by which campaigns deliver messages. During the 2020 election, nearly a quarter of all political advertising spending was devoted to digital advertising (Media Project, 2021). At one point during the 2020 Democratic primaries, the leading candidates were spending more on Facebook ads than television ads (Goldmacher and Bui, 2019). Indeed, the amount spent on digital advertising appears to be increasing with each passing election cycle (Homonoff, 2020).

The amount spent on digital advertising raises an obvious question: Do social media advertisements affect vote choice? Campaigns—particularly winning ones—seem to think so. Brad Parscale, Donald Trump’s 2016 campaign manager, credited Facebook advertisements for Trump’s victory (Beckett, 2017). From a campaign perspective, the allure of social media ads is considerable. As they allow for more fine-grained targeting and do not require large start-up costs, Facebook ads are used by a broader set of campaigns than traditional ads and appear especially appealing to down-ballot candidates (Fowler et al., 2021).

Direct evidence attesting to the influence of digital advertisement on election outcomes, however, is hard to come by. A recent study conducted in Germany suggests that such ads can motivate vote choice in the intended direction (Hager, 2019), though the estimates in that study are not statistically distinguishable from zero at conventional levels. In the U.S., previous work has relied on randomized exposure to ads, with effects on candidate preferences measured via follow-up phone surveys ostensibly unrelated to the ads (Broockman and Green, 2014; Turitto et al., 2014). Other studies have examined the effects of geographically targeted Facebook advertisements on voter turnout in 2012 and 2013, obtaining weakly negative and insignificant point estimates (Collins et al., 2014).

¹Yale University, New Haven, CT, USA

²Columbia University, New York, NY, USA

³George Washington University, Washington, DC, USA

Corresponding author:

Ethan Porter, George Washington University, 805 21st St NW, Washington, DC 20052-0086, USA.

Email: evporter@gwu.edu



Creative Commons Non Commercial CC BY-NC: This article is distributed under the terms of the Creative Commons Attribution-NonCommercial 4.0 License (<https://creativecommons.org/licenses/by-nc/4.0/>) which permits non-commercial use, reproduction and distribution of the work without further permission provided the original work is attributed as specified on the SAGE and Open Access pages (<https://us.sagepub.com/en-us/nam/open-access-at-sage>).

In this paper, we add to this small body of existing evidence a study that evaluates the effects of persuasive ads deployed via Facebook and Instagram on vote choice in the United States. In the weeks immediately prior to the 2018 elections, in partnership with a Democratic-leaning political action committee, we randomized exposure to advertisements on Facebook and Instagram at the ZIP code level in Florida. While the partner organization's advertisements were meant to help elect Democrats, neither of the ads explicitly mentioned a specific candidate or campaign.

Instead, both ads focused on the virtues of the Democratic Party and the deficiencies of the Republican Party. In this way, they reflect broader trends in Facebook ads, which tend to be more partisan than television ads (Fowler et al., 2021). The tested ads were also quite similar to the non-candidate specific, issue-focused ads that have recently grown in prominence because of changes to campaign finance regulations (Persily et al., 2018). Together, the tested ads accumulated more than 1.1 million impressions. Following a pre-registered, randomized field experimental design similar in spirit to Arceneaux (2005) and Hager (2019), we measured the impact of the ads on Democratic vote share at the precinct level.

Our findings echo the emerging consensus that persuasive political messaging often has limited effects in general elections (Coppock et al., 2020; Kalla and Broockman, 2018) and that any such effects dissipate rapidly over time (Gerber et al., 2011; Hill et al., 2013). There are of course exceptions (e.g., Spenkuch et al., 2018), but ours is not one of them. Our estimate of the effect of the ads on vote share is substantively small and estimated with reasonably high precision. Our results underscore the challenge faced by groups trying to influence vote choice online. Political persuasion is hard, and social media advertisements do not necessarily make it any easier.

