Political Advertising Online and Offline*

Erika Franklin Fowler†, Michael M. Franz‡, Gregory J. Martin§,
Zachary Peskowitz¶, and Travis N. Ridout‖

Abstract

Despite the rapid growth of online political advertising, the vast majority of scholarship on political advertising relies exclusively on evidence from candidates' television advertisements. The relatively low cost of creating and deploying online advertisements and the ability to target online advertisements more precisely may broaden the set of candidates who advertise and allow candidates to craft messages to more narrow audiences than on television. Drawing on data from the newly-released Facebook Ad Library API and television data from the Wesleyan Media Project, we find that a much broader set of candidates advertise on Facebook than television, particularly in down-ballot races. We then examine within-candidate variation in the strategic use and content of advertising on television relative to Facebook for all federal, gubernatorial, and state legislative candidates in the 2018 election. Among candidates who use both advertising media, Facebook advertising occurs earlier in the campaign, is less negative, less issue focused, and more partisan than television advertising.

*Except where noted in the text, analyses presented were preregistered (`https://osf.io/3px5b`) prior to the release of the Facebook ad library. The Wesleyan Media Project acknowledges funding from the John S. and James L. Knight Foundation and Wesleyan University. We are grateful to Laura Baum, Dolly Haddad, Colleen Bogucki, Mason Jiang and the numerous undergraduates across our institutions for their efforts on this project. We thank Amanda Wintersieck, Devra Moehler, and seminar participants at APSA, the Princeton CSDP American Politics seminar, the University of Maryland, and the Wesleyan Media Project Post-Election Conference for comments on previous versions.

†Associate Professor of Government, Wesleyan University

‡Professor of Government and Legal Studies, Bowdoin College

§Assistant Professor of Political Economy, Stanford Graduate School of Business

¶Associate Professor of Political Science, Emory University

‖Thomas S. Foley Distinguished Professor of Government and Public Policy, Washington State University

How does the medium of political communication affect the message, if at all? A glance at the landscape of US political media suggests some connection between the two, with right-wing outlets dominant on talk radio and cable news, and successful new digital-native outlets generally leaning left. In the comparative context, campaigns in democracies where broadcast media are more centralized and public-owned are more programmatic and party-centered than in those with more fragmented viewer markets (Plasser and Plasser 2002). Of course, these are pure correlations, and it is entirely possible that these associations between medium and content simply reflect the demographic profile of the audience,[1] or common consequences of varying political cultures.

Nonetheless, the dramatic technological changes experienced over the past 15 years have real potential to shift the strategic landscape of campaign communication, and thereby alter the content of campaign messaging that voters receive. In particular, the mass adoption of the Internet, smartphones, and social media have moved the technological frontier of mass communication in two strategically important ways. First, social media platforms substantially lower the cost of advertising,[2] expanding the set of candidates for whom advertising - and thus the potential to reach voters and seriously contest an election - is a real possibility. Second, and perhaps even more consequential, social media platforms offer vastly more precise targeting capabilities than legacy broadcast media. This feature of social media could allow campaigns to strategically tailor messages to narrowly-defined audiences, a capability with the potential to undermine democratic accountability.[3]

[1] Or perhaps some deeper psychological connection between preferences for medium and preferences for political ideology (Young 2019).

[2] The low cost to post ads on social media is not without some complicating factors. For example, some media coverage of the 2020 Democratic presidential primary noted that the competition among over 20 candidates for ad space on Facebook, in part driven by the need to meet unique donor thresholds to participate in early debates, meant that prices from Facebook were much higher than what many campaigns expected to pay. Those costs often meant that campaigns were spending more on social media than what those efforts were raising in online donations. Still, the price relative to TV remains much lower. See Egkolfopoulou (2019).

[3] For example, in the classic model of Ferejohn (1986), voters' ability to use the threat of losing reelection to control incumbent behavior hinges on their observing a common performance signal; if the performance signals are individual-specific, voters' power over incumbents evaporates. Wood and Ravel (2018) discuss the normative consequences of microtargeting with a particular emphasis on how democracy can be harmed when citizens are only exposed to political appeals from the candidates and campaigns that they are predisposed to support.

While there are clear theoretical reasons to think that the mass adoption of social media would alter equilibrium campaign behavior, the examples above illustrate that differentiating consequences from correlates of communication technology is difficult. This paper attacks this challenge by introducing a new dataset of candidate-sponsored advertising, covering all advertising on TV and on Facebook by the universe of US congressional, statewide and state legislative campaigns in 2018. We combine information from the Facebook Ad Library API, which archives all political advertisements run on Facebook since late May 2018 (Nicas 2018), and the Wesleyan Media Project (WMP) database of political ads on television. We compare, on multiple dimensions of content and quantity, advertising on the two media *by the same candidate in the same race*. The use of within-candidate comparisons allows us to hold fixed candidate attributes, the competitiveness of the electoral environment, constituency characteristics, and other covariates that might otherwise bias a comparison of content across media.[4]

Comparing content across media within the same electoral campaign allows us to assess whether and how candidates take advantage of three opportunities afforded by social media: to increase advertising quantity thanks to its lower costs of production and placement, to use advertising for other purposes – like fundraising – that are impractical on television, and to strategically adapt their self-presentation to match the preferences of finely-segmented audiences. Because the latter in particular may involve subtle changes that are difficult to detect at scale, we build a rich dataset of finely detailed advertising features – choices of words, images, facial expressions, and references to political figures – that are measured in a consistent way across media. In addition to providing a comprehensive description of the content of political advertising both online and offline, these data elucidate how the capabilities of social media alter candidates' choices of issue agenda, tone, and ideological positioning in their advertising.

Our findings offer some confirmation but also a number of surprises relative to our ex ante theoretical expectations.[5] Notably, Facebook ads engage in less attacking of the opponent

[4]As we show later, the composition of candidates who advertise using the two modes is quite different, implying that naïve comparisons of means will be strongly biased by the selection of candidates into communication media.

[5]We posted a preanalysis plan (`https://osf.io/3px5b`) specifying analyses and expectations prior to the release of the Facebook Ad Library API.

and more promotion of the sponsoring candidate, compared to the same candidate's ads on TV. This finding suggests that fear of a voter backlash (Roese and Sande 1993, Lau, Sigelman and Rovner 2007, Dowling and Wichowsky 2015) is not a significant constraint on the negativity of campaign advertising: campaigns could, if they chose, use Facebook's targeting capability to show negative ads only to supporters, and avoid exposing the swing voters or opponents' supporters who are likely to exhibit backlash. Candidates do not appear to be implementing this strategy in significant numbers. Our results are instead consistent with an account of negative ads as demobilizing to the supporters of the opponent (Krupnikov 2011), as the more selected audience for Facebook ads leads to less rather than more negativity compared to TV.

Facebook ads contain less issue content than television ads by the same candidate. This is true even for relatively niche issues, where one might expect the targeting ability and low production cost of Facebook to make viable the production of ads hitting a wider range of issues not of sufficiently mainstream interest to justify the cost of a TV spot. We speculate that the compressed format and reduced attention that viewers give to online communications (Dunaway et al. 2018) counteracts these forces for more varied issue discussion.

Facebook ads are, however, more easily identifiable as partisan and more ideologically polarized than their TV counterparts. This is true both in the aggregate and within-candidate. Candidates do appear to take advantage of finer targeting to deliver more partisan messaging, which suggests that the capabilities of social media push candidates toward using ads more for mobilization than for persuasion. We also find that the ideological positioning of candidate messaging is more variable within-candidate on Facebook than on TV. That is, candidates are better able to fine-tune their message to comport with audience preferences on Facebook. In ads run by the same candidate in the same race, both issue mentions and perceived partisanship correlate with the demographic composition of the audience.

On the extensive margin, the set of candidates who advertise on Facebook is much broader than those who advertise on TV. The ability of ad spots on Facebook to be geographically targeted to avoid wasting impressions on viewers outside of an electoral district matters especially for down-ballot candidates; at the state house level, more than 10 times as many candidates advertise on Facebook than advertise on TV.

Taken together, these findings suggest that communication media have substantial im-

pact on candidates' communication strategy. The primary impact of an increase in targeting precision appears to be to allow candidates to reach their supporters more efficiently. For lower-resourced candidates, this is the difference between advertising and not. For higher-resourced candidates, the change leads to a shift of advertising messages away from those focused on persuasion – taking popular issue positions, attacking the opponent, and downplaying partisan cues – and towards those focused on mobilization. The political diversity of television audiences compels candidates to engage in attempts at persuasion; absent this constraint, candidates prefer to abandon most discussion of issues or comparison with the opponent and instead activate preexisting partisan loyalties. Given the connection between candidates' campaign issues and legislative activity once in office (Sulkin 2011), the relative lack of issue content on Facebook may lead to reduced citizen knowledge of candidates' policy platforms as the use of social media for political communication rises. We take up this and other implications of our results in the concluding section.

Theory and Empirical Predictions

Our theorizing begins with the two strategically important differences between television and online ads. First, there is a difference in cost. Because digital ads can be displayed to individual users instead of the entire local audience for a television program, online advertisements can be purchased in much smaller increments of impressions. Unlike television ads, the audience for online advertising need not follow the boundaries of television media markets ("Designated Market Areas" or DMAs), a fact which is especially important for political advertisers attempting to reach electorates in districts whose boundaries may not align well with those of a DMA. This increase in geographic alignment has the effect of (sometimes dramatically) lowering the effective cost per impression, as candidates need not waste impressions on viewers who cannot vote in their district. Moreover, the cost of production of a digital advertisement can be much lower than that on television.

Second, the precision of audience targeting varies across television and online advertising. While television advertisers can select programs with particular demographic profiles (Lovett and Peress 2015) in an attempt to reach a desired audience, television programs provide a

far from perfect partition of the ideological or partisan spectrum.[6] Social media firms, on the other hand, have an unusually rich set of individual-specific information, including self-identified interests, demographics, and media consumption choices that can be used to target advertisements to precise audiences: a campaign could, for instance, run an advertisement only to users who self-identify as political moderates, or users who follow the page of a particular national politician. Facebook offers advertisers the ability to go even a step further by specifying their own "custom audiences," for example lists developed from voter files and turnout history, or from contacts at campaign events.

We develop a series of hypotheses about the impact of social media technology on advertising quantity and content on the basis of these two observations. While most of the theoretical and empirical work on campaign advertising to date has focused on television (Freedman and Goldstein 1999, Goldstein and Freedman 2000; 2002*a*, Sides and Vavreck 2013, Krasno and Green 2008, Kahn and Kenney 1999, Fowler, Franz and Ridout 2016), our research nonetheless speaks to three relevant literatures: the question of whether the Internet equalizes the playing field between well-known candidates with abundant resources and upstart candidates, the strategic use of different communication modes, and the literature on the content of messaging in elections. We take on each in turn.

Equalizing or Normalizing?

First, we situate our work in the on-going debate on the impact of new technologies on electoral competition. Do digital media and the internet help equalize electoral competition (Barber 2001, Gainous and Wagner 2011; 2014) by allowing poorly financed candidates to compete on a more level field, or merely reinforce existing resource inequities (Bimber and Davis 2003, Hindman 2008, Stromer-Galley 2014, Gibson et al. 2003)?

We are interested in whether Facebook allows candidates with fewer resources (most often challengers and candidates down-ballot) to overcome resource imbalances in airing relatively costly television ads at the media market-level. The cost to advertise on television is often

[6] In the left panel of Figure 1, Lovett and Peress (2015) show that the vast majority of television programs have net conservative identifiers between -0.1 and 0.3 and the most liberal show has a net conservative identifier level of -0.285 and the most conservative show has a net conservative identifier of 0.692, implying that all television programs in their sample have nontrivial liberal and conservative audiences.

cited as part of the incumbency advantage at the federal level (Prior 2006).

We start by asking whether and how online advertising broadens the set of candidates who advertise by comparing both extensive and intensive margins of advertising on television to that on Facebook. Of particular interest is the ability of challengers to level the electoral playing field by using Facebook advertisements in electoral environments where television advertising is feasible for incumbents, but too costly for challengers. We also ask whether the much lower entry cost of Facebook advertising enables candidates in down-ballot races who are priced out of the market for TV ads to reach voters. Taken together, these analyses examine whether more financially constrained candidates, specifically challengers and state legislative candidates, advertise relatively more on Facebook, compared to their incumbent and up-ballot counterparts.

When and Where do Candidates Advertise?

Online advertising can be tailored to achieve different campaign goals than traditional advertising on television. The low cost of online advertising and the ability to target has potential implications for both when candidates choose to advertise and where these ads are displayed. Facebook offers two potential targeting advantages relative to television that may affect how campaigns use the platform. First, behavioral information can be used to serve engagement-oriented advertisements to well-off users who have expressed an interest in politics and are particularly likely to donate to a campaign. Second, Facebook advertisements can be targeted to much lower levels of geographic aggregation, such as the zip code, than television advertisements, which can only be geographically targeted at the DMA level. These capabilities of online advertising have implications for both when in the campaign candidates serve online advertisements and the spatial location of these advertisements.

Campaigns can use Facebook advertisements to solicit campaign resources in a way that is infeasible with television advertisements. While television advertisements may incidentally increase campaign contributions,[7] online advertising is better suited to soliciting campaign

[7] Urban and Niebler (2014) show that advertising that spills over from media markets in competitive states into uncompetitive states increases the probability of receiving campaign contributions from residents of the uncompetitive state who reside in the media market relative to other residents of the uncompetitive state who are not exposed to the advertisements.

resources and measuring return on investment. Online advertisements might serve a similar function to direct mail as a cost-effective tool for generating campaign resources for candidates (Hassell and Monson 2014).

Previous content analyses of online advertisements suggest that campaigns do use these ads to recruit volunteers and donations. Campaigns often link their advertisements to landing pages where users can sign up for a mailing list, register to volunteer, or make contributions. Online advertisements allow users to immediately follow through by performing an action at the request of the campaign. One analysis of the 2016 presidential campaign found that fewer than half of the digital ads that were sampled had a goal of voter persuasion (Franz et al. 2019). Similarly, in their study online display ads from the 2012 presidential campaign, Ballard, Hillygus and Konitzer (2016) coded only 37 percent of the ads as focusing on undecided or persuadable voters.

Financial contributions and volunteers are more valuable earlier in the campaign when candidates still have time to build out campaign infrastructure and use these resources to mobilize and persuade potential voters. TV ads, on the other hand, are most useful to campaigns in the days leading up to the election. Gerber et al.'s (2011) field experiment demonstrated that television advertising has a measurable persuasive effect on citizens' political preferences, but that the effects are short-lived, lasting no longer than a week or two. This research suggests that ads that attempt to persuade will have higher electoral returns as the election date approaches. Based on this logic, we expect that Facebook advertising will be used earlier in the campaign than television.

The targeting ability of online ads also has implications for their spatial location, relative to TV. One dimension in which this difference may manifest itself is the distribution of online ads to users who are ineligible to vote in the candidate's election but may be willing to contribute resources to the candidate.[8] While the different motivations of online and offline advertising would lead to the prediction that a higher proportion of online ads are sent to out-of-state residents, a countervailing factor that increases the relative proportion of

[8]The ideal data to examine this issue would include information on whether the audience member resides outside the electoral constituency of the candidate. However, the public Facebook database includes only the state of the advertising audience, limiting our analysis to that level of geographic aggregation. We calculate the proportion of the advertisement audience that resides outside the candidate's state and then aggregate to the candidate level.

TV ads outside of the state is the spatial structure of media markets, which often cross state lines. Candidates in electoral constituencies with a DMA that crosses state boundaries are often forced to waste advertising dollars on out-of-state viewers. In some cases, the lack of congruence between an electoral district and the containing DMA makes the effective price of TV ads so high that candidates cannot afford to advertise at all. We use our data to ask whether the proportion of ads displayed to out-of-state residents differs across Facebook and television, and how this difference varies with the electoral district–media market congruence.

What Messages do Candidates Include in Advertising?

A final relevant literature considers the content of political advertising and its determinants. One possibility is that campaigns emphasize a similar message across modes, what Bode et al. (2016) call "a single coherent message strategy." Alternatively, campaigns might adapt their message to meet the expectations of the varied audiences across media. As noted, television audiences are more politically diverse than targeted online audiences, suggesting that TV ads may be used to persuade the median voter while online messages may be directed at those who share an ideological or partisan affinity with the candidate. These different audiences may have different issue priorities and different expectations of campaign tone. To that end, we examine both in our analyses.

