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.

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

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

[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.

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

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

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.

Recall that because the API data list spending per creative in bins, we used the midpoint of the bins

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

FIGURE A1

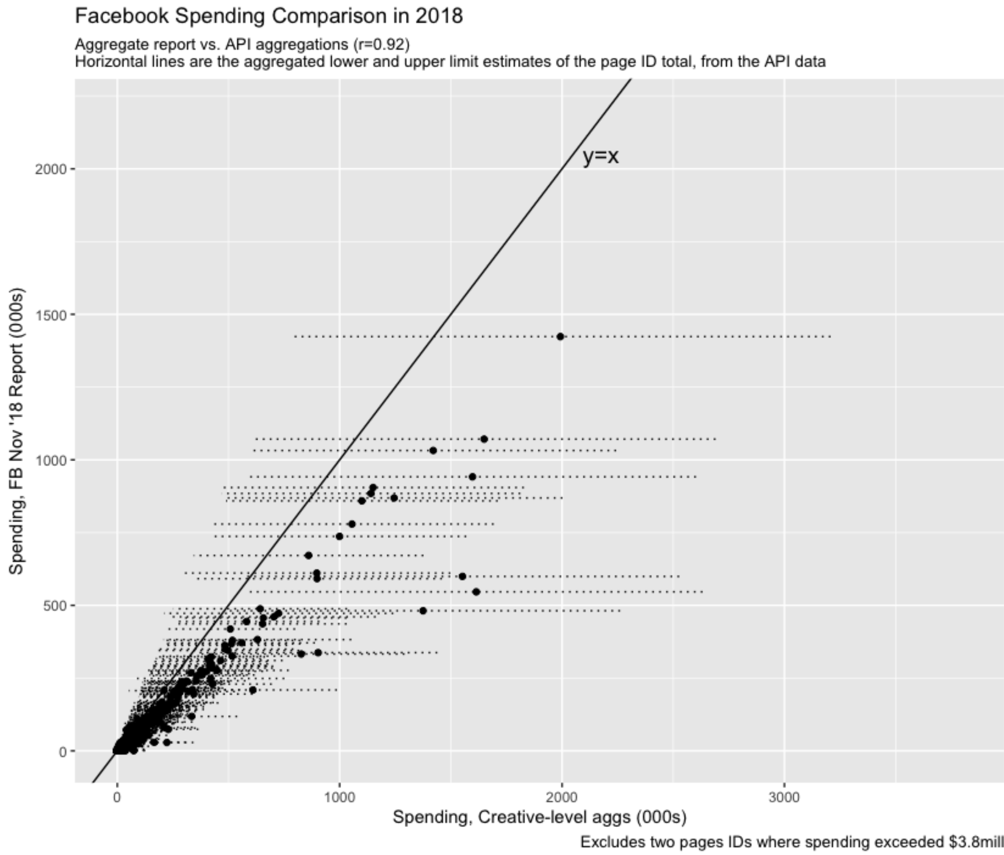

to estimate the creative-level spending. But sponsors may have paid on the upper or lower end of that bin, which is only an issue if that tended to happen systematically above or below the midpoint.

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.

A.6. Ad quantity summary statistics

Figures A2b and A2a show the density of candidate-level spending, by office sought, on television and Facebook, respectively. These plots condition on non-zero spending on the medium; e.g., a campaign is included in A2b only if it had strictly positive spending on TV.

FIGURE A2. Density of candidate-level spending on each medium, by office.