The paper is structured as follows. We begin by reviewing the available evidence on the effectiveness of political messaging and clarify our contribution. We then turn to describing the experimental design, as well as the context in which the experiment was administered. As we discuss, our design allows us to estimate a crucial quantity: the effect of the ads not on feelings or intentions, but vote choice itself. We present our results and contextualize our estimates through a meta-analysis of previous studies.

The challenge of political persuasion

Our goal is to understand the extent to which digital political advertising affects vote choice. On the one hand, campaigns and outside groups spend enormous sums of money to deliver advertisements to citizens as they use their digital devices. In 2018, the year this study was conducted, Facebook reported that \$400 million was spent on political advertising on its platform (Fowler et al., 2018). Across

platforms, the amount spent on digital political advertising continues to grow (Homonoff, 2020). On the other hand, the academic literature abounds with skepticism about the persuasive effects of political advertisements in general, regardless of format. The experimental evidence assembled by Kalla and Broockman (2018) shows that, across multiple channels of communication, the persuasive effects of political messaging are close to zero in the context of general elections. Kalla and Broockman's assessment is stark: "When we focus on the choices voters actually make on election day in a general election, we find that any early persuasion has decayed and that any persuasion near election day fails reliably." This summary conclusion is echoed by Coppock et al. (2020), who measure the (evidently small) effects of dozens of 2016 presidential television ads using survey experiments conducted over the course of the election.

All studies of the effects of advertising on vote choice face the fundamental data challenge that individual-level vote choice is not observable. Scholars have often addressed this difficulty by substituting survey measures for vote choice. For example, Gerber et al. (2011) examined the effects of television ads on favorability ratings of the advertising candidate and self-reported vote intention. Broockman and Green (2014) collaborated with political campaigns to evaluate the effectiveness of Facebook ads on attitudes toward the candidates. The treatments were deployed on Facebook, with outcomes measured by a polling company shortly thereafter. In two experiments, they find that ads for a Republican state legislative candidate have no discernible effects on voters' attitudes toward the candidate or vote intention (estimate: 1.6 points, SE: 1.4 points). Because these survey measures are likely correlated with actual vote choice, they come close to measuring the main object of interest—the effect of ads on vote choice—but do so only for the segment of voters who are willing to respond to surveys. Similarly, Turitto et al. (2014) use survey outcomes to measure the persuasive effects of a digital advertising campaign, cluster-assigned at the municipality level. Owing to the relatively small number of clusters, the point estimate (1.1 points, SE: 2.1 points) cannot be distinguished from zero.

As noted by Arceneaux (2005), since precincts are the lowest level of aggregation at which vote choice is observed, treatments must be assigned at the precinct level or higher to study the effects of treatments on vote choice measured at the ballot box. Arceneaux points out that some of the power loss due to cluster random assignment can be offset by adjusting estimates using the detailed pretreatment covariate information available about each precinct from historical election returns. In that study, voters in randomly selected precincts in Kansas City, Missouri were cluster-assigned to receive visits from door-to-door canvassers urging support for a ballot initiative. Hager (2019) used an analogous precinct-

randomized design to study the electoral effects of Christian Democratic Union (CDU) Facebook advertisements during Germany’s 2016 election. Treatment localities were assigned to either emotional or fact-based ads, while the control localities were assigned to no ad at all. Pooling over the treatment arms, the average treatment effect estimate on CDU vote share was 1.7 percentage points (SE: 1.2 points). Covariate adjustment does increase the precision of the estimate, but it diminishes in size and remains statistically insignificant at 0.9 points (SE: 0.6 points).