Scholars have long noted the potential of negative television advertisements to harm the sponsor of the advertisement, a backlash effect (Roese and Sande 1993). In their meta-analysis of 40 studies of negative campaigning, Lau, Sigelman and Rovner (2007) find citizens evaluate the sponsor of negative messages more negatively in 33 of the studies, and this effect is substantively large and statistically significant.[9] Because of the targeting limitations inherent on television, negative ads will be viewed by citizens who are favorably disposed toward the candidate who is attacked in the advertisement. As a result, these citizens may lower their evaluations of the sponsoring candidate and/or increase their likelihood of turning

[9]The magnitude of the backlash effect may be contingent upon advertising characteristics. Dowling and Wichowsky (2015) employ survey experiments to show that the sponsor of the advertisement conditions how respondents punish candidates for negative advertisements. When negative advertisements are sponsored by independent groups, opposing partisan voters do not punish the candidate as much as when the advertisement is directly sponsored by the candidate.

out. Inability to target the negative message to those citizens who will be most receptive to it increases the magnitude of the backlash effect. Thus, campaigns may allocate their negative messaging to online platforms where they can more precisely control who sees those ads, limiting the potential for a backlash. Our dataset thus provides an ideal setting to evaluate how constrained candidates are by fear of backlash effects. We ask: Do a higher proportion of ads exhibit a negative tone on Facebook relative to television?[10]

On the issue agenda of advertising, again expectations about the audience may drive the nature and level of issue discussion. Bode et al. (2016) found that Twitter was much less likely to provide discussion of issues than television advertising, but the study acknowledges that the character limitations of the medium (at the time 140) might restrict the capacity to raise issue or policy claims relative to other platforms. Still, issue discussion on Twitter does occur and Kang et al. (2018) found higher issue convergence within a campaign between advertising and Twitter and lower convergence between advertising and email in 2014 U.S. Senate Races. Twitter is closer to a broadcast medium than email, given that tweets are often seen and shared by journalists (and can therefore be seen by voters of different partisan and ideological dispositions). Email, in contrast, is targeted to individuals with a direct past tie to the campaign, either from a donation, sign-up, or request to receive emails. In that sense, email is conceptually more similar to Facebook advertising.[11] Our next research question, then, is: Do candidates discuss different policy issues on Facebook than on television?

We also investigate the degree of partisanship and polarization of ideological positioning in Facebook relative to television ads. On TV, candidates often downplay their partisan affiliation (Neiheisel and Niebler 2013) and, consistent with a goal of persuading on-the-fence swing voters, highlight issue stances where they are most different from their party (Henderson 2019). We ask: Does the more precise targeting afforded by Facebook give

[10]Initial work in this area has offered mixed support for this hypothesis. Roberts (2013), who focused on web-only videos posted during the 2004 and 2008 U.S. presidential campaigns, found more attacks online than on television. Anstead et al. (N.d.) found slightly more negativity (operationalized as mentions of another party) in the parties' Facebook ads than in their party election broadcasts during the 2017 general election in the United Kingdom. On the other hand, Bode et al. (2016) documents more negativity on television during the 2010 U.S. Senate campaigns than on campaigns' Twitter accounts, though the focus in that study was organic content instead of paid advertising.

[11]It is worth noting that the emergence of the online ad libraries makes the content of this advertising more publicly available, which may affect strategic behavior of campaigns.

candidates license to include more explicitly partisan messaging in their ads?

Finally, we investigate the effect on within-candidate variation in messaging. Narrowcast Facebook ads might allow the same candidate to offer different messages to different audiences, varying the content of advertising according to characteristics of the audience, rather than staking out a unified message strategy. We ask: Do Facebook ads have more within-candidate variation in ideological positioning than TV ads? Does the content of messaging correlate with characteristics of the audience, within-candidate?

Data and Methods

We draw on television and Facebook advertising data from all federal, statewide, and state legislative candidates. A challenge that has hampered the study of online political communication in the past, as Ballard, Hillygus and Konitzer (2016) discuss, is that many advertisements only appear briefly and are targeted to specific users in a way that is not visible to third parties. These limitations have prevented scholars from seeing the complete universe of campaign advertisements. Facebook, however, has recently released a database of information on the political advertisements run on its platform since May 2018 (Nicas 2018). We use this database to study campaign ads in the 2018 midterms. Although others have used these data (e.g. Edelson et al. (2019)), we believe ours to be the first study that examines not only the volume of spending but also the content of the ads, how candidate advertising strategies vary up and down the ballot, and when and where candidates deploy their advertisements.

Data on television advertising come from the Wesleyan Media Project (Fowler, Franz and Ridout 2016), which since 2010 has tracked political advertising on local broadcast, national broadcast and national cable television. The Wesleyan data rely upon ad tracking by a commercial firm, Kantar/CMAG, which detects and classifies ads aired in each of the 210 media markets in the United States. The data are at the level of the ad airing, so for each advertisement we observe the television station, media market, and time of day when the ads aired. The data also measure the estimated cost of each airing. In addition to these raw tracking data, Kantar/CMAG supplies Wesleyan with a video of each ad (the "creative"), and coders at the project classify each on a variety of characteristics, including its tone and

the issues that were mentioned.[12]

The Facebook Ad Library API includes a snapshot of the ad creative as it would have displayed to users, including any text, images, and video. It also reports the sponsor who financed the ad, the dates of the ad, the approximate number of impressions that the ad received, the cost of the ad, and aggregate demographic information on the age range, gender, and state of residence of the ad's audience.[13] Facebook includes both candidate and issue advertisements in this database. We focus on candidate-sponsored ads. The data were accessed via Facebook's API, which we had access to in Fall 2018.[14] The API allowed for bulk downloads of ad data based on a supplied list of key words or page IDs. We collected a comprehensive list of candidates' Facebook page IDs and downloaded all ads from these pages.

From the television and Facebook ad creatives, we extracted a large set of features by processing the ad's text, images, video, and audio through commercially available computer vision, audio transcription, and natural language processing software. The extracted features are described in complete detail in Appendix B. Features include word frequencies in text and transcribed audio, descriptive tags for images included in the ad, and attributes such as emotion classifications and predicted ages of human faces detected in the ad's images.

We use these features to detect the occurrence of negative advertisements and advertisements by issue area. We had research assistants classify a training sample of the Facebook advertisements on issue and tone dimensions and then used these classifications, along with classifications of the TV ads in the Wesleyan data, as the basis for a supervised learning classification procedure, described in detail in Appendix B. The fitted model from this process then produces predicted values for tone and issue content, which we use as our measure of these quantities for all ads in the data set. Using the same process and the same ad features, we also produced predictions of the partisanship and campaign finance-based ideology score

[12]For more information on the details of human coding, including reliability statistics, see Appendix A.

[13]The ad audience information is based on impressions, not targeting decisions by the ad buyer. This is unlike the data released by Google about political ads purchased on its own platform (`https://transparencyreport.google.com/political-ads/home`).

[14]Facebook also has a publicly searchable web-based library, located here: `https://www.facebook.com/ads/library/report/`, although the public version and the API appear to operate independently and therefore our results may differ from what is available through the publicly available web-based version.

of the ad sponsor (Bonica 2014).[15] To aggregate these ad-level measures to the candidate level, we calculated expenditure-weighted averages of message content for each candidate.[16]

We have also gathered information on the partisanship, incumbency status, and campaign resources of the federal and state-level candidates from the two major parties.[17] The final dataset contains 7298 total candidates: 1032 who advertised on both Facebook and television, 242 who advertised only on TV, and 6024 who advertised only on Facebook. Additional summary statistics are in Appendix A.

Which Campaigns Advertise Online and Offline?

We first provide some descriptive analysis of the aggregate use of both Facebook and television media, by office sought. In Figure 1, we show the distribution of spending by congressional, gubernatorial, other statewide,[18] and state legislative candidates, for both Facebook (1a) and television (1b).[19] The densities include only candidates with positive spending on the indicated mode; we examine the mass at zero in Figure 2, described in the next paragraph. On average, races for governor and US Senate saw the most Facebook spending, with spending by the median candidate[20] around $100K. Some candidates, however, spent

[15]Training data for the prediction of sponsor party comes from the candidate registration statements and thus is defined for all ads, not just those in the subset that were coded by WMP. Training data for prediction of the Bonica CFscore comes from the 2018 update of the DIME data (Bonica 2013) and is defined for all federal-level candidates with at least 25 unique contributors. Hence, predicted partisanship scores are in-sample predictions for all ads in the dataset, and predicted CFscores are in-sample predictions for most federal candidates.

[16]The pre-analysis plan specifies two methods of weighting, using expenditures and impressions. Our TV data, however, does not contain a measure of impressions but only an estimate of the total cost of the spot. As all of our analyses compare across media, we require a consistently available weighting across both modes, and hence we focus on the expenditure-weighted values. Expenditures and impressions are very highly correlated and there are unlikely to be substantial differences between the two.

[17]This candidate-level information is drawn from the OpenSecrets.org and FollowTheMoney.org databases for federal and state candidates, respectively.

[18]The "other statewide" category includes all non-gubernatorial offices elected on a statewide basis. Examples of such offices are secretaries of state, attorneys general, or utility commissioners.

[19]Facebook reports spending levels at the creative-level in bins rather than exact amounts. Thus, to estimate the total spending level for each candidate, we aggregated the midpoint of the reported spending level for a given advertisement across all advertisement that the candidate sponsors. Appendix A shows that this ad-level approximation method produces total spending numbers that are consistent with the actual page-level aggregate totals reported by Facebook.

[20]Among candidates who spent anything on the mode.

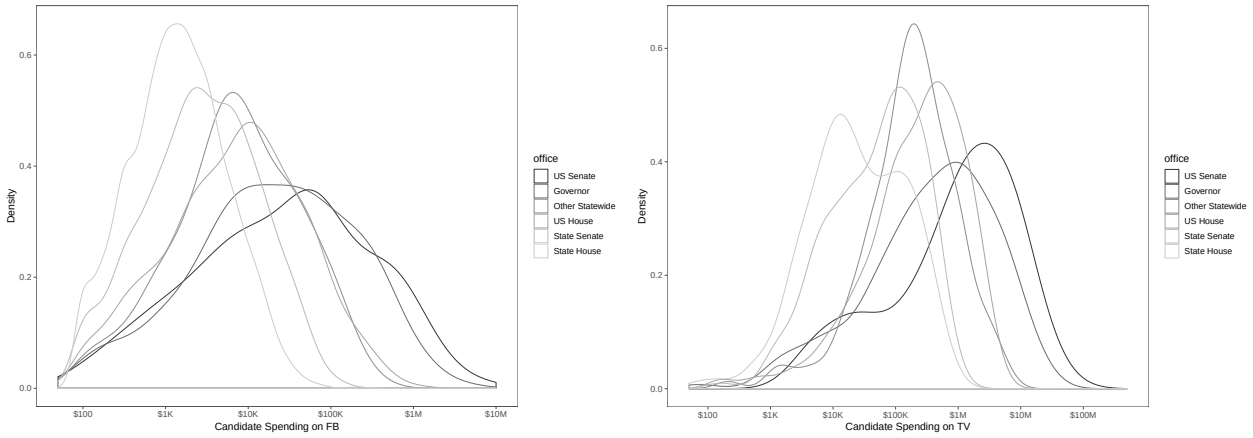

(a) Facebook (b) TV

Figure 1: Density of candidate-level spending on each medium, by office. Plots condition on non-zero spending on the indicated medium (i.e., they exclude the point masses at zero).

up to $10 million on ads on the platform. Unsurprisingly, total spending declines as we go down the ballot, with state house candidates spending the least. The same pattern also holds for TV, but with typical spending numbers increased by an order of magnitude or more: for Senate and governor races, median TV spending was in the neighborhood of $1M. Candidates for all levels of office spent more on television than on Facebook ads. A relative comparison between the two panels reveals that the cross-office differential is compressed on the Facebook platform relative to television: the difference in typical spending between Senate or governor and state house races on Facebook is about two orders of magnitude, compared to closer to three on television.

Figure 2 examines the extensive rather than intensive margin of advertising, by medium. The panels plot the proportion of all candidates with non-zero spending on Facebook (2a) and TV (2b) ads, by office and incumbency status. The effect of Facebook's relatively low cost in expanding access to advertising is clearly evident in the down-ballot races: less than 10% of state house and senate candidates advertised on television, whereas closer to 40% advertised on Facebook. Facebook also appears to narrow the incumbent-challenger gap in access in most offices. In fact, in the two farthest down-ballot categories, challengers were

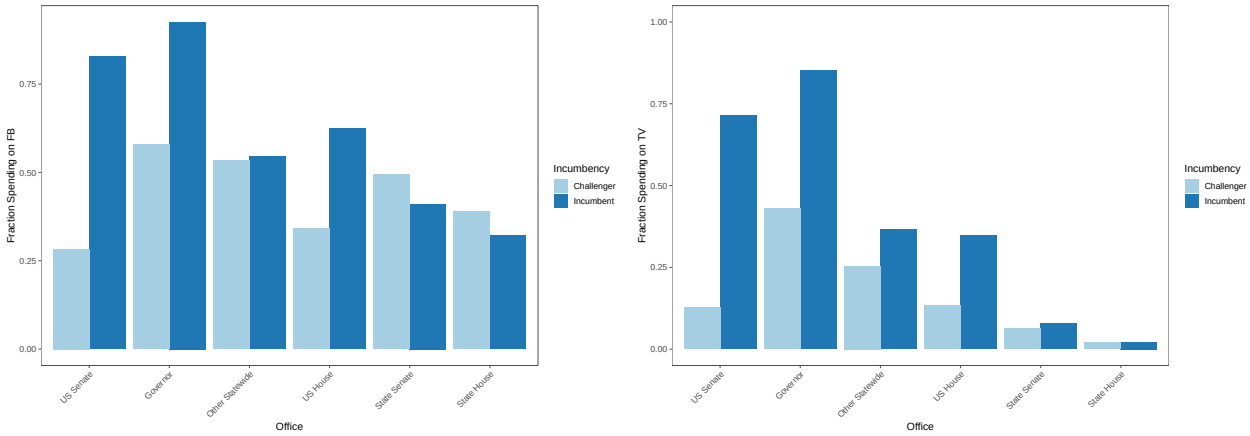

(a) Facebook (b) TV

Figure 2: Fraction of candidates with positive spending on each medium, by office and incumbency status.

more likely to advertise on Facebook than their incumbent counterparts.[21]

As outlined earlier, our first research question concerns whether and how online advertising broadens the set of candidates who advertise. Figures 1 and 2 make clear that both the composition of candidates who advertise, and the level of expenditures they invest, are quite different across media. We show this in regression form in Table 1, in which the dependent variable is total advertising spending on Facebook or television between May 24, 2018, and Election Day.[22] We estimate the following regression with candidate fixed effects:

$$AdSpending_{ik} = \alpha_i + \gamma Facebook_k + Facebook_k CandCovar_i \delta + \epsilon_{ik} \tag{1}$$

The dataset for this regression contains one observation for each candidate's spending on television advertisements and one observation for each candidate's spending on Facebook

[21] The "challenger" categories here include candidates who ran and lost in a primary. This is the primary reason that the advertising rates for US Senate challengers are so low: many non-viable candidates file to run in Senate primaries but raise and spend very little money.

[22] As Martin and Peskowitz (2018) show, candidate expenditures are almost never made directly to television stations, but are instead mediated by political consultants. Our primary interest in this study lies in the intensity and use of advertising after this intermediation occurs, so we estimate our models with the direct cost of television and online advertising instead of adding the markup that political consultants extract from their clients.

advertisements. The α_i are candidate-specific fixed effects, $Facebook_k$ is a binary indicator for whether the particular observation corresponds to Facebook advertising expenditures, and ϵ_{ik} is an idiosyncratic error term. The inclusion of the candidate fixed effects means that our estimates use only within-candidate variation to identify the Facebook effect γ. $CandCovar_i$ is a row vector of candidate covariates: an indicator for whether the candidate is a challenger, and indicators for the office sought by the candidate. These candidate covariates cannot be directly included in the regression specification, because none of these characteristics vary within candidates and we include the candidate-specific fixed effects α_i, in the equation. We can, however, interact these covariates with the $Facebook$ indicator to determine how these covariates are associated with the intensity of using Facebook advertisements relative to television advertisements. In this and all regressions reported in the paper, we cluster standard errors at the level of the candidate.

Table 1: Within-candidate regressions of spending levels on FB indicator.

	Spending ($K)		
	(1)	(2)	(3)
Facebook	−120.81***	−114.15***	−1,266.21***
	(14.97)	(18.11)	(258.17)
x Incumbent		−22.97	
		(32.05)	
x US Senate			−458.18
			(595.77)
x Other Statewide			1,056.44***
			(262.47)
x US House			1,055.55***
			(259.17)
x State Senate			1,255.96***
			(258.18)
x State House			1,265.11***
			(258.17)
Candidate FE:	Y	Y	Y
N	14,530	14,530	14,530
R^2	0.57	0.57	0.63

*p < .1; **p < .05; ***p < .01
Robust standard errors (clustered by candidate) in parentheses. An observation is a candidate x medium. The excluded category in column (3), which includes office interactions, is gubernatorial candidates.