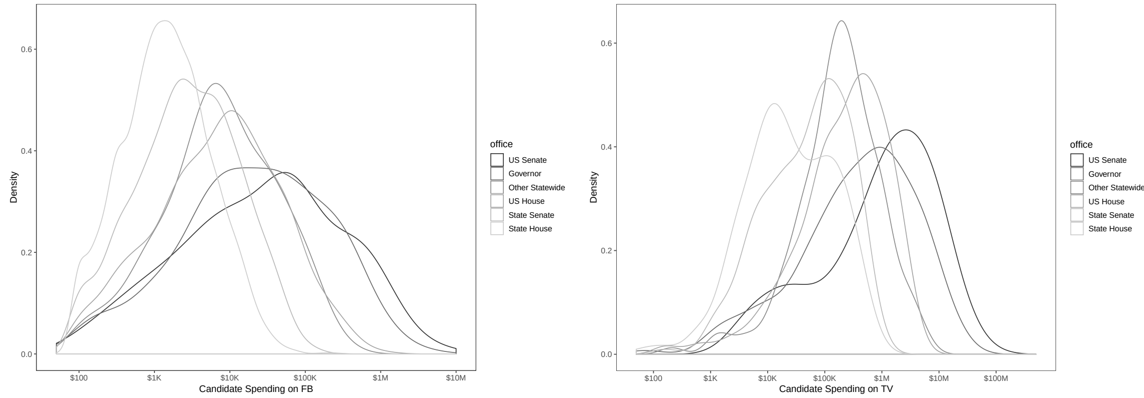

(a) Facebook **(b) TV**

Note: Plots condition on non-zero spending on the indicated medium (i.e., they exclude the point masses at zero).

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.

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

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 et al. (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 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 et al. (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

FIGURE B.1. Prediction accuracy and error rates for the dropout-regularized logistic classifier, by outcome.