Like the [Broockman and Green \(2014\)](#), [Turitto et al. \(2014\)](#), and [Hager \(2019\)](#) studies, the present study also examines the effects of political advertisements delivered via social media. Other studies conducted on Facebook and Twitter explicitly leverage the network structure of these platforms to understand social interaction. For example, [Bond et al. \(2012\)](#) randomly assigned Facebook users to receive either informational or social messages regarding election turnout; [Eckles et al. \(2016\)](#) administered an encouragement design to uncover the role that peer effects play in Facebook behavior; [Haenschen \(2016\)](#) recruited Facebook users to participate in experiments on the capacity of social pressure on the platform to increase voter turnout; [Coppock et al. \(2016\)](#) randomly encouraged petition signers to tweet a petition endorsement to their followers. By contrast, our study leaves the social aspect of these platforms entirely to the side. We test the effect of ads delivered via Facebook and Instagram because many Americans spend large portions of their days on these platforms—indeed, the high volume of traffic on these sites is presumably why they are such attractive targets for political advertising in the first place.

Field experiment: Florida advertisements

We administered a randomized field experiment just prior to the November 2018 general elections to measure the effect of online pre-roll ads designed to encourage voting for Democratic candidates in four Florida congressional districts. We registered our pre-analysis plan at <https://egap.org/registration/5312> and include it in the [supplemental materials](#) for reference.

Our experimental units are the 210 ZIP codes associated with Florida congressional districts 15, 16, 26, and 27. ZIP codes are typically much larger than voting precincts, with the average ZIP code containing approximately 10–20 voting precincts. We associate precincts with ZIP codes using common membership on the voter file. This process creates many false negatives because voters’ addresses can be out of date with the precinct location on the file; we drop precinct-ZIP code pairs that account for less than 10% of voters putatively residing within a given ZIP code. The matching process is further complicated by the fact that precincts do not nest perfectly within ZIP codes.

In the main analysis, we exclude precincts that overlap ZIP code boundaries. We followed a blocked-and-clustered random assignment procedure. We grouped ZIP codes into matched trios according to their precinct size ([Moore, 2012](#)). Blocking on cluster size reduces the threat of bias when cluster sizes are uneven ([Imai et al., 2009](#)). ZIP codes could be assigned to one of three conditions: a control condition ($N = 104$) or one of two advertisement conditions ($N = 53$ each). After excluding precincts that span ZIP codes and matching to 2018 election returns, our dataset consists of 853 voting precincts fully nested within 154 ZIP codes. [Table 1](#) offers details on ZIP codes, voting precinct levels and treatment assignment.

Substantively, the two treatment ads were very similar. Both featured pro-gun control messages that criticized the Republican Party for its support of anti-gun control policies. The ads attempted to persuade viewers to vote for Democrats in the 2018 midterm elections because of their position on guns. In one of the ads, guns conspicuously dangle from teachers’ hips as Trump is heard to proclaim “the end of gun-free school zones.” In the other ad, a student texts his mother while his school is locked down during an active shooter incident. Neither of the ads reference specific campaigns or candidates; instead, they direct viewers’ attention to the partisan differences on gun control policy.

Because the advertisements are not candidate-specific, they are equally applicable to voters in different House districts. We deployed the ads across the four congressional races spanned by our study ZIP codes. Changes to campaign finance regulations have coincided with dramatic increases in the amount spent on non-candidate expenditures ([Persily et al., 2018](#)),

Table 1. Sample description at the ZIP code and voting precinct levels.

| | No exclusions | | Only nested precincts | | With outcome data | |
|-------------------|---------------|-----------|-----------------------|-----------|-------------------|-----------|
| | Precincts | ZIP codes | Precincts | ZIP codes | Precincts | ZIP codes |
| Control | 4683 | 104 | 524 | 84 | 420 | 76 |
| Treatment video 1 | 2246 | 53 | 285 | 43 | 220 | 41 |
| Treatment video 2 | 2306 | 53 | 287 | 42 | 213 | 37 |
| Total | 9235 | 210 | 1096 | 169 | 853 | 154 |

Table 2. Treatment advertisements.