The first column of estimates in Table 1 shows that spending on Facebook ads is significantly less than spending on television ads. As the specification includes candidate fixed effects, this is not simply an artifact of differences in financial resources across the pools of

candidates who advertise on each medium. The mean within-candidate difference is on the order of $100K. The second column interacts the Facebook indicator with a dummy for incumbency; the point estimate is negative, indicating that the TV-Facebook gap in spending is larger for incumbents than for challengers, but this difference is not statistically different from zero. The final column of estimates reveals a clear gradient from top to bottom of the ballot; Senate and Governor candidates spend well over $1M more on television than on Facebook on average; the gap is closer to $200K for US House and non-Governor statewide candidates, and zero for state house and state senate candidates.

Consistent with the idea of Facebook providing a large effective cost reduction, the most financially constrained candidates rely on Facebook more, relatively speaking, than candidates with typically less binding financial constraints. The existence of online advertising allows down-ballot candidates to make appeals to the voting public that they cannot afford to make on television. The existence of this platform, then, with a wide reach and low cost to entry, has facilitated new means of connecting with potential supporters.

Advertisement Timing and Geographic Targeting

We next examine how candidates differentially time the release of and geographically target their advertising on the two media. On timing, evidence suggests the persuasive effects of advertising are short-lived (Gerber et al. 2011) and thus advertisements whose goal is to persuade voters will have higher electoral returns as the election date approaches. Facebook ads may be used for a more diverse range of goals - such as fundraising - than are TV ads, and thus may have higher value earlier in the campaign than TV ads.

Before moving to regression analysis, it is instructive to examine time trends in the raw data. Figure 3b shows the timing of advertising on TV (in dollar terms) between June 1 and Election Day. There is a steady ramp-up of spending as the election approaches.[23]

Figure 3a shows the same time trends on Facebook. The overall level is much lower, with even late campaign spending on Facebook lower than television spending in the summer months of 2018. But the relative pattern is even more skewed toward the end of the campaign

[23]The bump in governor spending in June-August is due to late primaries in a few states. The regular dips down from the overall upward trend are weekends: television viewing, particularly for the kinds of programs on which political ads run, drops substantially on weekends.

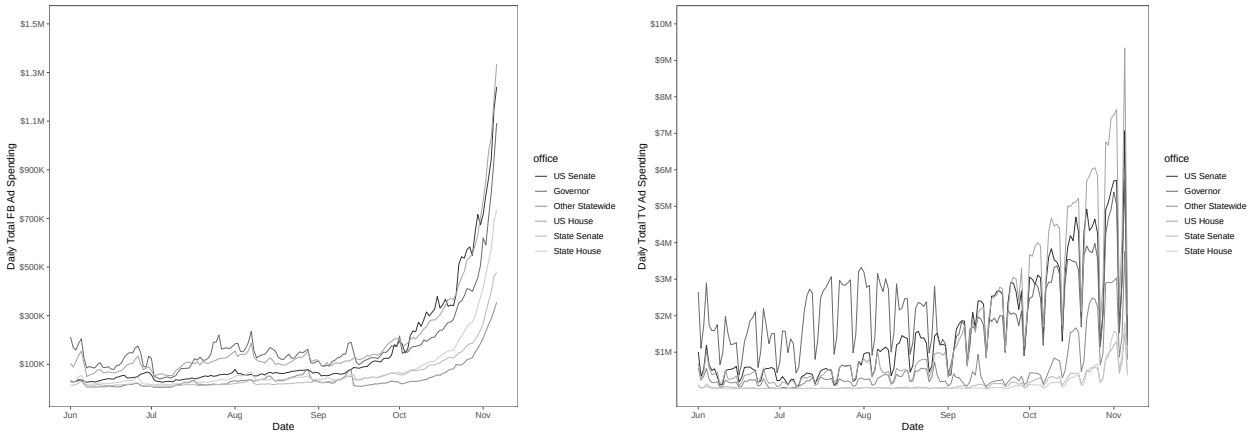

(a) Facebook (b) TV

Figure 3: Daily spending by office over the course of the campaign. Plots include spending on the indicated mode by all 7298 candidates in the data.

than that on television. Across all offices, daily spending is flat from June until the end of September. Only in October does spending accelerate before reaching its peak on Election Day. Television spending, in comparison, begins its rise more than a month earlier.

One possibility is that congestion due to the fixed number of TV ad spots available in the later days of the campaign pushes TV spending earlier; congestion on Facebook is much less binding because the online platform does not have the requirement that all viewers on the platform at a given time see the same content. Another possibility is that the apparent pattern is due to compositional changes over time; perhaps the kind of campaigns that engage in both TV and Facebook advertising indeed use Facebook relatively early and TV relatively late, but there is a large group of Facebook-only advertisers who enter at the end of the campaign.

We address this question with a regression of the timing of campaign spending by medium, controlling for candidate fixed effects. We regress the quantity of spending on candidate fixed effects plus our medium indicator interacted with a full set of time-to-election dummy variables, defined weekly. The regression specification is described by:

$$AdvSpending_{iwk} = \alpha_i + Facebook_k \eta_w^{FB} + (1 - Facebook_k)\eta_w^{TV} + \epsilon_{iwk} \tag{2}$$

The two sets of week fixed effects η^{FB} and η^{TV} correspond to advertising on Facebook and television respectively, and allow for general time-patterns that flexibly differ between the two modes. This specification allows us to determine how online advertising's relative intensity varies as the general election date approaches. Results are displayed graphically in Figure 4 and demonstrate that the TV/Facebook ratio is indeed increasing over time, as predicted; however, TV advertising dominates at all stages of the campaign. In other words, within candidate, TV advertising accelerates in the final months of the campaign at a faster rate than spending on Facebook.[24] This result suggests that the pattern in Figure 3 is less a function of differences in congestion across medium and more the result of over-time changes in the set of candidates advertising on each. Facebook-only advertisers also tend to be relatively light advertisers, and candidates with relatively low advertising budgets focus their spending (on all modes) toward the end of the campaign.

Next, we examine the spatial distribution of political advertisements, specifically the proportion of each candidate's ads that are viewed by out-of-state residents. If Facebook ads are used for purposes other than voter persuasion or mobilization, then candidates may be more likely to use Facebook ads to target out-of-state voters, who cannot vote for the candidate but can contribute in other ways. At the ad level, we compute the fraction of impressions seen by users in the state in which the candidate is running for office.[25] We then aggregate to the candidate level by computing a weighted average, weighting by expenditures. Our estimating equation is:

$$PropOutState_{ik} = \alpha_i + \gamma Facebook_k + \epsilon_{ik} \tag{3}$$

The timing and spatial targeting effects might interact with one another. Campaigns may deploy their advertisements early and outside of their electoral constituencies in order

[24] Week fixed effects here are weeks to the general election date. There are some primary elections in some states that occur in June or later of 2018; we do not differentiate here between spending targeted toward the general or the primary.

[25] Facebook's API provides a breakdown of impressions by state for each ad. For TV, we use the fraction of DMA population living in zip codes that are in the state in which the race was held as our proxy for share of impressions in state. This is a proxy and not an exact measure because the composition of viewers can vary by time and program; we do not have sufficient information on the geographic distribution of program-level viewing to estimate the in-state proportion at date, time, or program level.

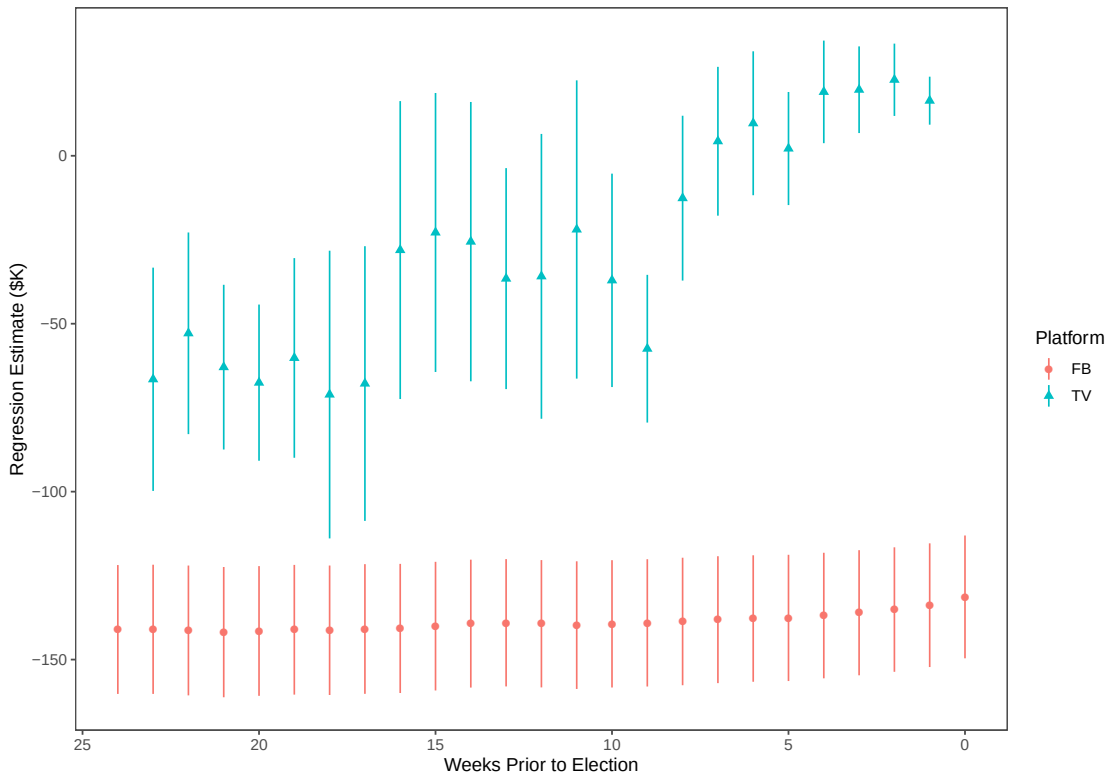

Figure 4: Regression estimates of weeks-to-election effects, by medium. Coefficients are extracted from the regression model specified in equation 2. The regression includes candidate fixed effects and separate sets of weeks-to-election dummies for every week out to 24 weeks prior to election day. Bars are asymptotic 95% confidence intervals, using standard errors clustered at candidate level.

to generate campaign resources. To investigate this possibility we estimate the following regression:

$$\begin{aligned} PropOutState_{itk} =& \alpha_i + \beta_1 DaysUntilGeneralElection_t + \\ & \beta_2 DaysUntilGeneralElection_t Facebook_k + \\ & \gamma Facebook_k + \epsilon_{itk} \end{aligned}$$

Results, in Table 2, show that in fact Facebook ads are *less* likely to be seen by viewers outside the candidate's state. This is true throughout the campaign, as the days-to-election

trend is tiny and statistically insignificant.[26] Although some candidates are certainly using Facebook to appeal for donations from out of state residents, it appears that such candidates are a relatively small minority. The dominant effect of Facebook is that, by providing finer-grained geographic targeting than television media markets allow, candidates can waste fewer impressions over state lines.

Table 2: Within-candidate regressions of in-state proportion on medium.

	Proportion In-State	
	(1)	(2)
Facebook	0.073***	0.047***
	(0.021)	(0.011)
Days to Election		0.0001
		(0.0001)
FB x Days to Election		−0.0001
		(0.0001)
Candidate FE:	Y	
N	8,081	483,589
R^2	0.867	0.665

*p < .1; **p < .05; ***p < .01
Robust standard errors (clustered by candidate) in parentheses. An observation is a candidate x medium in column (1), and a candidate x medium x day in column (2). Proportion in-state is the expenditure-weighted average fraction of impressions reaching viewers in the state of the election.

Finally, we examine how the level of congruence between a candidate's electoral constituency and DMA influences the allocation of advertising across television and Facebook. Candidates who run in low congruence districts waste a larger portion of their television impressions when they advertise to audience members who cannot vote in the election than candidates who run in high congruence districts. As a result, we expect that candidates in low congruence districts will allocate more of their advertising expenditures to Facebook.[27]

[26] Although the point estimate is negative in magnitude, implying that Facebook ads are (slightly) more likely to be seen out-of-state in the early days of a campaign.

[27] We define congruence as the share of the DMA's population that is located in the relevant congressional district or state. In cases where the electoral district includes multiple DMAs, we define this variable as the maximum value of congruence across all of the DMAs. This definition of congruence is slightly different from the definition used by Snyder and Strömberg (2010) in their analysis of the effects of newspaper circulation congruence with congressional districts. Snyder and Strömberg (2010) weight market share by reader share to arrive at their measure of congruence. Lacking information on the spatial distribution of the television audience, we instead opt for a less refined measure.

Because of the difficulty of calculating congruence at the state legislative district level, we restrict the analyses here to the sample of congressional and gubernatorial candidates.

The estimates in Table 3 indicate that the greater the congruence between the media market and the candidate's electoral district, the less the candidate spends on Facebook, which is consistent with our expectations. Magnitudes are such that DMA congruence explains essentially all of the TV-Facebook differential estimated in Table 1 for Congressional candidates: a congressional candidate running in a zero-congruence district would be predicted to spend about the same on both modes, whereas a candidate running in a perfectly-congruent district would be expected to spend about $890K less on Facebook. The large effect of congruence on spending suggests that television and Facebook advertising are close substitutes, as the effective price differential that candidates face explains a large amount of the variation in usage.

Table 3: Within-candidate regressions of spending levels on DMA congruence.

	Spending ($K)
Facebook	−28.516
	(62.723)
x DMA Congruence	−888.766***
	(211.147)
Candidate FE:	Y
N	3,718
R^2	0.592

*p < .1; **p < .05; ***p < .01
Robust standard errors (clustered by candidate) in parentheses. An observation is a candidate x medium. Sample is restricted to US House, Senate and statewide candidates.

Advertisement Content

In this section, we move from utilization of advertising, in dollar terms, to the actual content of ads. We investigate the effects of the lower production costs and greater precision in audience targeting on the message that candidates present to voters. Our general expectations are that the first will allow for more experimentation and variation in messaging; the second will allow candidates to offer more polarizing messages.

Tone As we noted earlier, scholars have long noted the potential of negative television advertisements to harm the sponsor through backlash effects (Roese and Sande 1993). One reason why this might be the case is that the negative advertisements are viewed by citizens who are favorably disposed toward the candidate who is attacked in the advertisement. The differential ability to target online and offline advertisements raises the possibility that candidates may allocate their negative messaging to online platforms where they can more precisely control the audience for their messages.

To examine the tone of advertisements across television and online, we operationalize negativity through references to an opponent where ads that solely mention an opponent save for the sponsor name are classified as attack, ads that solely reference the favored candidate are positive, and ads that mention both candidates are contrast (Goldstein and Freedman 2002*b*). We estimate the following regression, with dependent variable $Tone_{ik}$ equal to the candidate-medium average tone from the predictive model detailed in Appendix B:

$$Tone_{ik} = \alpha_i + \gamma Facebook_k + Facebook_k CandCovar_i \delta + \epsilon_{ik} \tag{4}$$

Again, the inclusion of candidate fixed effects (α_i) eliminates differences in message content due to candidate-level fixed attributes such as district partisanship and demographics or race competitiveness, partisanship, and so on, any of which might correlate with the candidate's propensity to use Facebook advertising. Because, as Figure 4 shows, the relative usage of the media also differs over the campaign and message content may evolve secularly over campaign time, we also estimate versions of the specification that control for candidate-week rather than candidate fixed effects, thus eliminating any confounding by within-candidate time trends of general form. We also estimate a version with fixed effects at the candidate-election (where election can be either the 2018 primary or general) level, controlling for possible confounding due to correlation of the primary season with Facebook use.[28]

We are primarily interested in how the relative intensity of advertising tone differs across Facebook and television advertisements for the same candidate, which is captured by the coefficient γ. The interaction effects (δ) capture how this varies with candidate characteristics

[28]Only the candidate-fixed-effect version was specified in the pre-analysis plan. We include the candidate-week and candidate-(general/primary) fixed effect versions because of the possibility that timing drives the result in the main specification. As the results show, estimates are very consistent across the three versions.