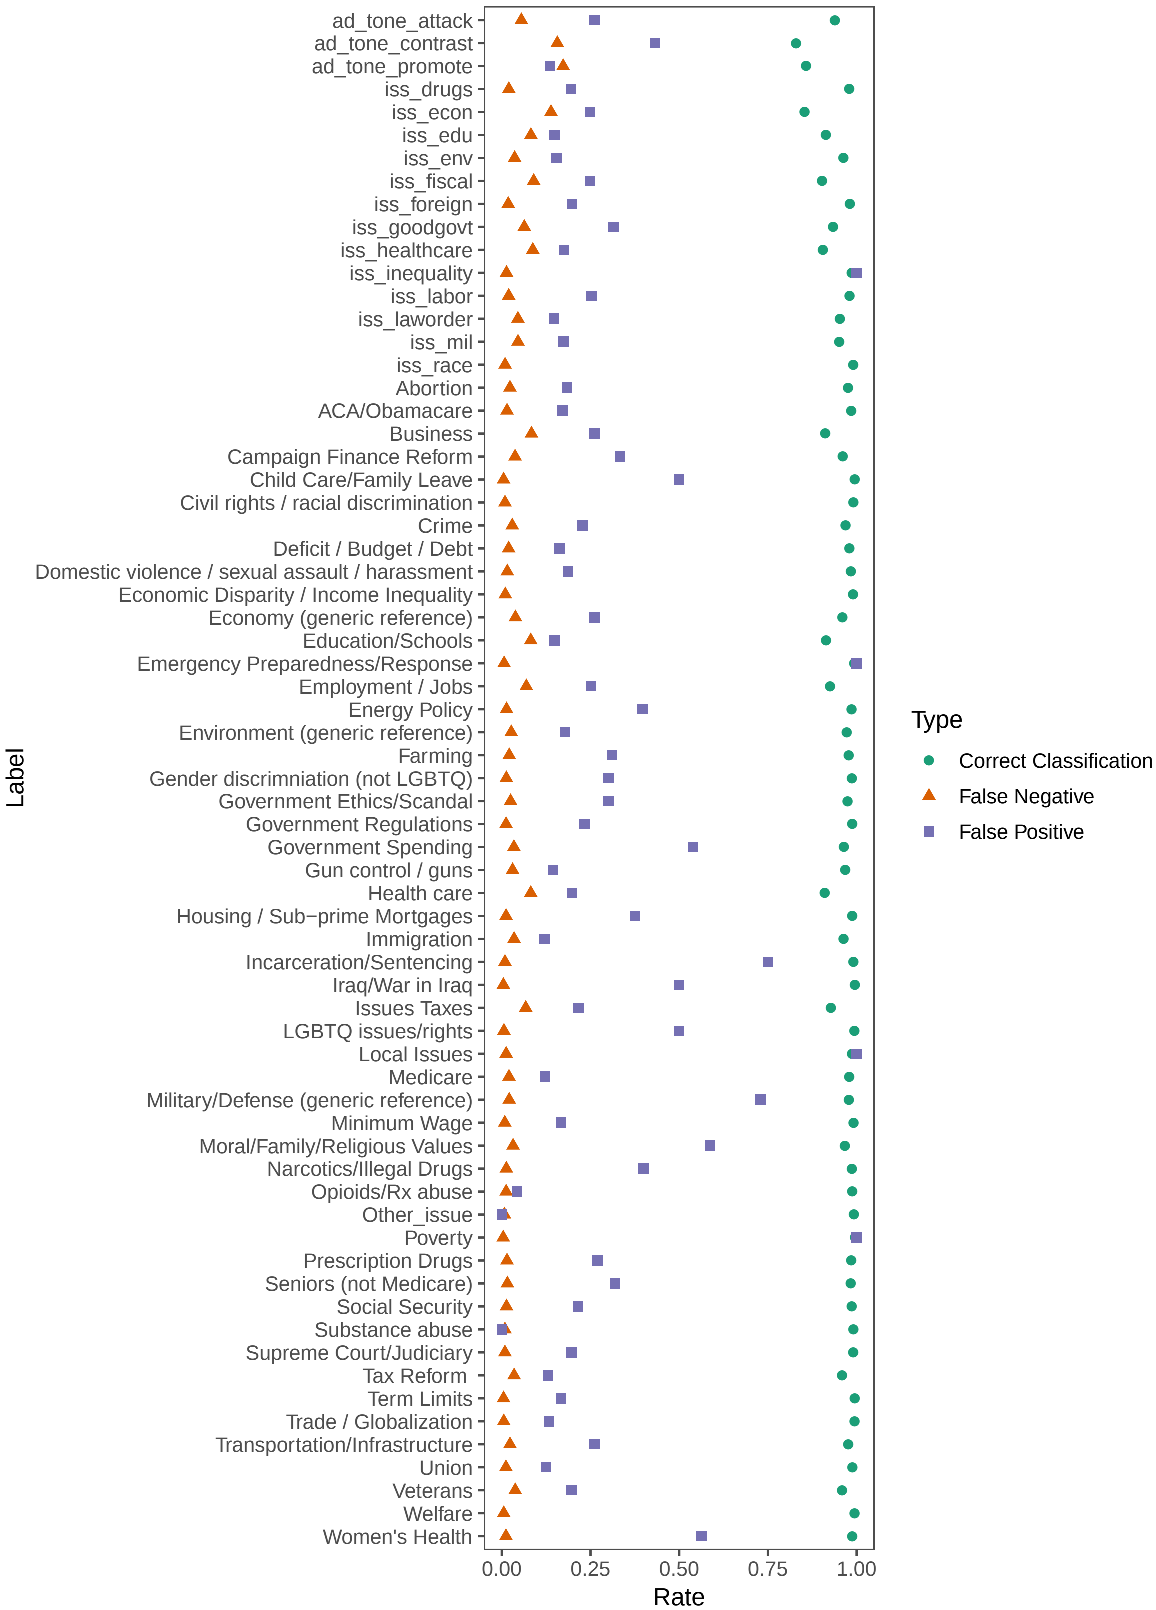

Note: Rates are estimated by averaging over five cross-validation folds.

FIGURE B.2. Density of party score predictions, by party.