| Ad 1: Armed Teachers | Ad 2: Lockdown |
|---|---|
| Bell rings | Ominous music |
| Teacher with handgun hands back papers | On-screen text messages: Mom. I'm in social studies. they're putting us on lock down |
| Whistle blows | What? Did they say why? |
| Gym teacher with handgun leads drills | Someone has a gun and they can't find him |
| Trump audio: "We're getting rid of gun-free on lock down zones" | A gun???? Are you ok? |
| Display text: Republicans want to arm up to 40% of teachers | Yeah but I'm so scared. I wish you were here |
| November 6th: Vote Democrat | I'm on my way |
| | Try to stay quiet |
| | [Text bubbles appear and disappear] |
| | Sweetie? |
| | Keep talking to me. I'll be there very soon. |
| | Voice-over: The NRA has given Republicans in Congress nearly \$10 Million. 95 percent of Senate Republicans have an "A" grade from the NRA. |

Stills were captured at 0:03 seconds into advertisements.

which makes an experimental test of the effect of ads that do not mention candidates especially interesting. Table 2 shows how the ads appeared at 0:03 s and provides the full transcripts.

The ads were deployed at the same time (starting 25 October 2018), and each of the ads was run until its budget of \$30,000 was exhausted (\$60,000 in total). Viewers could not skip the ad until they had watched at least 3 seconds. Table 3 shows summary data from the advertisement vendor on the number of three-second views, full views, and clicks generated by each advertisement. In total, the experimental intervention generated more than 1.1 million three-second exposures and over 100,000 full views.

Our primary outcome measure is the precinct two-party vote share for Democrats in the 2018 congressional elections, obtained from the Florida Secretary of State. Following our pre-registered analysis plan, we include the two-party vote share in the 2012, 2014, and 2016 general elections as covariates in order to increase the precision of our estimates. We cluster our standard errors at the unit of assignment, which is the ZIP code. Again following our analysis plan, we conduct one-tailed hypothesis tests using a randomization inference procedure under the sharp null hypothesis of no effect.

Table 4 displays effects on vote share. The first column shows the unadjusted difference-in-means estimate of the

Table 3. Treatment exposure information.

| | Clicks | Full views | Three-second views |
|------|--------|------------|--------------------|
| AD 1 | 24,832 | 70,266 | 611,732 |
| AD 2 | 17,185 | 47,647 | 516,746 |

Source: Ad vendor.


treatment effect of any video: a 2.1 percentage point increase in Democratic vote share, though this estimate is uncertain, as evidenced by its large standard error of 3.0 percentage points (one-tailed randomization inference p -value = 0.471). The inclusion of covariates (column 2) dramatically increases precision: the adjusted estimate is -0.04 percentage points with a standard error of just 0.85 percentage points. Model two is clearly more informative and suggests that the average treatment effect is very close to zero (p -value = 0.508). Columns 3 and 4 disaggregate the "any video" treatment so that each video's effect is estimated separately; the effects of the first are estimated to be mildly negative and the effects of the second to be mildly positive.

In the supplementary materials, we report a series of additional analyses for the interested reader. In particular, we report the estimated effects on vote margin

Table 4. Effects on vote share.

| | Model 1 | Model 2 | Model 3 | Model 4 |
|------------------------------|---------------------------------|----------------------------------|---------------------------------|----------------------------------|
| Any treatment video | 0.0021 (0.0225) | -0.0004 (0.0085) | | |
| Treatment video 1 | | | -0.0205 (0.0250) | -0.0016 (0.0109) |
| Treatment video 2 | | | 0.0255 (0.0291) | 0.0007 (0.0096) |
| Two party vote share (2016) | | 0.8482 ^a (0.0613) | | 0.8483 ^a (0.0616) |
| Missingness indicator (2016) | | 84.3343 ^a (6.0944) | | 84.3359 ^a (6.1268) |
| Two party vote share (2014) | | 0.1489 ^a (0.0511) | | 0.1484 ^a (0.0521) |
| Missingness indicator (2014) | | 14.8150 ^a (5.0869) | | 14.7661 ^a (5.1776) |
| Two party vote share (2012) | | 0.0531 (0.0307) | | 0.0525 (0.0315) |
| Missingness indicator (2012) | | 5.3026 (3.0541) | | 5.2424 (3.1297) |
| Intercept | 0.4937 ^a (0.0157) | 0.0135 (0.0122) | 0.4937 ^a (0.0157) | 0.0140 (0.0128) |
| R ² | 0.0000 | 0.8356 | 0.0100 | 0.8357 |
| Num. Obs | 853 | 853 | 853 | 853 |
| N Clusters | 154 | 154 | 154 | 154 |