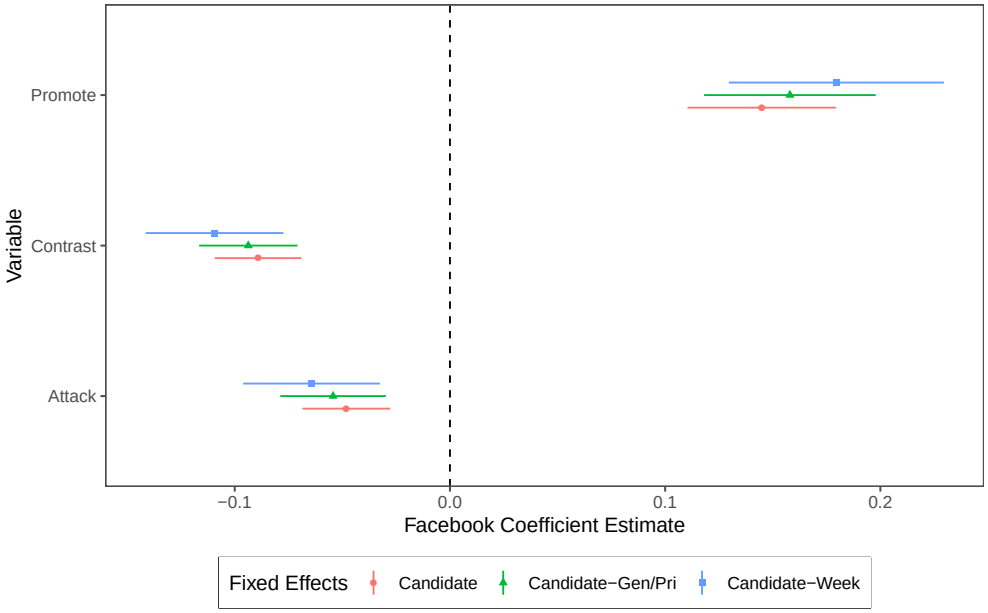

Figure 5: Effect of Facebook on ad tone, within candidate. Bars are asymptotic 95% confidence intervals, using standard errors clustered at candidate level.

such as office or incumbency status.[29]

We indeed find differences in tone across media, with Facebook ads significantly more positive than television ads (Figure 5). The magnitude of the effects are consistent across all three specifications of fixed effects, though standard errors widen as we get to finer-grained specifications. Furthermore, television ads are significantly more likely to be contrast or attack ads than are ads on Facebook. Advertising on Facebook is clearly more positive, even within the same candidate at the same time in the campaign cycle.

This result is more consistent with an account of negative ads as demobilizing to swing voters or supporters of the opponent (Ansolabehere and Iyengar 1996, Krupnikov 2011) than with backlash effects. Because Facebook ads are often run to custom audiences that the campaign generates from their own lists of contributors and volunteers, the audience is likely to be friendlier on average to the candidate than a television audience. The fact that usage of attack ads declines rather than increases in this context implies that candidates

[29] Across the board, we find interaction terms to be noisily estimated and insignificant, and we hence focus in the main text on the main effects. For completeness, interaction effects are presented in Appendix C.

prefer to show attack ads to opponents rather than to supporters, which comports with the demobilization but not the backlash account of negative ads.

Issue content We use the same specifications to analyze the issue content of advertising across media. As detailed in the theory section, we expect that the ability to target ads to a narrower group of viewers than television allows may induce campaigns to message on more niche issue areas that would go unmentioned in a broad-audience ad. We focus on the set of issue areas defined by the WMP[30] and estimate regressions of the form:

$$IssueScore_{ik}^{j} = \alpha_i^j + \gamma^j Facebook_k + Facebook_k CandCovar_i \delta^j + \epsilon_{ik}^j \tag{5}$$

where j indexes issue areas, and $IssueScore_{ik}^{j}$ is the (expenditure-weighted) average predicted probability of mention of issue j for ads sponsored by candidate i on medium k. As in the tone regressions, we also run analogous specifications where fixed effects are included at the candidate-week or candidate-election level.

Figure 6 shows the impact of medium on the likelihood that a variety of specific issues are mentioned in advertising. Estimates are almost uniformly negative. In any case where we can reject the null hypothesis of no difference at the 5% level, the difference is negative, and point estimates in the baseline specification with candidate fixed effects are positive (but substantively small) for only one issue category, the environment. For important issue areas like the economy, health care, immigration, and education, the magnitudes are substantively large, in the range of 3-6 percentage points. This effect size is roughly a third to a half of the baseline predicted mention rate of these categories in the Facebook data (see Appendix A for summary statistics).

We also construct summary measures of the "issue diversity" of a candidate's advertising, and the total share of advertising that references any policy issue (see equation 6) (as opposed to advertising focused purely on candidate characteristics or experience). To measure issue

[30] We collapsed some detailed issue categories into broader composite issues to improve statistical power of the classifiers. For instance, the WMP issue categories "Taxes" and "Deficit / Budget / Debt" are combined into a single Fiscal Policy category; an ad in the training data is tagged as Fiscal if it mentions either of these sub-issues. Additionally, we exclude some issues for which the WMP human codes had low inter-coder reliability ($\kappa < 0.7$). Appendix A provides a detailed accounting of our choices of which issue classifications to include.

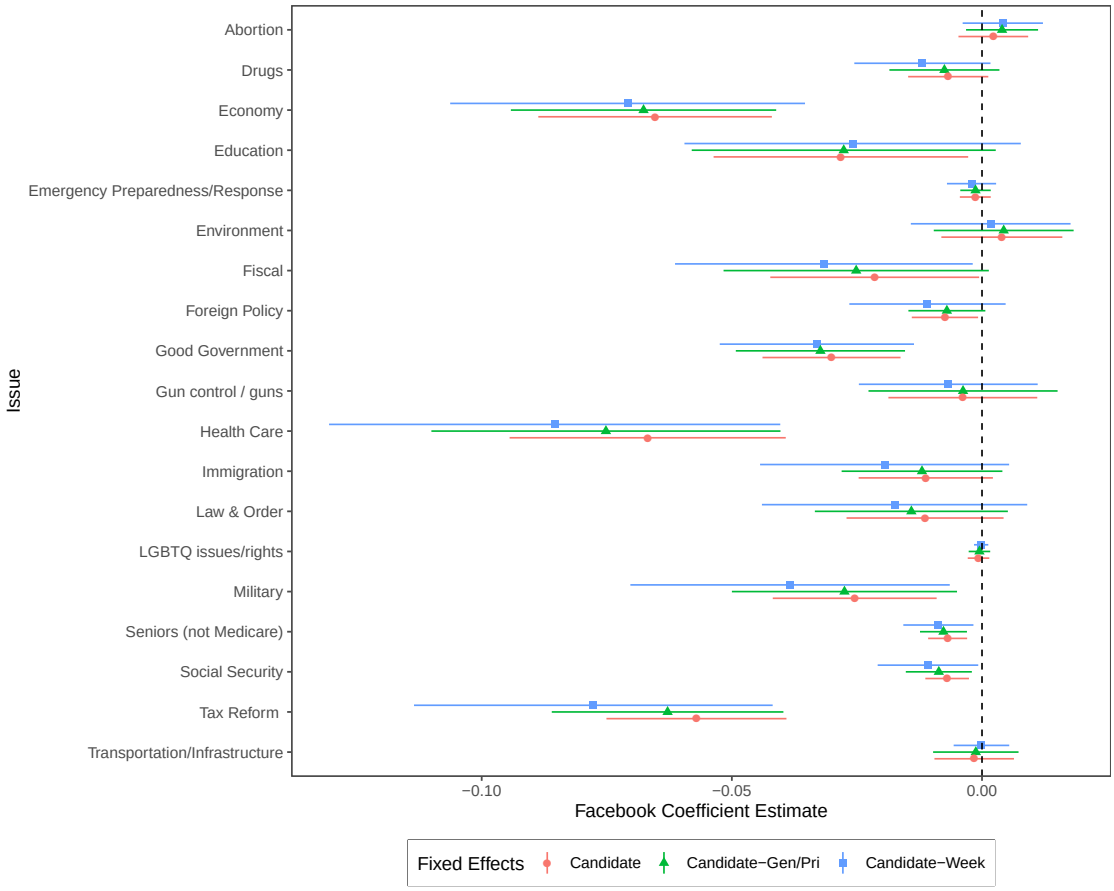

Figure 6: Effect of Facebook on mention of specific issues, within candidate. Bars are asymptotic 95% confidence intervals, using standard errors clustered at candidate level.

diversity, we construct the Herfindahl-Hirschman index of a candidate's advertising, which is the sum of squared shares of a candidate's advertising devoted to each issue (expressed in equation 7).

$$AnyIssue_{ik} = \frac{\sum_l \left(\max_j IssueScore^j_{ikl}\right) Expenditure_{ikl}}{\sum_l Expenditure_{ikl}} \tag{6}$$

$$IssueHHI_{ik} = \sum_j \left(\frac{\sum_l IssueScore^j_{ikl} Expenditure_{ikl}}{\sum_l Expenditure_{ikl}}\right)^2 \tag{7}$$

We regress these measures on the same right-hand side variables described in the issue-specific regressions. We find that Facebook ads are approximately 10 percentage points less likely to mention one of our issue areas than are television ads. However, the within-candidate issue HHI does decline by a small amount, indicating that Facebook ads have lower issue concentration (i.e., greater issue diversity) than do television ads by the same candidate (Figure 7).

Taken together, results on issue content suggest that Facebook does allow candidates to broaden the set of issues they touch on in their advertising, but that this effect is swamped by an overall decline in total issue content. As the proportion of attack advertising also declines, this result is consistent with the Geer (2006) result on the greater factual content of negative ads. It appears that candidates use Facebook's targeting capabilities not to take positions on controversial public policies for narrowly-targeted audiences, but instead to focus on purely promotional, valence-oriented ads aimed at mobilizing their base of existing supporters.[31]

The reasons for this are unclear, though we can speculate. One possibility is that with TV ads, campaigns get 30 seconds of a viewer's attention whereas with Facebook ads, which users can easily scroll past, a campaign may only have a few seconds to capture the viewer's attention, and thus it may be difficult to deliver more complex and issue-focused messages.[32] It is also possible that the diversity of goals on Facebook (e.g., email acquisition and fundraising) ends up watering down the issue content.

Party / Ideology We next examine the effect of Facebook on the partisanship and ideological polarization of messages contained in campaign ads, using the same within-candidate design as used to examine effects on the other content outcomes.[33] Numerous popular accounts and some scholarly research (Lelkes, Sood and Iyengar 2017) point to internet access and online communication as generative of a more polarized and aggressively partisan political discourse. We have the opportunity to test whether candidate-sponsored messaging is

[31] Our measure of "attack" ads require the ad to specifically attack the candidate's opponent. This does not rule out "going negative" in a more general sense: negative attacks against the opposing party or an opposing party leader are not counted by this measure. Our subsequent analysis of partisan content, however, is likely to pick up these references, as they are highly indicative of party affiliation.

[32] Note, however, that a substantial minority ($\sim 35\%$) of Facebook ads include embedded video in similar lengths to TV ads.

[33] Note: analyses in this sub-section were not included in the pre-analysis plan.

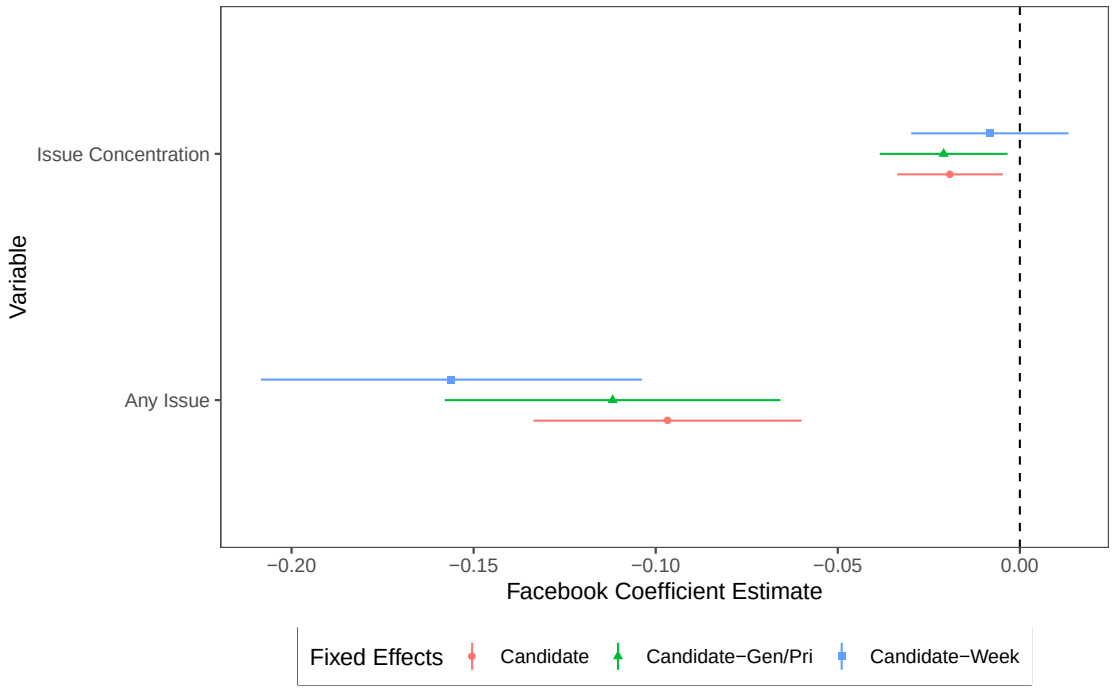

Figure 7: Effect of Facebook on issue diversity, within candidate. Bars are asymptotic 95% confidence intervals, using standard errors clustered at candidate level.

more clearly partisan or polarized on ideological lines online (on Facebook) as compared to TV, holding candidate attributes fixed.

Political ads do not, of course, generally come with an ideological label; the ad's ideological location must be inferred from its content. Candidates more often than not avoid explicit party labels in advertising (Neiheisel and Niebler 2013), but voters can use other cues to infer partisanship (Henderson 2019). Analogous to the application in Gentzkow, Shapiro and Taddy (2019) to politicians' speech in Congress, we seek to measure the distinctiveness of an ad's content along party lines. Does it use words or phrases that are used disproportionately by elected officials of one party? Do its choices of images and political references make it easy or difficult for viewers to infer the party or left-right positioning of the sponsor?

To operationalize this idea, we fit classification models of the party label of an ad's sponsor, and the ad's donation-based ideology score (CFscore), on the same set of ad features we used to predict issue content and tone. This is a much easier problem than predicting

issue content, since the party label is observed for all candidates in the case of party, and nearly all federal candidates in the case of CFscore. The predicted value from these models become the basis for outcome variables in within-candidate regressions. The interpretation of these variables is simple: a score of 0.99 on our party measure, for instance, indicates that our model is almost certain that the ad was run by a Republican candidate. A score of 0.5 on our CFscore prediction indicates that the model expects on the basis of the ad's features that the ad sponsor has CFscore of 0.5.

To measure if Facebook encourages candidates to take more partisan or ideologically extreme stances in advertising, we take the absolute value of the party / CFscore predictions and average within candidate-medium (again weighting by expenditures). We also compute the standard deviation of the party and CFscore predictions within candidate (also weighted by expenditure) as a measure of the degree of within-candidate heterogeneity in presentation. A candidate with a consistent ideological message throughout all his/her ads will have low standard deviation of these measures, whereas a candidate offering a liberal-friendly message to liberal audiences and a conservative-friendly message to conservative audiences will have high standard deviation. We estimate the same within candidate (or within candidate-week or candidate-election) specifications as on the other outcome measures to rule out the possibility that the mixture of advertisers differs across media on the ideological dimension.

Our regression results, displayed in Figure 8, show that Facebook increases both the extremism and the variability of ideological positioning within candidate on both measures. The substantive size of the effect on the extremism measures is fairly large. On CFscore, the difference between co-chair of the House Progressive Caucus Pramila Jayapal (CFscore = -1.59) and co-chair of the House Problem Solvers Caucus Josh Gottheimer (CFScore = -0.94) is about 0.65 points. The estimated Facebook effect in our main specification is about 0.125 points, or roughly 20% of this difference between prominent members of the progressive and moderate wings of the Democratic caucus. We emphasize that this is a within-candidate effect.[34]

The party score effect is smaller and the confidence interval overlaps zero. For comparison, Jayapal's Facebook ads have average predicted probability of Republican sponsorship of 0.02,

[34]The difference in partisan and ideological distinctiveness is also visible in the aggregate distribution of predicted scores; see Figures B.2 and B.3 in Appendix B.