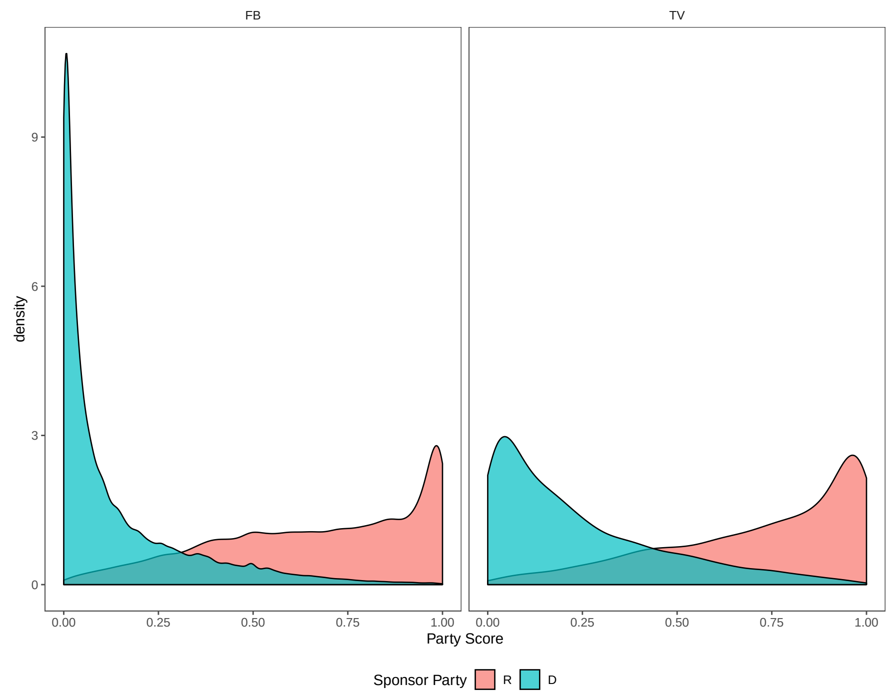

Note: The red curve is the distribution of ads run by Republican candidates; the blue curve is the distribution of ads run by Democratic candidates. Left panel is ads on Facebook; right panel is ads on TV.

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.

B.5. Main results using only the human-coded subsample

As a robustness check, we estimate our primary models using only the human-coded subsample of ad creatives. Results of this exercise are shown in Figure B.4. The point estimates for the tone, single issue, any issue, and party / ideological extremity models are similar to our main specifications, though confidence intervals, unsurprisingly, widen.

Estimates of the Facebook effect in models of issue diversity and variance in ideological positioning differ from the main specification. This is because our human-coded sample covers the universe of TV ads but only a small fraction of Facebook ads. Hence, many candidates have little or no observed

FIGURE B.3. Predicted versus actual CFScore, by advertisement.

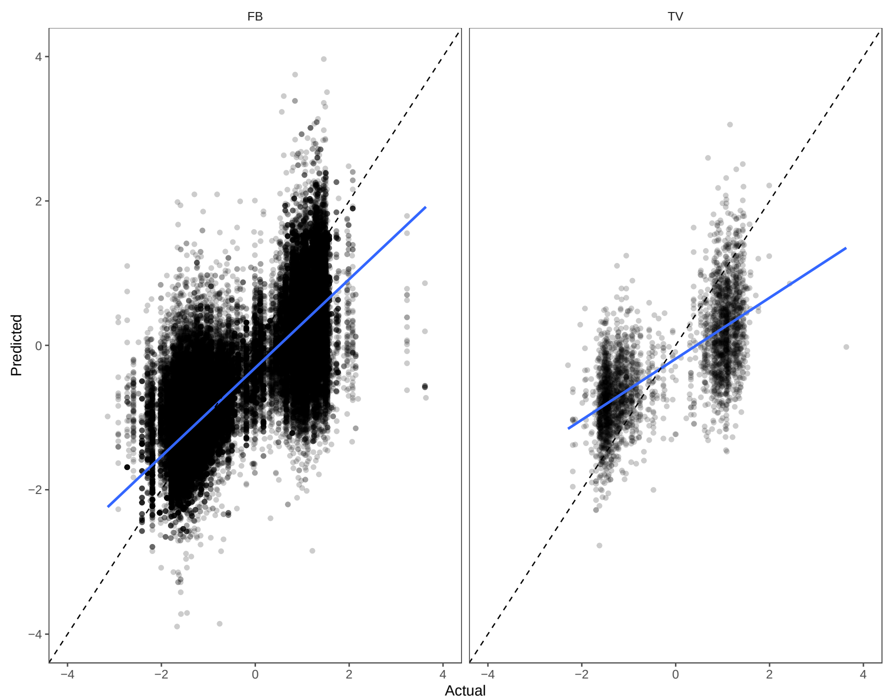

Note: Each point is an individual ad creative; the horizontal axis shows the actual CFScore of the sponsor, and the vertical axis shows the predicted value from our model. The dashed black line is the 45-degree line; the solid blue line is the OLS fit. Left panel is ads on Facebook; right panel is ads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

FIGURE B.4. Models of issue content, tone, and partisanship on medium estimated on the subsample of human-coded ad creatives.