^ap < 0.05. CR2 Cluster-robust standard errors are in parentheses.

**Figure 1.** Bayesian integration of the research literature on the effects of digital ads on vote choice.

and voter turnout. These estimates are substantively small and statistically nonsignificant. We also include an alternative estimation approach in which aggregate precinct-level outcomes to the ZIP code level, then compare treated ZIP codes to untreated ZIP codes. This approach has the advantage that we can include all 210 randomized ZIP codes in the analysis, but relies on a proportional allocation rule for handling precincts that span ZIP code boundaries. The substantive results in terms of magnitude and significance on vote share, vote margin, and turnout are very similar to those reported here in the main text.

We close our empirical section with a formal integration of our findings with the existing research literature on the effects of digital advertising on vote choice. To do so, we adopt a Bayesian framework in which diffuse priors are updated in the wake of each experiment. For ease of exposition, we assume that prior beliefs about the size of the average treatment effect are distributed normally, and we update these priors assuming that each experiment's results have a normal sampling distribution. The normality assumption allows us to apply Bayes' Rule by weighting the priors and each experimental result by the inverse of its squared standard error.

We begin with a diffuse prior centered on zero with a standard deviation of 5 percentage points. This prior distribution is depicted in the leftmost graph of Figure 1 and represents a considerable amount of prior uncertainty about the plausible effects of digital advertisements on vote choice. The three field experiments to date on this topic are Broockman and Green (2014), Turitto et al. (2014), and Hager (2019). Figure 1 shows the evolution of posteriors since 2014, as each study updates the priors formed by the studies that preceded it. After these first three studies, the posterior (panel 4) has a mean of 1.0 percentage points with a standard deviation of 0.5 percentage points.

How does the present study build upon the previous literature? Panel 4 shows the state of the literature before our study, while panel 5 shows how the posterior distribution looks after our study is included. Contributing an estimate centered precisely on zero, the Florida study shrinks the posterior to a mean of 0.7 points with a standard deviation of 0.4 points. Despite the fact that each of the four studies returned statistically insignificant results, the accumulation of evidence over time yields a relatively sharp picture of small positive effects.

In the appendix, we conduct a design diagnosis (Blair et al., 2019) that shows how, on its own, our experiment is relatively underpowered for the 1.0 percentage point average effect implied by the previous literature (power: 21%). However, if we conceive of the design as first obtaining an experimental estimate, then combining it with prior evidence to produce a Bayesian posterior, the power is much stronger (89%). As this exercise demonstrates, in research settings with small effects, individual experiments may be underpowered, but placing them within a cumulative research program can still advance knowledge incrementally.

Discussion

To our knowledge, the present study is among the first to test the effects of digital advertising on precinct-level vote choice in the U.S. Given the growing enthusiasm for digital advertising on large social media platforms, the fact that the digital ads we tested had no apparent effect on Democratic vote share is bracing. In light of prior experimental evidence, it appears that the average effect of digital advertising is weakly positive. We consider four possible explanations that may guide future research.