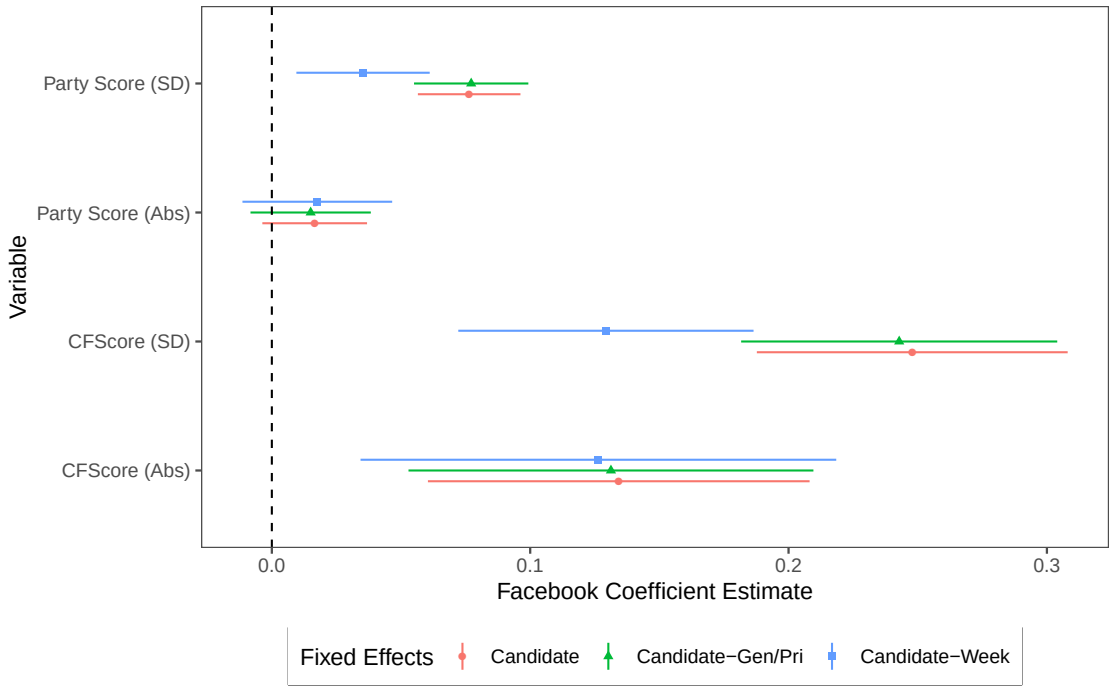

Figure 8: Effect of Facebook on predictions of party and campaign-finance based ideology score, within candidate. Rows labeled (Abs) are the absolute value of the indicated variable (averaged by candidate). For Party Score, we use the absolute value of the predicted probability that the ad was run by a Republican candidate minus one-half. For CFScore we use the absolute value of the predicted CFScore of the sponsoring candidate. Rows labeled (SD) are the within-candidate standard deviation of the indicated variable. Bars are asymptotic 95% confidence intervals, using standard errors clustered at candidate level.

translating to party extremism score of $\text{abs}(0.02 - 0.5) = 0.48$. Gottheimer's Facebook ads have corresponding probability of 0.21 or extremism score of 0.29,[35] for a difference of 0.19. Our point estimate of the effect on the party extremism score is about 0.02 or about 10% of the Jayapal-Gottheimer difference.

[35]Gottheimer's TV ads had average party extremism score of 0.25, lower than his score on Facebook; Jayapal did not advertise on television in 2018. (Note that because Jayapal did not advertise on television, her ads do not inform the within-candidate estimates; we use her example merely to illustrate the size of the within-candidate effect relative to the overall scale of the measure.)

Message Targeting on Facebook Finally, we examine how candidates varied their messages with characteristics of the audience on Facebook.[36] We ask whether, holding the candidate sponsor fixed, issue content, tone, or ideological positioning vary according to the audience receiving the message. Although the Facebook database provides only a fairly crude set of audience characteristics - age, gender and state of residence - these nonetheless correlate with issue positions, issue interest and attention, and ideological or partisan preferences (Aldrich et al. 2019). We estimate specifications of the form:

$$y_{il} = \alpha_i + \beta' x_{il} + \epsilon_{il} \tag{8}$$

Where y is an outcome variable (one of the issue, tone, or ideological predicted values introduced previously), l indexes ad spots, and i indexes candidates. α_i is a candidate fixed effect, and x_{il} is a vector of audience impression shares across demographic groups. β is the vector of coefficients of interest capturing the correlation between, for example, the share of the audience for an ad[37] that is female and between the ages of 18-25, and the ad's predicted probability of mentioning an education issue. Our specification of x is maximally flexible, given the data available: we allow for separate coefficients for each Gender-Age cell.

Several interesting patterns emerge. Candidates discuss education issues more to users, and especially female users, in the 25-44 age range. Candidates discuss health care more prominently in ads targeted to female users, and to users in either the two oldest or two youngest age cohorts. Economic and fiscal policy issues get more mention in ads targeted to male users, particularly those in the middle age cohorts. Coefficient magnitudes can be interpreted as predicted change in message for a 0-1 change in the audience share of the corresponding demographic cell. E.g., an ad whose audience was exclusively men ages 18-25 would be expected to be about 10 percentage points less likely to mention health care than an ad whose audience was exclusively women ages 18-25. These effects are quite large relative to the mean incidence of the issue tags in the data.

Again, these estimates all include candidate fixed effects, so we are not simply picking up differences in constituency characteristics (e.g. that candidates representing older districts

[36]Note: analyses in this sub-section were not included in the pre-analysis plan.

[37]Audience shares are measured as fractions of the total ad impressions viewed by users in the given demographic cell.

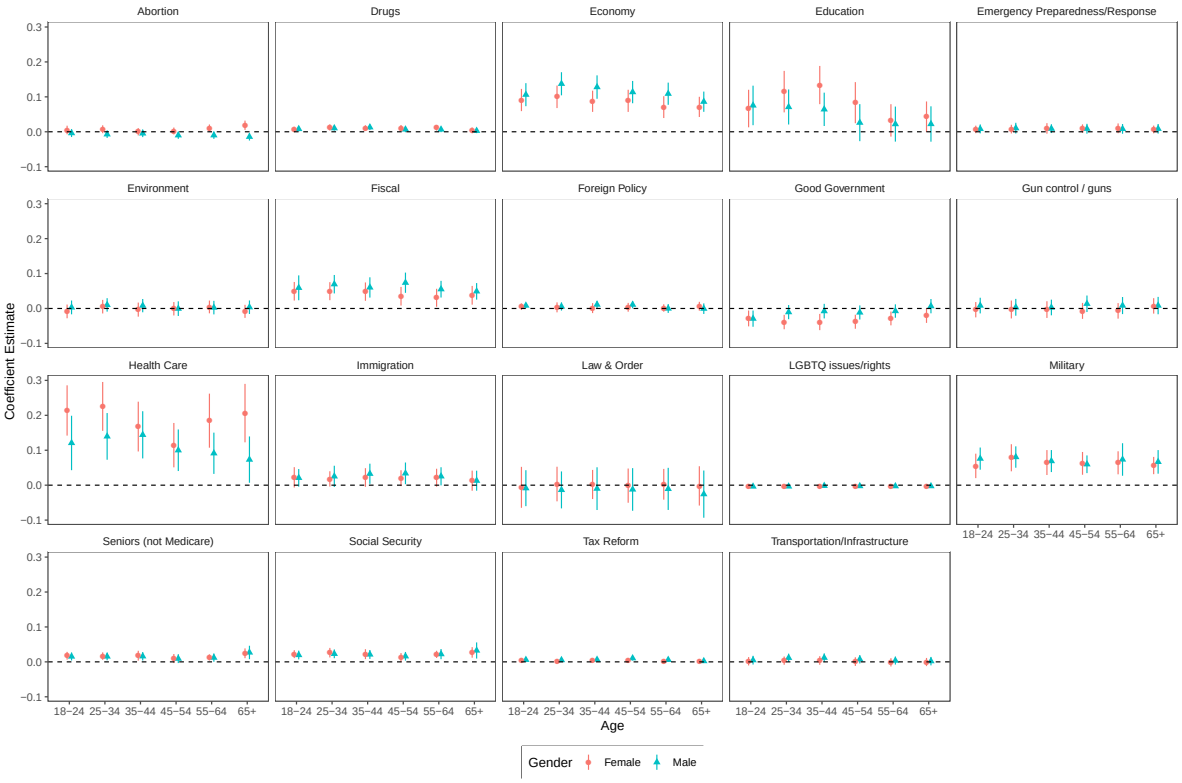

Figure 9: Regression coefficients of predicted ad issue content on audience demographic, within candidate. The sample is all Facebook advertisements. Bars are asymptotic 95% confidence intervals, using standard errors clustered at candidate level. The omitted category on both dimensions is users whose corresponding demographic variable (Age or Gender) is unknown.

might run more ads mentioning Social Security). These are differences in the way that the same candidate selectively presents his/her agenda to voters, depending on the kinds of voters being targeted.

We run the same analysis on ad-level estimated partisanship. Results are presented in the same format in Figure 10. Results indicate that candidates present themselves as more right wing (higher values on either the CFScore or Republican probability scale) to more male audiences. Given the gender gap in partisanship (Aldrich et al. 2019), this suggests candidates are pandering to audience preferences on this dimension. However, the relationship to audience age, which also correlates with party ID and voting, is weak.

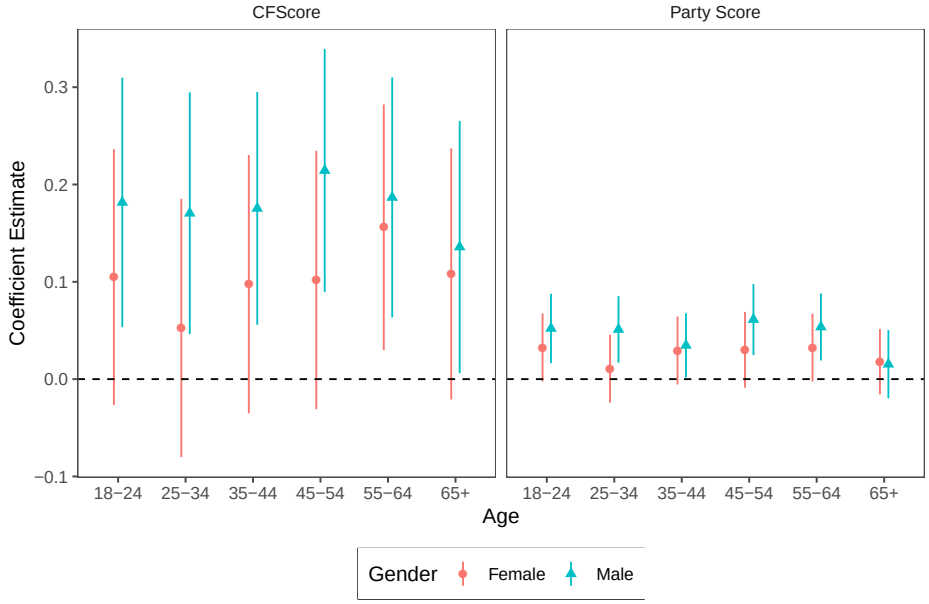

Figure 10: Regression coefficients of predicted ad ideology and partisanship on audience demographic, within candidate. The sample is all Facebook advertisements. Bars are asymptotic 95% confidence intervals, using standard errors clustered at candidate level. The omitted category on both dimensions is users whose corresponding demographic variable (Age or Gender) is unknown.

Implications

This analysis is the first comprehensive accounting of advertising on Facebook by political candidates up and down the ballot and the first to examine how campaigns use Facebook and television advertising. We examined not only the aggregate differences across candidates but also within-candidate differences in spending and content across online and television media. Our analysis reveals important differences in how campaigns use these media. We conclude by briefly discussing the implications of our main findings for American democracy and articulating an agenda for future research.

The ability to create and deploy advertisements in small cost increments online has a dramatic impact on which candidates use paid advertising. Our findings tend to support the equalization side of the debate over whether new technologies enable less-resourced candi-

dates to compete with those who have traditionally had more resources. Facebook does not enable challengers to compete ad for ad with incumbents, especially in the races at the top of the ballot, but it does seem to create a more even playing field than television. Voters see disproportionately more Facebook ads from challengers and down-ballot candidates relative to television. Moreover, the finding that candidates rely on Facebook more when the television media market is incongruent with their district shows that citizens who reside in these districts learn more about these candidates than if TV were the only medium available. Our analysis suggests that voters are exposed to advertising from a wider set of candidates than if Facebook did not exist. Facebook appears to foster more intense electoral competition, which may increase citizen awareness of state and local candidates and candidates running for office in electoral constituencies that are incongruent with television markets. These are largely positive developments for American democracy.

Our findings also speak to the tone of the campaign to which voters are exposed. In spite of the Internet's reputation as an uncivil cesspool, a world of online advertising does not necessarily mean a more negative politics. In fact, advertising on Facebook engages in significantly less attacking of the opponent than advertising on television. There is a robust scholarly debate about the consequences of negative advertising for citizen knowledge, participation in politics, and attitudes toward the political process. For those who ascribe to the view that negative advertisements demobilize and increase cynicism among voters, the lower level of opponent attack advertising on Facebook is reassuring. The decrease in negativity, however, comes with a cost: online advertising has less issue discussion than on television. Advertisements are an important tool for increasing citizens' issue knowledge and holding politicians accountable for their policy choices in office. A shift away from television and towards social media advertising may thus reduce this component of voters' knowledge.

Less issue discussion does not necessarily rule out different issue discussion, and we suggested initially that Facebook might allow campaigns to emphasize different issues in their online and TV spots. Campaigns might promote different issue priorities to different audiences. We find little evidence of a "these issues here" and "those issues there" approach to TV and Facebook; across all issues that we examined, discussion is no greater on Facebook than television. For many issues, the difference is strictly negative and substantively large. It might be axiomatic that issue discussion in candidate ads is better for voters than issue-less

appeals, and so there is some reason for concern that Facebook does not contribute to the information environment in ways that allow voters to make decisions based on candidate policy proposals.

Though less issue-centered, the messages that candidates choose to include in their Facebook ads are more easily identifiable as partisan and more clearly ideological than those they include in TV spots. These three differences - reduced negativity, lower issue content, and increased partisanship - all point toward the use of social media ads for mobilization of existing supporters as opposed to persuasion of marginal voters. Social media can thus be expected to increase the polarization of the information environment that voters experience in campaigns, with Republican-leaning voters seeing mostly pro-Republican ads with little attempt to engage the opponent's positions, and vice versa for Democratic voters.

On both the issue coverage and partisanship dimensions, there is more variation in message content within candidate on Facebook than on TV. Candidates use the targeting capability of Facebook to tailor their messages to different audiences, which is difficult to do on TV thanks to the diverse audience of most TV programs. This increase in within-campaign variation in messages, and the fact that message content correlates fairly strongly with viewer attributes, implies that Facebook contributes to a fragmentation in the information environment across the electorate.

The normative implications of political advertising on social media, then, are mixed. Social media lowers barriers to entry and thereby exposes voters to information about a broader set of candidates and offices. On the other hand, for already well-funded campaigns, it shifts campaign strategy away from moderation in service of persuading voters on the fence and towards mobilizing the base. TV ads are still by far the dominant mode of communication and are unlikely to disappear in coming campaigns, and so the effects of the introduction of social media will not be felt immediately but will take time to play out. And, the rise of addressable technologies on television means that TV may become more similar to Facebook over time. Still, scholars should take advantage of the difference in targeting capability while it lasts by continuing to compare how campaigns use paid advertising on social media and on television and by documenting how the use of these tools change in coming electoral cycles.

As campaigns learn about the capabilities offered by digital campaigning and targeting technology is improved for online platforms and rolled out for television (Bruell 2018), we

expect that campaigns will continue to develop new approaches to persuade and mobilize voters. We also believe that researchers will be able to build off our work to better understand the causes and consequences of different advertising strategies on online and TV modes. Incorporating information on the cost of buying narrowly-targeted advertisements and the choice space of targeting options that advertising platforms offer campaigns will help researchers understand how campaigns trade off the electoral benefits of targeting with the increased cost. Comparing how citizens respond to messages delivered online and on television will better inform our understanding of the effects of online advertising platforms on citizen participation, election outcomes, and attitudes toward the political system.

References

Aldrich, John H, Jamie L Carson, Brad T Gomez and David W Rohde. 2019. *Change and Continuity in the 2016 and 2018 Elections*. CQ Press.

Ansolabehere, Stephen and Shanto Iyengar. 1996. *Going negative: how political advertisements shrink and polarize the electorate*. New York: Free Press.

Anstead, Nick, João Carlos Magalhães, Richard Stupart and Damian Tambini. N.d. "Political Advertising on Facebook: The Case of the 2017 United Kingdom General Election." London School of Economics Working Paper.

Ballard, Andrew O., D. Sunshine Hillygus and Tobias Konitzer. 2016. "Campaigning Online: Web Display Ads in the 2012 Presidential Campaign." *PS: Political Science & Politics* 49(3):414–419.

Barber, Benjamin R. 2001. "The Uncertainty of Digital Politics: Democracy's Uneasy Relationship with Information Technology." *Harvard International Review* 23(1):42–47.

Benoit, Kenneth, Kohei Watanabe, Haiyan Wang, Paul Nulty, Adam Obeng, Stefan Müller and Akitaka Matsuo. 2018. "quanteda: An R package for the quantitative analysis of textual data." *Journal of Open Source Software* 3(30):774.
URL: *https://quanteda.io*

Bimber, Bruce and Richard Davis. 2003. *Campaigning Online: The Internet in US Elections*. New York: Oxford University Press.

Bode, Leticia, David S. Lassen, Young Mie Kim, Dhavan V. Shah, Erika Franklin Fowler, Travis Ridout and Michael Franz. 2016. "Coherent Campaigns? Campaign Broadcast and Social Messaging." *Online Information Review* 40(5):580–594.