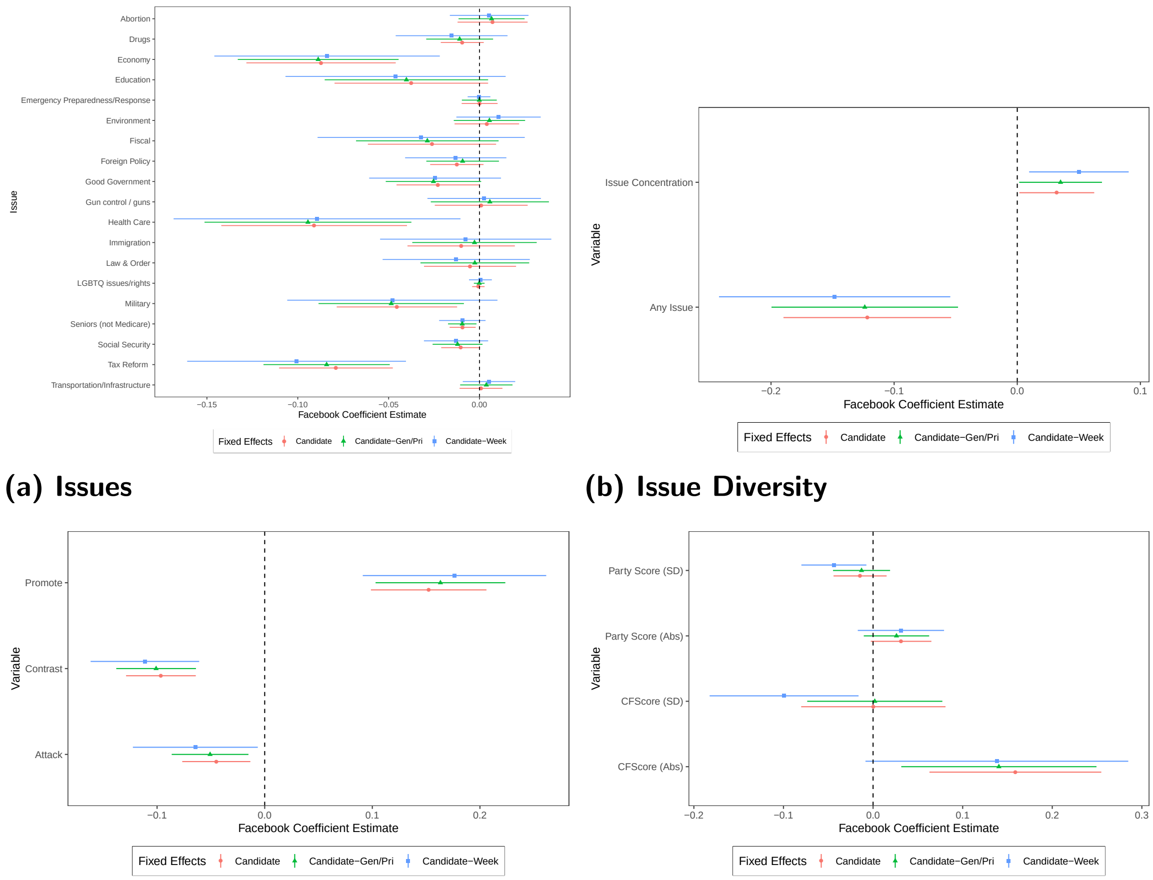

(a) Issues **(b) Issue Diversity**

(c) Tone **(d) Party / Ideology**

Note: Plots show the estimated coefficient on the Facebook dummy, in models with candidate (or candidate-week or candidate-election) fixed effects. These models use only the human-coded ad creatives to compute candidate-level averages.

variance in ideological positioning or in issues covered on Facebook, merely because most of their Facebook ads are missing in the human-coded subsample. These models illustrate the utility of the machine coding approach, which allows us to get a fuller picture of advertising on Facebook than would be feasible with human coders alone.

B.6. Estimating measurement error influence

In addition to estimates which restrict to human-coded ads in the previous section, we also conducted a resampling exercise to estimate the degree of potential influence of misclassification on our regression results. This exercise involved randomly perturbing our machine-classified data set according to the estimated item-specific false positive and false negative rates presented in Figure B.1. Specifically, we reversed[6] the predicted score for a fraction p^c_{FP} of ads predicted to contain the characteristic c and p^c_{FN} of ads predicted not to contain the characteristic c.[7] We then aggregated to candidate × medium level, weighting by expenditure, and recomputed our fixed-effects models. We repeated this process 500 times and recorded the standard deviation of the resulting coefficient estimates on the Facebook dummy in each model.