First is the possibility that digital advertising generally has little effect on voting behavior. Although the number of publicly accessible studies remains small, it is noteworthy that neither studies of vote choice (Broockman and Green, 2014; Turitto et al., 2014; Hager, 2019) nor studies of voter turnout (Collins et al., 2014) lend clear support to the thesis that digital advertising persuades or motivates. That said, if one were to take the mean of our final Bayesian posterior (0.7 percentage points) at face value, it would imply a cost per vote of \$8.68, a figure that compares favorably to other

campaign tactics (Green and Gerber, 2019).¹ Bear in mind, however, that the point estimate from our single study is very slightly negative, which would imply an infinite cost per vote.

Second, it may be that digital ads work when voters pay attention to them, but voters disregard them en route to the content that attracts them to social media sites. It is telling that, for every voter who watched our ads all the way through, roughly 10 viewers skipped our ads after 3 seconds. This brings us to the third explanation, which is that our ads were deficient. Although one can never rule out the possibility that others ads would have performed better, it should be noted that the advertising campaign was one that was actually deployed by an organization seeking to persuade the electorate; the ads were very much the kind of issue-based messaging that has grown in prominence in the wake of *Citizens United* and other Court decisions. However, we concede that future work should investigate whether the failure to directly mention and endorse candidates undercuts the influence of political advertising.

Finally, it may be that advertising's influence on vote choice is attenuated during the final days of a closely contested general election. In part, as Kalla and Broockman (2018) contend, persuasion is harder during a general election, when party labels overwhelm others considerations. Weak effects may also reflect the sheer volume of competing campaign messages, drowning out the experimental ad's message. Of course, this explanation raises the question of why digital advertising is relied on so heavily during the waning days of a general election, when the marginal returns may be attenuated and the marginal costs of advertising are at their maximum.

Each of these explanations lends itself to further research along the lines of our experiment. To ascertain whether digital advertising is capable of swaying votes, randomized trials involving a broad range of advertising content and volume are necessary. Of particular interest in the context of American campaign finance regulations are head-to-head comparisons between candidate-focused ads and issue-focused ads. To assess whether advertising effectiveness is contingent on context, it would be instructive to compare the effectiveness of presidential ads deployed in battleground and non-battleground states, perhaps leveraging adjacent media markets that attract markedly different numbers of ads.

Acknowledgments

We thank Frank Chi, Jesse Ferguson, Conor Gaughan and Tara McGowan for facilitation.

Declaration of conflicting interests

The author(s) declared no potential conflicts of interest with respect to the research, authorship, and/or publication of this article.

Funding

The author(s) received no financial support for the research, authorship, and/or publication of this article.

Ethical approval

This research was reviewed and approved by the Institutional Review Board of George Washington University (IRB#NCR202186).

Supplemental material

Supplemental material for this article is available online.

ORCID iDs

Alexander Coppock  <https://orcid.org/0000-0002-5733-2386>

Ethan Porter  <https://orcid.org/0000-0002-9530-6955>

Note

1. This calculation is based on delivering advertisements to 822,783 voters living in targeted ZIP codes, at a total cost of \$100,000, inclusive of production and distribution costs. $100,000 / (822783 * 0.007 * 2) = 8.68$.