Bonica, Adam. 2013. "Database on Ideology, Money in Politics, and Elections." `http://data.stanford.edu/dime`.

Bonica, Adam. 2014. "Mapping the Ideological Marketplace." *American Journal of Political Science* 58(2):367–386.

Bruell, Alexandra. 2018. "NBCU Joins 'OpenAP' TV Ad Consortium, Licenses Data Assets." Wall Street Journal. Published: 2018-04-19. Accessed: 2018-04-23.
URL: *https://www.wsj.com/articles/nbcu-joins-openap-tv-ad-consortium-licenses-data-assets-1524132000*

Dowling, Conor M. and Amber Wichowsky. 2015. "Attacks Without Consequences? Candidates, Parties, Groups, and the Changing Face of Negative Advertising." *American Journal of Political Science* 59(1):19–36.

Dunaway, Johanna, Kathleen Searles, Mingxiao Sui and Newly Paul. 2018. "News attention in a mobile era." *Journal of Computer-Mediated Communication* 23(2):107–124.

Edelson, Laura, Shikhar Sakhuja, Ratan Dey and Damon McCoy. 2019. "An Analysis of United States Online Political Advertising Transparency." NYU Working Paper.
URL: *https://arxiv.org/abs/1902.04385*

Egkolfopoulou, Misyrlena. 2019. "Facebook Is Big Winner in Democrats' 2020 Presidential Debates." Bloomberg. Published: 2019-07-22. Accessed: 2019-08-05.
URL: *https://www.bloomberg.com/news/articles/2019-07-22/social-media-ads-could-bleed-some-2020-democratic-candidates-dry?srnd=politics-vp*

Ferejohn, John. 1986. "Incumbent performance and electoral control." *Public Choice* 50:5–25.
URL: *http://www.jstor.org/stable/30024650*

Fowler, Erika Franklin, Michael M. Franz and Travis N. Ridout. 2016. *Political Advertising in the United States*. Boulder, Colo.: Westview Press.

Franz, Michael M, Erika Franklin Fowler, Travis Ridout and Meredith Yiran Wang. 2019. "The Issue Focus of Online and Television Advertising in the 2016 Presidential Campaign." *American Politics Research* p. 1532673X19875722.

Freedman, Paul and Ken Goldstein. 1999. "Measuring media exposure and the effects of negative campaign ads." pp. 1189–1208.

Gainous, Jason and Kevin M. Wagner. 2011. *Rebooting American Politics: The Internet Revolution*. Lanham, Maryland: Rowman & Littlefield.

Gainous, Jason and Kevin M. Wagner. 2014. *Tweeting to Power: The Social Media Revolution in American Politics*. New York: Oxford University Press.

Geer, John Gray. 2006. *In defense of negativity: Attack ads in presidential campaigns*. Chicago: University of Chicago Press.

Gentzkow, Matthew, Jesse M. Shapiro and Matt Taddy. 2019. "Measuring Group Differences in High-Dimensional Choices: Method and Application to Congressional Speech." *Econometrica* 87(4):1307–1340.

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

Gibson, Rachel K., Michael Margolis, David Resnick and Stephen J. Ward. 2003. "Election Campaigning on the WWW in the USA and UK: A Comparative Analysis." *Party Politics* 9(1):47–75.

Goldstein, K. and P. Freedman. 2000. "New evidence for new arguments: Money and advertising in the 1996 Senate elections." 62(4):1087–1108.

Goldstein, Ken and Paul Freedman. 2002*a*. "Campaign advertising and voter turnout: New evidence for a stimulation effect." 64(3):721–740.

Goldstein, Ken and Paul Freedman. 2002*b*. "Lessons learned: Campaign advertising in the 2000 elections." *Political Communication* 19(1):5–28.

Hassell, Hans J.G. and J. Quin Monson. 2014. "Campaign Targets and Messages in Direct Mail Fundraising." *Political Behavior* 36(2):359–376.

Henderson, John. 2019. "Issue Distancing in Congressional Elections.".
URL: *http://www.jahenderson.com/IssueDistancing-WP-2019.pdf?attredirects=0*

Hindman, Matthew. 2008. *The Myth of Digital Democracy*. Princeton, N.J.: Princeton University Press.

Kahn, Kim Fridkin and Patrick J. Kenney. 1999. "Do negative campaigns mobilize or suppress turnout? Clarifying the relationship between negativity and participation." 93(4):877–889.

Kang, Taewoo, Erika Franklin Fowler, Michael M. Franz and Travis N. Ridout. 2018. "Issue Consistency? Comparing Television Advertising, Tweets, and E-mail in the 2014 Senate Campaigns." *Political Communication* 35(1):32–49.

Krasno, J. S. and D. P. Green. 2008. "Do televised presidential ads increase voter turnout? Evidence from a natural experiment." 70(1):245–261.

Krupnikov, Yanna. 2011. "When Does Negativity Demobilize? Tracing the Conditional Effect of Negative." *American Journal of Political Science* 55(4):797–813.

Lau, Richard R., Lee Sigelman and Ivy Brown Rovner. 2007. "The Effects of Negative Political Campaigns: A Meta-Analytic Reassessment." *Journal of Politics* 69(4):1176–1209.

Lelkes, Yphtach, Gaurav Sood and Shanto Iyengar. 2017. "The hostile audience: The effect of access to broadband internet on partisan affect." *American Journal of Political Science* 61(1):5–20.

Lovett, Mitchell and Michael Peress. 2015. "Targeting Political Advertising on Television." *Quarterly Journal of Political Science* 10(3):391–432.

Martin, Gregory J. and Zachary Peskowitz. 2018. "Agency Problems in Political Campaigns: Media Buying and Consulting." *American Political Science Review* 112(2):231–248.

Neiheisel, Jacob and Sarah Niebler. 2013. "The Use of Party Brand Labels in Congressional Election Campaigns." *Legislative Studies Quarterly* 38(3):377–403.

Nicas, Jack. 2018. "Facebook to Require Verified Identities for Future Political Ads." New York Times. Published: 2018-04-06. Accessed: 2018-04-06.
URL: *https://www.nytimes.com/2018/04/06/business/facebook-verification-ads.html*

Plasser, Fritz and Gunda Plasser. 2002. *Global political campaigning: A worldwide analysis of campaign professionals and their practices.* Greenwood Publishing Group.

Prior, Markus. 2006. "The Incumbent in the Living Room: The Rise of Television and the Incumbency Advantage in U.S. House Elections." *Journal of Politics* 68(3):657–673.

Roberts, Chris. 2013. "A Functional Analysis Comparison of Web-Only Advertisements and Traditional Television Advertisements from the 2004 and 2008 Presidential Campaigns." *Journalism & Mass Communication Quarterly* 90(1):23–38.

Roese, Neal J. and Gerald N. Sande. 1993. "Backlash Effects in Attack Politics." *Journal of Applied Social Psychology* 23(8):632–653.

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

Snyder, Jr., James M. and David Strömberg. 2010. "Press Coverage and Political Accountability." *Journal of Political Economy* 118(2):355–408.

Stromer-Galley, Jennifer. 2014. *Presidential Campaigning in the Internet Age*. New York: Oxford University Press.

Sulkin, Tracy. 2011. *The Legislative Legacy of Congressional Campaigns*. New York: Cambridge University Press.

Urban, Carly and Sarah Niebler. 2014. "Dollars on the Sidewalk: Should U.S. Presidential Candidates Advertise in Uncontested States?" *American Journal of Political Science* 58:322–336.

Wager, Stefan, Sida Wang and Percy Liang. 2013. Dropout Training As Adaptive Regularization. In *Proceedings of the 26th International Conference on Neural Information Processing Systems - Volume 1*. NIPS'13 USA: Curran Associates Inc. pp. 351–359. eventplace: Lake Tahoe, Nevada.
URL: *http://dl.acm.org/citation.cfm?id=2999611.2999651*

Wood, Abby K. and Ann M. Ravel. 2018. "Fool Me Once: Regulating Fake News and Other Online Political Advertising." *Southern California Law Review* 91(6):1223–1278.

Young, Dannagal Goldthwaite. 2019. *Irony and Outrage: The Polarized Landscape of Rage, Fear, and Laughter in the United States*. Oxford University Press.

Appendices

A Data collection procedures and summary statistics

This section describes the procedures used in collecting advertising data on TV and on Facebook for major party candidates at the federal and state levels.

A.1 TV data collection

Data on television advertising comes from Kantar/CMAG, which is available through the Wesleyan Media Project, and includes the most comprehensive information available on local broadcast, national network and national cable advertising in each of the 210 media markets in the United States from January 1, 2017 through Election Day 2018.[1] For this analysis, we rely on Kantar/CMAG's classification of sponsor (to identify all of the candidate-sponsored advertisements) and their classification of level of focus (to identify all of the federal and state-level advertisements). All federal, gubernatorial, other statewide offices and state legislative ads were human content coded. In the modeling section, we restrict the analysis to ads that aired on or after May 24, 2018 to match the Facebook timeframe.

Table A1 shows the resulting numbers of unique ad creatives, candidates, and races in the Kantar/CMAG dataset for the period from May 24, 2018 through November 6, 2018.

Item	Count
N Creatives	5,199
N Candidates	1,289
N Races	730

Table A1: Counts of unique creatives, candidates, and races in the TV ads data.

[1] Although the company deploys discovery technology to identify new ads in every state and has tracking technology in all 210 media markets, not every media market has a tracking device capable of recognizing new ads. This means that ads for down ballot races likely to air only in smaller markets without discovery technology may be missed. See `http://mediaproject.wesleyan.edu/discovery-markets/` for more information.

A.2 Facebook data collection

Advertising data on Facebook was extracted from the Facebook Ad Library API, to which we had access in Fall 2018. The Ad Library includes all ads run on the platform that were tagged as political, beginning in late May of 2018. Facebook uses a combination of self-reports by advertisers, algorithmic detection, and user reports to flag ads as political. Despite evidence of instability on other issues in the beta API, we were unable to locate any candidates known from other sources to be advertising on Facebook who did not appear in the library. The more common problem we encountered was false positives. Some examples are ads run by nonprofit foundations of former politicians (e.g. the Carter Foundation), university programs in public policy, or news outlets. All of these kinds of ads were frequently tagged as "political" though they are not advertising on behalf of a candidate, party, or interest group. There are also pages that masquerade as candidate pages that actually attack the candidate; for example, in 2018, House Majority PAC ran ads on a page called "Meet the Real Troy Balderson."

Since the 2016 election, Facebook has required all ads run on the platform to be associated with a defined Facebook page. There are verification requirements associated with creating a page and running ads on its behalf, including verifying a physical address. We use this requirement to associate ads with candidates and to extract the universe of ads run by a given candidate on the platform. Specifically, we located page IDs associated with candidate pages, and then requested from the library API all ads associated with that page ID. To collect page IDs, we used the API search function to search for every candidate name appearing in our set of candidates. We manually examined the results, extracting page IDs that appeared to be official candidate pages and excluding third party groups. The mapping of candidates to pages is not 1 to 1; one candidate can have multiple pages, although the vast majority have just one. To summarize, the sampling process was four-step:

1. Generate a list of candidates by combining all unique candidates appearing in the FEC candidate master file[2] or in the FollowTheMoney.org databases of statewide candidate fundraising and state legislative candidate fundrasing.[3]

[2] `https://www.fec.gov/data/advanced/?tab=bulk-data`
[3] `https://www.followthemoney.org/`

2. Search for candidate names using the Facebook Ad Library API search function. We used a variety of transformations of candidates' name, state and office sought to form search strings, e.g. `firstname lastname`, or `lastname state`.[4] extract all unique page names and page IDs from the resulting ads.

3. Examine the resulting page IDs and manually confirm that they correspond to a candidate-sponsored page. Limit to manually verified page IDs.

4. Extract all ads associated with identified page IDs from the Ad Library API.

Table A2 shows the resulting numbers of unique ad creatives, pages, candidates, and races that this process produced.

Item	Count
N Creatives	359451
N Pages	7108
N Candidates	7056
N Races	3732

Table A2: Counts of unique creatives, pages, candidates, and races in the Facebook ads data.

A.3 Comparing API data with Facebook Aggregate Report

Our data come from the Facebook API, as noted, but Facebook also makes available an aggregated report (now published daily but previously weekly) that lists the to-date totals for all sponsors of ads on their platform. The post-election weekly report from November 2018 listed page name and disclaimer as the unit of analysis, without the page's unique numeric ID code. Variations of page name spelling and ad disclaimers would produce multiple rows of data in the aggregate report. We appended page ID onto as many rows as we could in the aggregate report. Then, we aggregated the creative level API data to the page ID and merged with the aggregate report. The goal is to compare the API estimate of each page's spending with Facebook's disclosed actual spending for each page name.

[4]The search function uses fuzzy rather than exact matching.

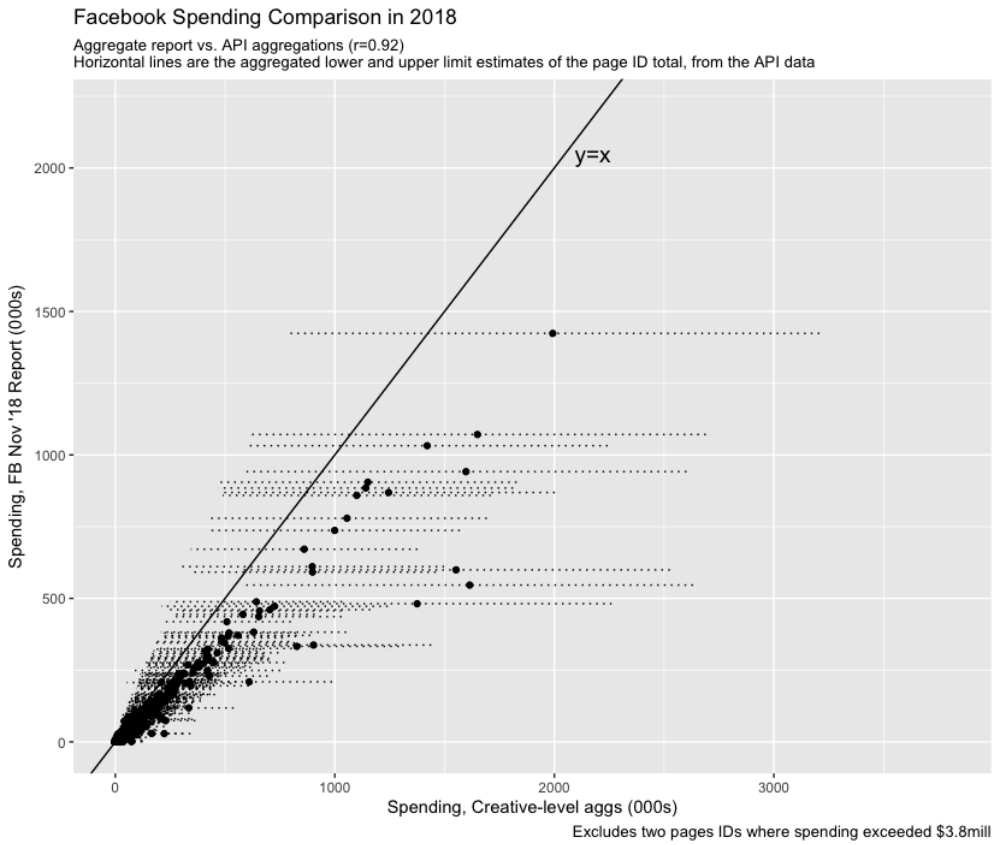

Figure A1

Recall that because the API data list spending per creative in bins, we used the midpoint of the bins to estimate the creative-level spending. But sponsors may have paid on the upper or lower end of that bin, which is only an issue if that tended to happen systematically above or below the midpoint.

Figure A1 plots the estimates by page ID that we obtain from the API with the totals as reported in the aggregate report. The two estimates are very highly correlated (r=0.92), but using the midpoint of the range on creative cost results in higher page ID-level estimates than the FB report. This suggests that sponsors tend to buy ads on the lower end of the binned totals. Still, the lower bound on those estimates always intersects with the FB aggregate report total, which gives us high confidence in using the estimates from the API.

A.4 Issue selection and consolidation

Human content coding was performed by research assistants at four different institutions. Training and supervision was provided by the same staff and coders went through multiple rounds of content coding and assessment to ensure consistency across coders and institutions. Overall, the team double-coded 1,595 television ads and 576 Facebook ads, which were used to calculate inter-coder reliability (ICR) statistics. Table A3 shows the complete list of issues coded by WMP, and the composite issue area to which the detailed issue is assigned, if any. The table also displays which issues had sufficiently high inter-coder reliability to include in our issue-by-issue and issue-diversity analyses.