Table B.1 shows that the influence of measurement error on the variance of our estimates is small - at least an order of magnitude smaller than the sampling variation we report in the coefficient plots for our main specifications. Our aggregation to candidate level is helpful here, as even if some ads are misclassified, their contribution to candidate-level averages is not too large to have serious influence on our estimates. We conclude that our classification accuracy is sufficient for our purposes of estimating candidate-level regressions. More caution, and additional analysis, would be appropriate before using our classifications in ad-level regressions.

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 (5) and (6). 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.

[6]By replacing the score x with $1 - x$.

[7]Because our predicted scores are continuous and range from 0 to 1, we consider any ad with predicted score greater than 0.5 to contain the characteristic and any ad with predicted score less than 0.5 to not contain it.

FIGURE C.1. Interaction terms with Facebook indicator in Tone models.

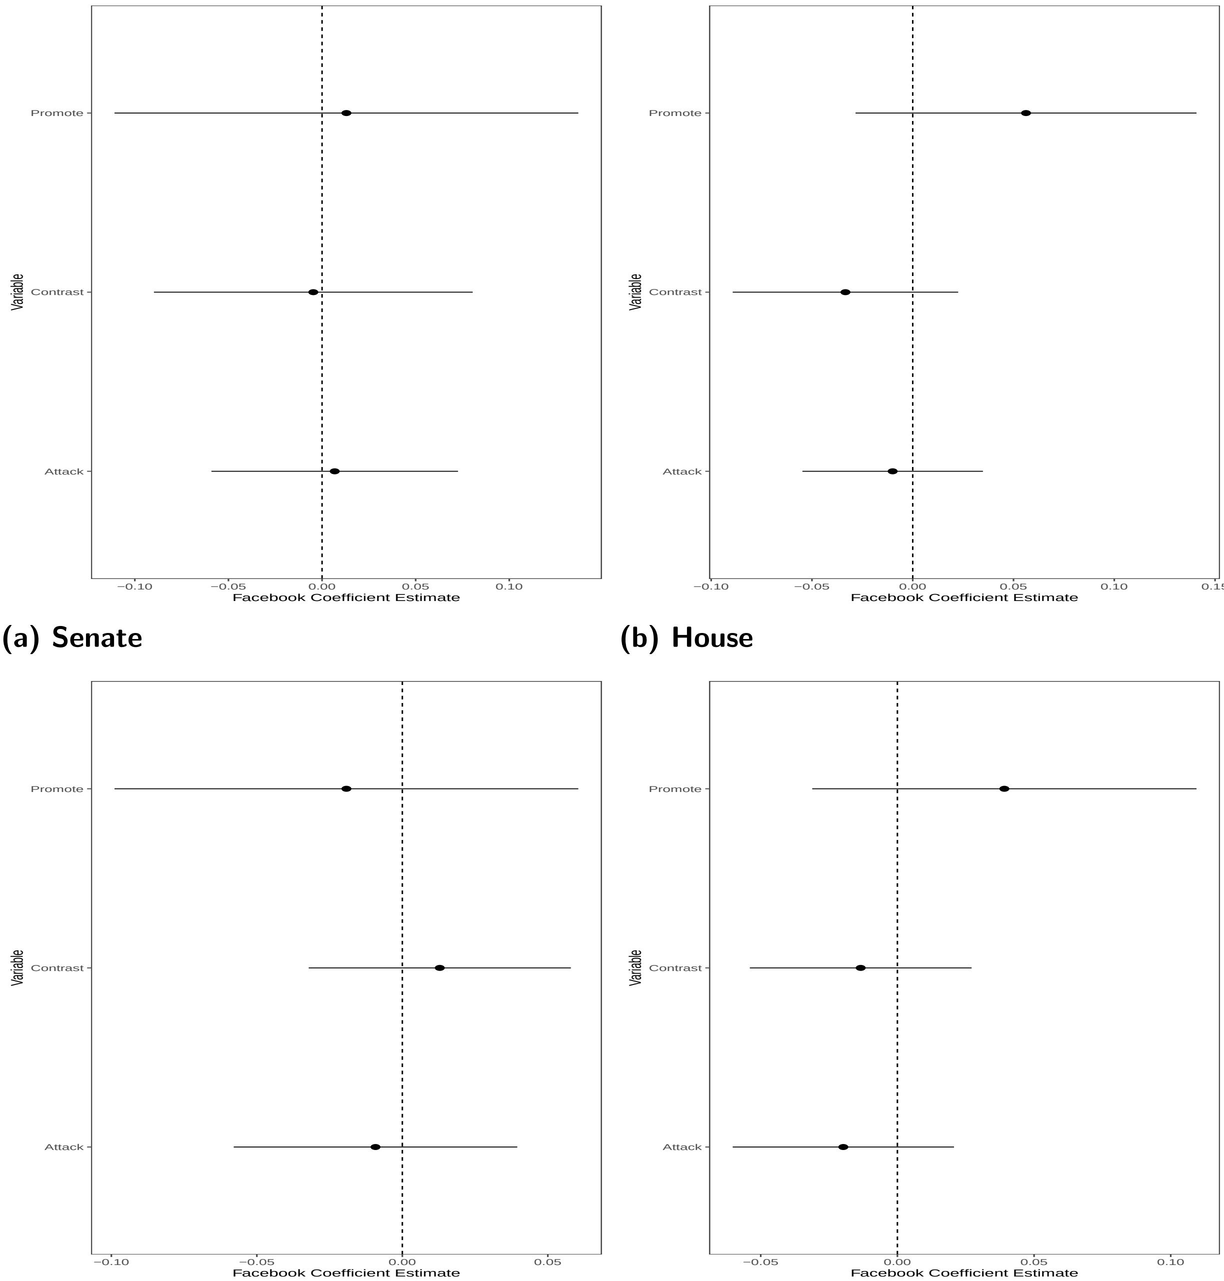

(a) Senate **(b) House**

(c) Incumbent **(d) Republican**

Note: 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.

FIGURE C.2. Interaction terms with Facebook indicator in Issue Diversity models.