References

- Arceneaux K (2005) Using cluster randomized field experiments to study voting behavior. *The ANNALS of the American Academy of Political and Social Science* 601(1): 169–179.
- Beckett L (2017) *Trump Digital Director Says Facebook Helped Win the White House*. London, UK: The Guardian. Available at: <https://www.theguardian.com/technology/2017/oct/08/trump-digital-director-brad-parscale-facebook-advertising>.
- Blair G, Cooper J, Coppock A, et al. (2019) Declaring and diagnosing research designs. *American Political Science Review* 113(3): 838–859.
- Bond RM, Fariss CJ, JonesJonesKramer JJ, et al. (2012) Cameron Marlow, Jaime E. Settle and James H. Fowler A 61-million-person experiment in social influence and political mobilization. *Nature* 489: 295–298.
- Broockman DE and Green DP (2014) Do online advertisements increase political candidates' name recognition or favorability? Evidence from randomized field experiments. *Political Behavior* 36(2): 263–289.
- Collins K, Kalla J and Keane L (2014) *Youth Voter Mobilization through Online Advertising: Evidence from Two GOTV Field Experiments*. Paper Presented at the Annual Meeting of The American Political Science Association. Washington, DC: The American Political Science Association. Unpublished Manuscript.
- Coppock A, Guess A and Ternovski, J (2016) When Treatments are Tweets: A Network Mobilization Experiment over Twitter. *Political Behavior* 38(1): 105–128.
- Coppock A, Hill SJ and Lynn V (2020) Persuasive Effects of Presidential Campaign Advertising: Results of 53 Real-Time Experiments in 2016. *Science Advances*.
- Eckles D, Kizilcec RF and Bakshy E (2016) Estimating peer effects in networks with peer encouragement designs. *Proceedings of the National Academy of Sciences* 113(27): 7316–7322.
- Fowler EF, Franz M and Ridout T (2018) *The Big Lessons of Political Advertising in 2018*. Available at: <https://theconversation.com/the-big-lessons-of-political-advertising-in-2018-107673>
- Fowler EF, FranzMartin MM, Gregory J, et al. (2021) Political Advertising Online and Offline. *American Political Science Review* 115(1): 130–149.
- Gerber AS, Gimpel JG, Green DP, et al. (2011) How Large and Long-lasting Are the Persuasive Effects of Televised Campaign Ads? Results from a Randomized Field Experiment. *American Political Science Review* 105(1): 135–150.
- Goldmacher S and Bui Q (2019) *Facebook Political Ads: What the 2020 Candi- Dates' Campaign Spending Reveals*. New York, NY: The New York Times. Available at: <https://www.nytimes.com/interactive/2019/10/14/us/politics/democrats-political-facebook-ads.html>.
- Green DP and Gerber A (2019) *Get Out the Vote: How to Increase Voter Turnout*. New York, NY: Brookings Institution Press.
- Haenschen K (2016) Social pressure on social media: using facebook status updates to increase voter turnout. *Journal of Communication* 66(4): 542–563.
- Hager A (2019) Do online ads influence vote choice? *Political Communication* 36(3): 376–393.
- Hill SJ, Lo J, Vavreck L, et al. (2013) How quickly we forget: the duration of persuasion effects from mass communication. *Political Communication* 30(4): 521–547.
- Homonoff H (2020) *2020 Political Ad Spending Exploded: Did it Work?* Available at: <https://www.forbes.com/sites/howardhomonoff/2020/12/08/2020-political-ad-spending-exploded-did-it-work/?sh=4dd5324d3ce0>.
- Imai K, King G and Nall C (2009) The essential role of pair matching in cluster- randomized experiments, with application to the Mexican universal health insurance evaluation. *Statistical Science* 24(1): 29–53.
- Kalla JL and Broockman DE (2018) The minimal persuasive effects of campaign contact in general elections: evidence from 49 field experiments. *American Political Science Review* 112(1): 148–166.
- Media Project Wesleyan (2021) *Political Ads in 2020: Fast and Furious*. Wesleyan University. Available at: <https://mediaproject.wesleyan.edu/2020-summary-032321/>.
- Moore RT (2012) Multivariate continuous blocking to improve political science experiments. *Political Analysis* 20(4): 460479.
- Persily N, Bauer RF and Ginsberg BL (2018) *Campaign Finance in the United States: Assessing an Era of Fundamental Change*. Bipartisan Policy Center.
- Spenkuch JL, Toniatti D and Toniatti D (2018) Political Advertising and Election Results*. *The Quarterly Journal of Economics* 133(4): 1981–2036.
- Turitto C, Green DP, Brian S, et al. (2014) *Testing the Persuasive Effects of Digital Media: A Cluster Randomized Field Experiment*. Available at: https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3537287. Unpublished Manuscript