A.5 Ad content summary statistics

Table A4 shows summary statistics of the content features of advertising on Facebook in our sample. Statistics are reported for *candidate*-level averages. For example, ads from the candidate whose advertising is maximally weighted to the Foreign Policy issue area have average score of 0.73 in the Foreign Policy domain.

The first row of A4 reports the fraction of ad impressions viewed by users in the same state in which the candidate was seeking office. 94% of the average candidate's ad impressions reach users in the candidate's state. There are, however, a small number of candidates who use Facebook advertising to reach primarily or even exclusively out-of-state users, perhaps for purposes of soliciting donations.

The next three rows report statistics for our tone classifications. Most candidates' Facebook advertising leans heavily toward the promotional category. Finally, the remaining rows report statistics for issue classifications. The most common issue areas on Facebook are Education, Economy, Fiscal Policy and Health Care.

B Machine classification of ad content

Our analyses of ad issue content and tone require a measure of these attributes that is consistently defined across media. For the TV data, WMP human coders classified every federal, gubernatorial and state legislative ad in the sample according to the 2018 WMP

codebook. The Facebook data, however, contain nearly 400,000 distinct creatives, an order of magnitude larger than the number of unique television creatives. With limited resources, a complete manual coding approach was infeasible. Instead, we implemented a supervised learning approach which uses the classifications of human coders to train a model that predicts these classifications from ad attributes. We then use the fitted model to predict content of all ads, including the "unlabeled" examples that human coders did not evaluate, in both TV and Facebook domains. We used the fitted values as our measure of content in all regressions of advertising content. We describe the method in the following subsections.

B.1 Training data

The training dataset (ads that were reviewed and classified by a human coder) contains all TV ads run by federal, gubernatorial and state legislative candidates in our sample. There are a total of 5,569 creatives in this set. In addition, we selected a random sample of Facebook ads to manually code. The randomization used in constructing the training sample blocked on state, party, and office to ensure broad coverage across these dimensions. In total, the issue content and tone of 9,073 Facebook ads were manually coded by WMP coders according to the same codebook applied to television ads. Hence the training dataset consists of 14,642 advertisements, each with a classification of tone and every issue in our issue battery.

The final issue and tone predictions we use in our regression analyses are generated from a model fit to the full training dataset. For validation and performance testing, we applied standard 5-fold cross-validation (withholding 1/5 of the training data, fitting a model on the remaining 4/5, and evaluating performance on the held-out 1/5 of examples), averaging estimates of correct classification and error rates across each fold.

B.2 Feature construction and selection

Every ad creative was run through a set of processing steps to extract relevant features on which to fit our classification models. There are four basic types of content that ads can contain: text, still images, video, and audio. Both TV and Facebook ads can and do contain all four types: TV ads often overlay text (such as a quote from a candidate or an endorser) over an image, and Facebook ads often contain embedded videos. The latter in particular is

quite common: about 35% of Facebook ads in our sample contain embedded video. To get a full picture of what a user would extract from an ad, we need to deal with all four types of data.

Video Video (from all TV ads, and the subset of Facebook ads with embedded video) was processed by 1) extracting the audio channel and passing to the audio processing step described below, and 2) sampling still frames from the video and passing to the image processing step described below. We sampled one frame at random for every 15 seconds of video, plus one frame each in the first and last two seconds and, for web videos, the display frame that shows before a user clicks play.

Audio The full audio track associated with a TV or online video was processed using Amazon's AWS Transcribe speech-to-text software. The resulting text was processed according to the text processing step described below.

Images We processed all images associated with an ad (including frames extracted from video as described above) using the Google Computer Vision API. The process extracts 1) all embedded text in an image, which was passed to the text processing step described below; 2) all human faces detected in the image, which were passed to the face processing step described below; and 3) image tags which describe the contents of an image in one or two words.

An indicator for each unique image tag that appeared in at least 0.01% of ad creatives AND in creatives associated with at least 10 distinct candidates is included as a feature in the matrix of predictors. There are 1,369 image tags that survive this check.

Faces Faces extracted from images were processed through the AWS Rekognition API. Rekognition outputs, for each face, estimates of the person's age and gender along with the image brightness and sharpness; whether the eyes and mouth are open or closed; whether the person is smiling; the presence of a beard, mustache or sunglasses; and "emotion" scores for seven attributes: CALM, HAPPY, SURPRISED, SAD, DISGUSTED, ANGRY, and CONFUSED.

We convert Rekognition's continuous scores into binary features by cutting off at thresholds. Specifically, we define binned indicators for each age bin of the set <18, 18-35, 35-50, 50-65, 65+. We construct indicators for each quantile of Sharpness and Brightness. We define a Male indicator if the gender score is greater than 0.6 and a Female indicator if the score is less than 0.4; we apply the same thresholds to create two indicators each for the Mouth Open, Eyes Open and Smile scores. Finally, the remaining scores are converted to single indicators for the score exceeding one-half.[5]

Face variables are aggregated to ad level by summing over all faces appearing in the ad. E.g. if Rekognition extracted two faces with Gender scores of 0.75 from an ad, the `face:gender_male` variable in our dataset for that ad would be equal to 2. There are a total of 29 face features in the final matrix of predictors.

Text All text associated with an ad (the concatenation of text extracted from the display text, embedded text in images, and transcribed text from the audio portion of any video) was processed by removing stopwords and stemming and then tokenizing using the `quanteda` package in `R` (Benoit et al. 2018). We included as tokens unigrams (single words) plus anything `quanteda`'s Named Entity Recognition (NER) functionality detected as a person, organization, or geographic place. This second type of token ensures that, for example, "Joe Biden" is counted as one instance of "Joe Biden" rather than one instance of "joe" and one of "biden." We again apply the frequency criteria that the token must appear in at least 0.01% of ads and in ads associated with at least 10 distinct candidates. A total of 6,683 words and 2,272 named entities survive these checks.

The final predictor matrix has a total of 10,353 features (columns) and 373,452 ads (rows), of which 14,642 have tone and issue classifications.

B.3 Classification method

We use the dropout-regularized logistic regression technique of Wager, Wang and Liang (2013) to classify the tone and issue content of the untagged ads. This method was chosen for three reasons. One, the number of unique ad creatives is very high because there are

[5] With the exception of age, the raw scores all range from 0 to 1.

many minor variations of the same ad: a candidate might experiment with changing a word or two in the headline, or altering the background color of the same image. Each of these variants will be stored as distinct creatives in Facebook's database. Dropout was designed precisely to be insensitive to small deletions of features, making it ideal for this application. Two, the Wager, Wang and Liang (2013) method makes use of information on the joint distribution of features in the untagged data to adjust the regularization penalty, which can improve performance relative to other regularization methods that use a constant penalty. Our application gives us a huge amount of untagged data to work with, maximizing the potential of this feature of dropout. Three, the use of a penalized logistic regression, unlike more complicated "deep learning" methods, gives easily interpretable coefficients that can be inspected and checked for logical validity. In cross-validation tests, dropout consistently outperformed other common logistic-regression-based methods like ridge, lasso or elastic net.

The final models used to produce tone and issue prediction use tuning parameters of $p = 0.5$ (dropout probability), $a = 0.1$ (weight on untagged data), and a small ridge penalty with $\lambda = 0.01$. These were selected by five-fold cross-validation on the negative tone outcome. We estimate one model per issue or tone category; hence each is a binary classification.

B.4 Model fit and error rates

Using a five-fold cross validation procedure, we evaluated the model's prediction accuracy and error rates. Results are displayed in Figure B.1. Correct classification rates are extremely high across the board: the worst-performing model is the "Contrast" tone model, where out-of-sample predictions are correct a little more than 80% of the time. The large majority of models achieve out-of-sample correct classification rates of 95% or more.

However, this statistic is somewhat misleading here because many of the issues in question are very rare: for example, the "Welfare" issue occurs in less than 1% of ads in the training data. Thus, even a constant model (predict 0 for every ad) can achieve very high correct classification rates on these issue categories.

More informative are the false negative and false positive rates. These are computed as, respectively, the fraction of model predictions of 0 (1) where the human classification is 1 (0). False negative rates are low everywhere, indicating the vast majority of the time

that the model says an ad is not, for example, an attack ad, human coders agree with this classification. False positive rates - the more difficult criterion given the rarity of the tags - are higher but still generally below 0.25, particularly for our composite issue categories (displayed at the top of the figure). Performance degrades somewhat for the more detailed individual categories: e.g. the "Law and Order" composite issue tag has false positive rate of about 0.2 whereas the "Incarceration / Sentencing" detail issue tag which it contains has false positive rate closer to 0.7. We focus in our analyses on the composite issue areas and the single issue categories (such as "Gun control / guns" that are sufficiently frequent in the training data to yield reasonably accurate predictions.

For our prediction models of party and CFScore, we show measures of model fit in Figures B.2 and B.3, respectively. Figure B.2 shows the ad-level distribution of ads by party of sponsor. There is evident separation between parties, though the separation is substantially greater for ads on Facebook than on television. Figure B.3 plots predicted against actual CFscore of the ad sponsor, by ad, and overlays the regression line to show the relationship between the two. The overall correlation is 0.8. Again, the fit is noticeably better on Facebook ads compared to television ads.

C Additional regression results

In this section we show coefficient estimates for interaction terms of the main variable of interest - $Facebook_{ik}$ - in our regression specifications (4) and (5). We cannot reject the null that these are zero across the board, and thus focus discussion in the main text on the main effects. We present these here for consistency with the pre-analysis plan.

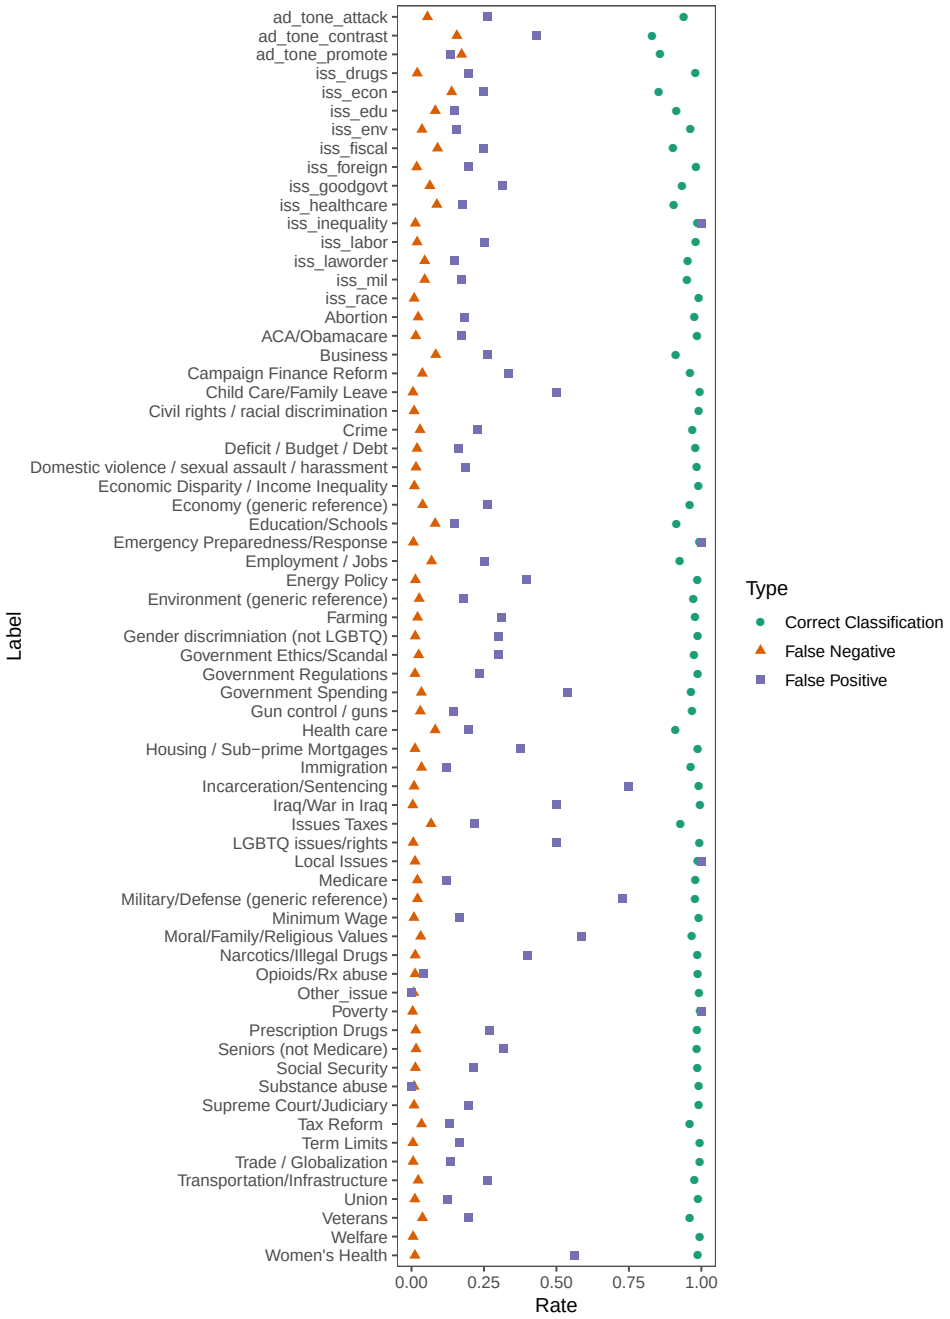

Figure B.1: Prediction accuracy and error rates for the dropout-regularized logistic classifier, by outcome. Rates are estimated by averaging over five cross-validation folds.

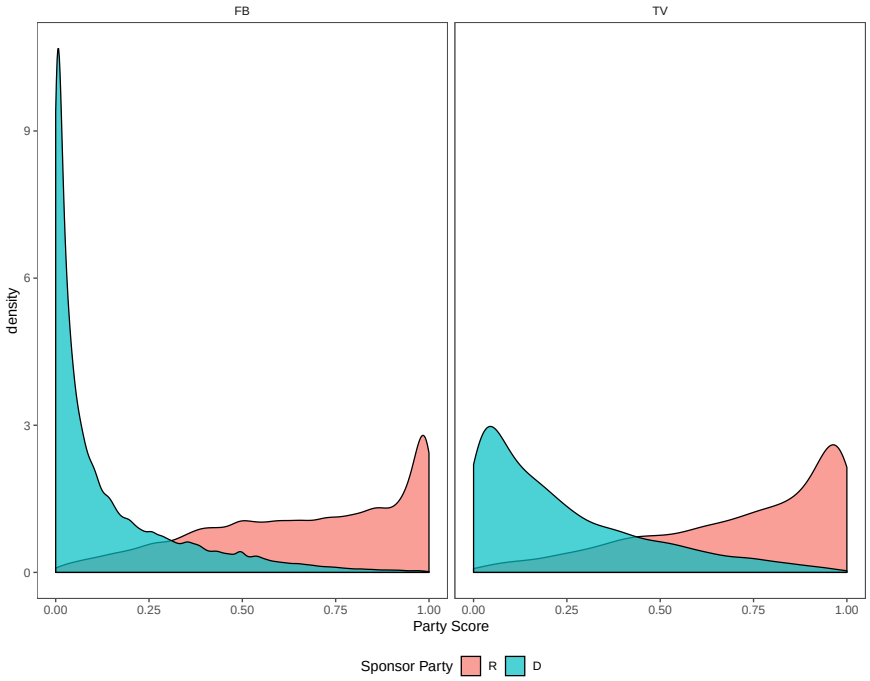

Figure B.2: Density of party score predictions, by party. The red curve is the distribution of ads run by Republican candidates; the blue curve is the distribution of ads run by Democratic candidates. Left panel is ads on Facebook; right panel is ads on TV.