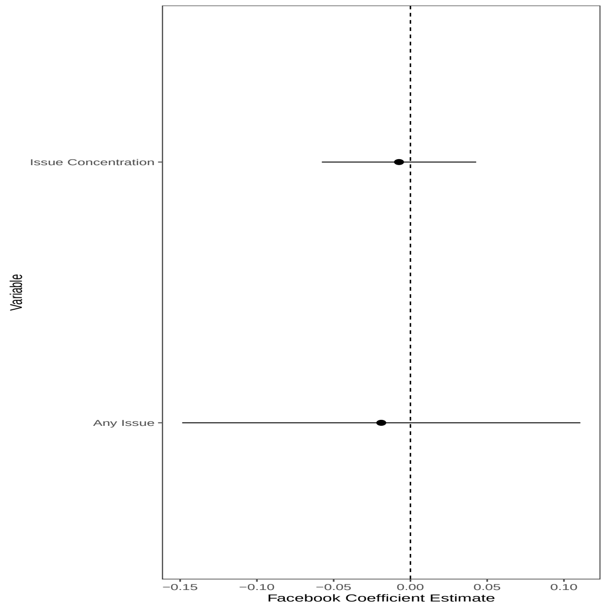

(a) US Senate

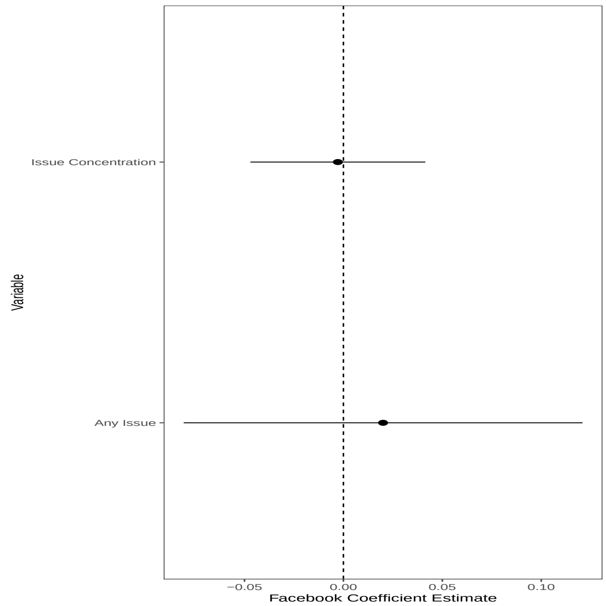

(b) US House

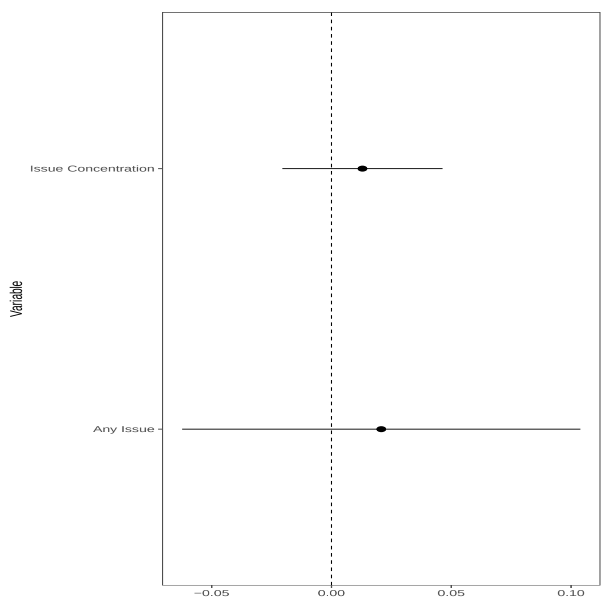

(c) Incumbent

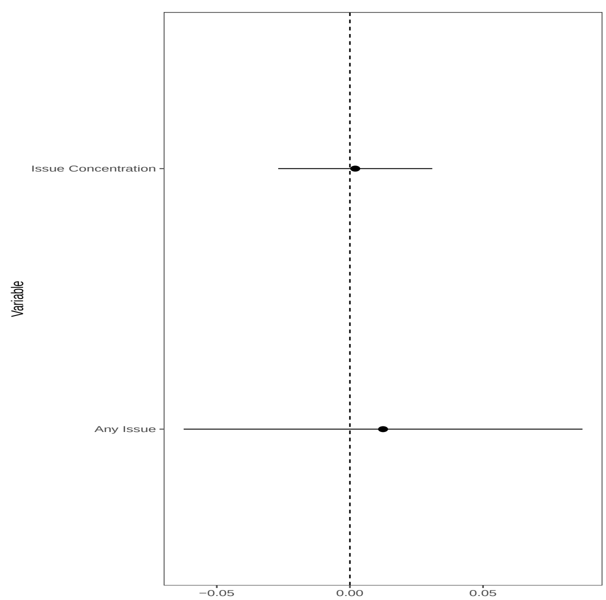

(d) Republican

Note: 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.

TABLE A3. Issue tags in the WMP data.

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

Note: 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.

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

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 B.1. Standard deviation of point estimates of Facebook coefficient in resampling exercise.

Label	Candidate	Candidate-Gen/Pri	Candidate-Week
Attack	0.003	0.003	0.002
Contrast	0.004	0.004	0.003
Promote	0.003	0.003	0.002
Any Issue	0.004	0.004	0.003
Issue Concentration	0.001	0.001	0.001
Abortion	0.002	0.002	0.001
Drugs	0.002	0.002	0.001
Economy	0.003	0.003	0.003
Education	0.003	0.003	0.002
Emergency Preparedness/Response	0.001	0.001	0.001
Environment	0.002	0.002	0.002
Fiscal	0.003	0.003	0.002
Foreign Policy	0.002	0.002	0.001
Good Government	0.003	0.003	0.002
Gun control / guns	0.002	0.002	0.002
Health Care	0.003	0.003	0.002
Immigration	0.002	0.002	0.002
Law & Order	0.003	0.003	0.002
LGBTQ issues/rights	0.001	0.001	0.001
Military	0.003	0.003	0.002
Seniors (not Medicare)	0.002	0.002	0.001
Social Security	0.001	0.001	0.001
Tax Reform	0.002	0.002	0.002
Transportation/Infrastructure	0.002	0.002	0.001

Note: See text for details. Columns indicate the level at which fixed effects are included in the regression model.