Label	High Kappa?	Composite Issue
Substance abuse	N	Drugs
Narcotics/Illegal Drugs	N	Drugs
Prescription Drugs	Y	Healthcare
Opioids/Rx abuse	Y	Drugs
Marijuana	N	Drugs
Drugs-Issues Tobacco	N	Drugs
Economy (generic reference)	Y	Economy
Issues Taxes	Y	Fiscal
Tax Reform	Y	NA
Deficit / Budget / Debt	Y	Fiscal
Government Spending	N	Fiscal
Recession / Economic Stimulus	N	Economy
Trade / Globalization	N	Economy
Employment / Jobs	Y	Economy
Business	Y	Economy
Union	N	Labor
Minimum Wage	N	Labor
Economic Disparity / Income Inequality	N	Inequality
Poverty	N	Inequality
Farming	N	NA
Housing / Sub-prime Mortgages	N	NA
Education/Schools	Y	Education
Lottery for Education	Y	Education
Child Care/Family Leave	N	NA
Environment (generic reference)	N	Environment
Climate Change / Global Warming	Y	Environment
Energy Policy	N	Environment
Keystone XL Pipeline	N	Environment
Campaign Finance Reform	N	Good Government
Government Ethics/Scandal	N	Good Government
Corporate Fraud	N	Good Government
Military/Defense (generic reference)	N	Military
Foreign Policy (generic reference)	N	Foreign Policy
Veterans	Y	Military
Foreign Aid	N	Foreign Policy
Nuclear Proliferation	Y	Foreign Policy
September 11th	Y	Foreign Policy
Terror/Terrorism/Terrorist	Y	Foreign Policy
Middle East	N	Foreign Policy
Afghanistan/War in Afghanistan	Y	Foreign Policy
Iraq/War in Iraq	Y	Foreign Policy
Israel	N	Foreign Policy
Iran	Y	Foreign Policy
ISIL/ISIS	Y	Foreign Policy
Syria	N	Foreign Policy
Russia / Putin	N	Foreign Policy
North Korea / Kim Jong Un	Y	Foreign Policy
China	Y	Foreign Policy
Health care	Y	Healthcare
ACA/Obamacare	Y	Healthcare
Women's Health	N	Healthcare
Medicare	Y	Healthcare
Crime	N	Law & Order
Incarceration/Sentencing	N	Law & Order
Supreme Court/Judiciary	Y	Law & Order
Capital Punishment	N	Law & Order
Police brutality / racial violence	N	Law & Order
Domestic violence / sexual assault / harassment	Y	Law & Order
Immigration	Y	NA
Abortion	Y	NA
Moral/Family/Religious Values	N	NA
Gun control / guns	Y	NA
Seniors (not Medicare)	Y	NA
Social Security	Y	NA
Welfare	N	NA
LGBTQ issues/rights	Y	NA
Gender discrimniation (not LGBTQ)	N	NA
Civil Liberties/Privacy	N	NA
Civil rights / racial discrimination	N	Race
Affirmative Action	N	Race
Gambling	N	NA
Assisted Suicide/Euthanasia	N	NA
Term Limits	Y	Good Government
Pledge of Allegiance (restrictions on)	N	NA
Local Issues	N	NA
Government Regulations	N	NA
Government Shutdown	N	NA
Emergency Preparedness/Response	Y	NA
Transportation/Infrastructure	Y	NA
Other_issue	N	NA

Table A3: Issue tags in the WMP data. Column 1 is the underlying issue tag. Column 2 indicates whether human coders had sufficiently high inter-coder reliability for inclusion in the issue-by-issue analyses. Column 3 is the composite issue to which the detail issue is assigned, if any. Composite issues are included in our issue-by-issue analyses if at least one of their sub-issues has high enough inter-coder reliability.

Content	Min	Mean	Max
Fraction Impressions In-State	0.00	0.94	1.00
Tone: Attack	0.00	0.03	0.66
Tone: Promote	0.00	0.88	1.00
Tone: Contrast	0.00	0.09	0.88
Tax Reform	0.00	0.01	1.00
Immigration	0.00	0.02	1.00
Abortion	0.00	0.02	0.99
Gun control / guns	0.00	0.03	1.00
Seniors (not Medicare)	0.00	0.01	0.91
Social Security	0.00	0.01	0.41
LGBTQ issues/rights	0.00	0.00	0.47
Emergency Preparedness/Response	0.00	0.00	0.89
Transportation/Infrastructure	0.00	0.02	0.96
Drugs	0.00	0.01	0.63
Fiscal	0.00	0.10	0.99
Economy	0.00	0.15	1.00
Military	0.00	0.04	1.00
Education	0.00	0.16	1.00
Law & Order	0.00	0.04	1.00
Foreign Policy	0.00	0.01	0.73
Health Care	0.00	0.09	1.00
Environment	0.00	0.04	0.95
Good Government	0.00	0.04	0.85

Table A4: Summary statistics of Facebook advertising content, at candidate level.

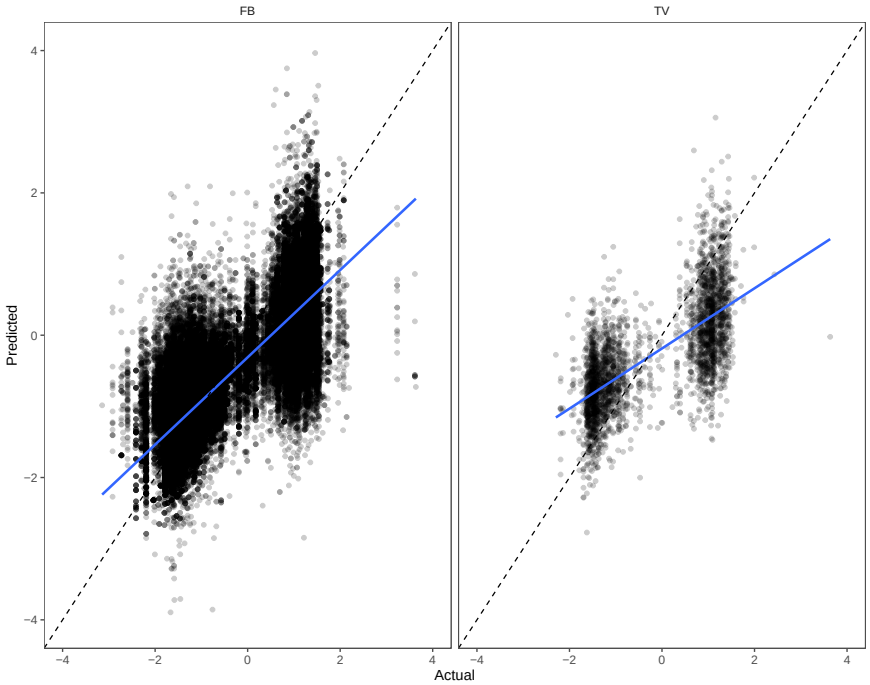

Figure B.3: Predicted versus actual CFScore, by advertisement. Each point is an individual ad creative; the horizontal axis shows the actual CFScore of the sponsor, and the vertical axis shows the predicted value from our model. The dashed black line is the 45-degree line; the solid blue line is the OLS fit. Left panel is ads on Facebook; right panel is ads on TV.

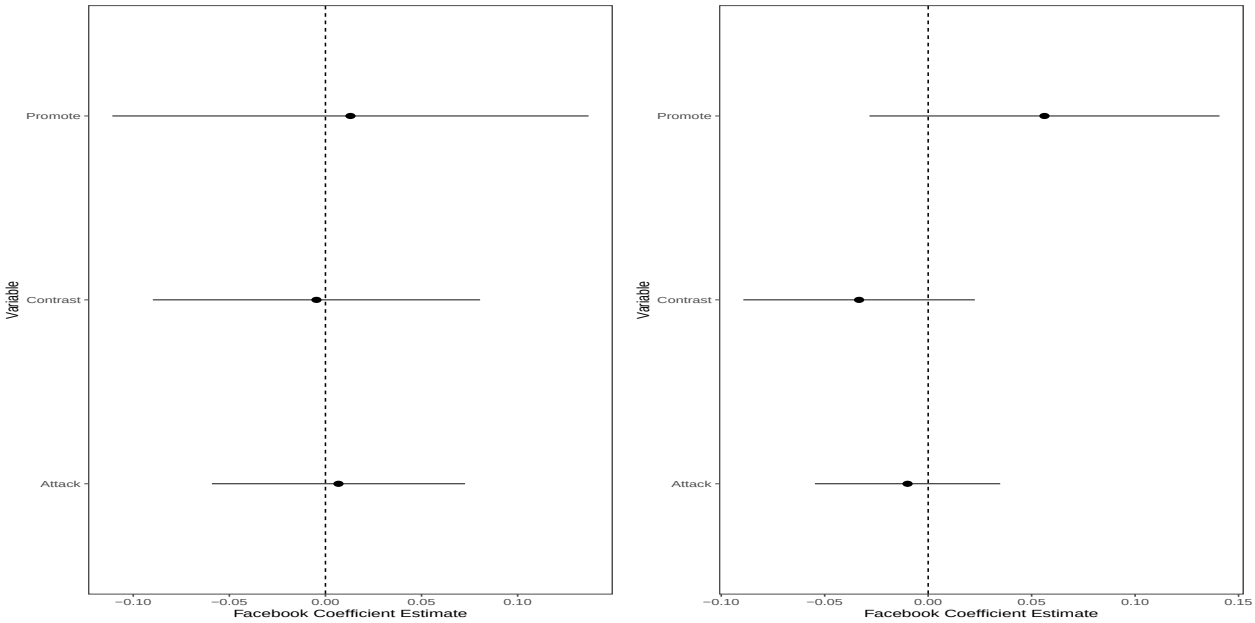

(a) Senate

(b) House

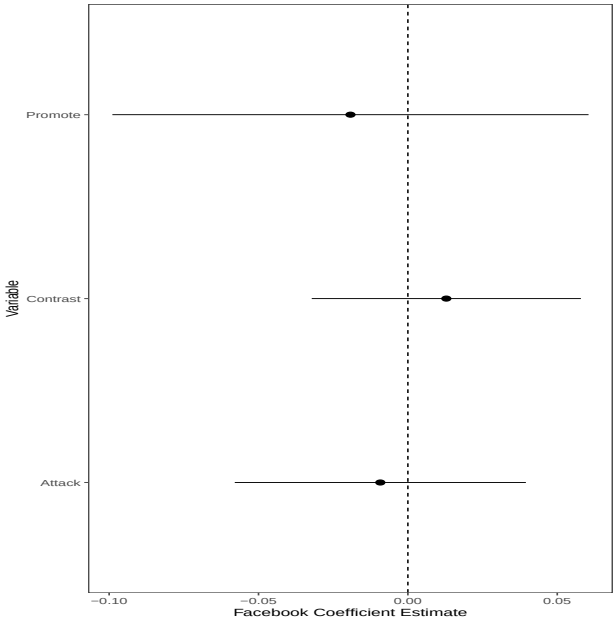

(c) Incumbent

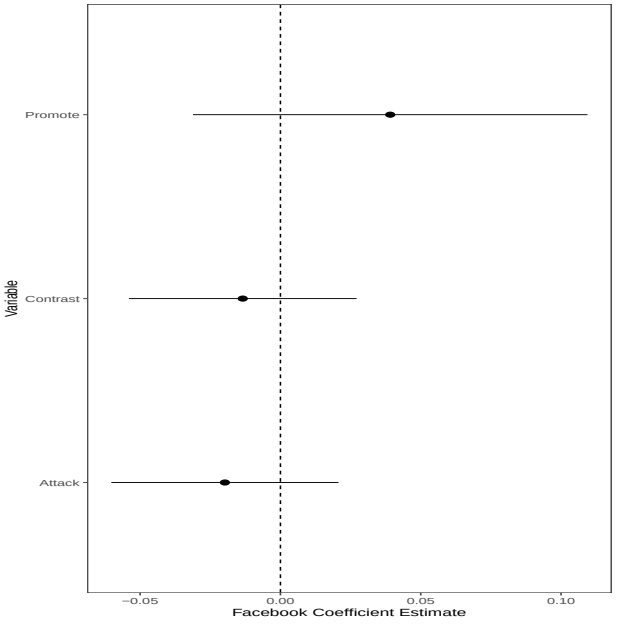

(d) Republican

Figure C.1: Interaction terms with Facebook indicator in Tone models. Plots display coefficient estimate and 95% asymptotic confidence interval (using standard errors clustered at candidate level) for interactions of the specified candidate attribute with the Facebook indicator. All interacted variables are binary indicators; e.g. in panel a) the displayed coefficient is the interaction of an indicator for the candidate running for the US Senate with the indicator for Facebook ads.

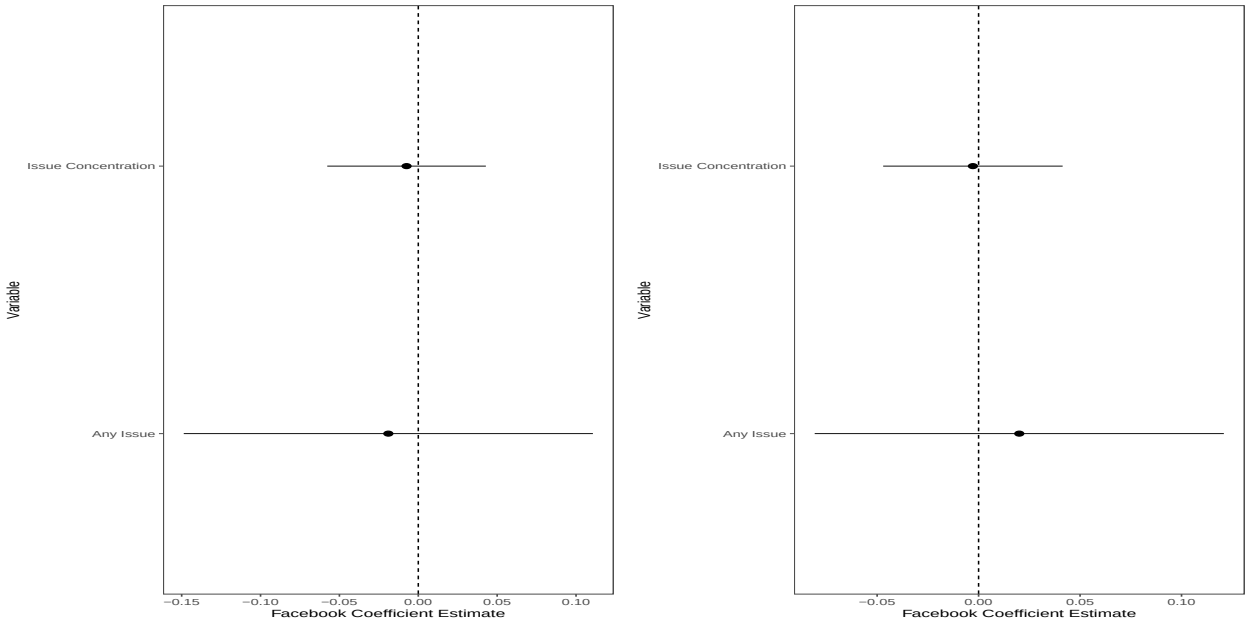

(a) Senate

(b) House

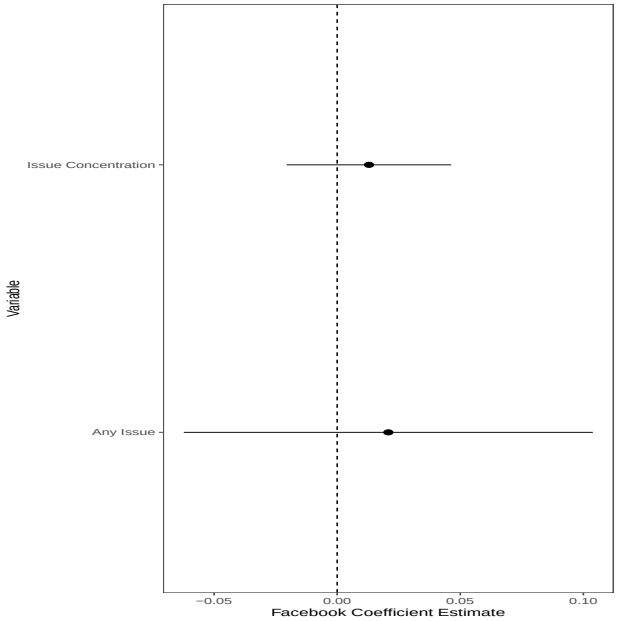

(c) Incumbent

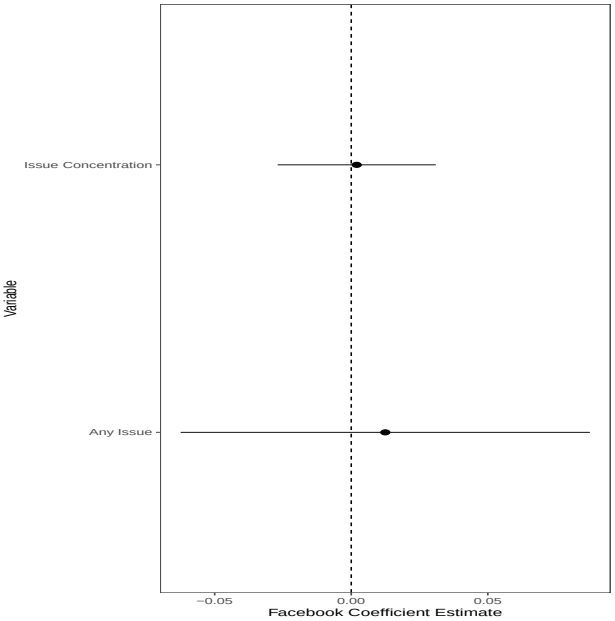

(d) Republican

Figure C.2: Interaction terms with Facebook indicator in Issue Diversity models. Plots display coefficient estimate and 95% asymptotic confidence interval (using standard errors clustered at candidate level) for interactions of the specified candidate attribute with the Facebook indicator. All interacted variables are binary indicators; e.g. in panel a) the displayed coefficient is the interaction of an indicator for the candidate running for the US Senate with the indicator for Facebook ads.