Social Science Computer Review
2024, Vol. 42(5) 1181–1203

DOI: 10.1177/08944393241258767
journals.sagepub.com/home/ssc
Sage

Platform Convergence or Divergence? Comparing Political Ad Content Across Digital and Social Media Platforms

Travis N. Ridout[1], Markus Neumann[2], Jielu Yao[3], Laura Baum[3], Michael M. Franz[4], Pavel Oleinikov[5], and Erika Franklin Fowler[5]

Abstract

When it comes to the study of the messaging of online political campaigns, theory suggests that platform divergence should be common, but much research finds considerable convergence across platforms. In this research, we examine variation across digital and social media platforms in the types of paid campaign messages that are distributed, focusing on their goals, tone, and the partisanship of political rhetoric. We use data on the content of paid election advertisements placed on YouTube, Google search, Instagram, and Facebook during the 2020 elections in the United States, examining all federal candidates who advertised on these platforms during the final 2 months of the campaign. We find that YouTube is most distinct from the other platforms, perhaps because it most resembles television, but convergence better describes the two Meta platforms, Facebook and Instagram.

Keywords

digital platforms, political advertising, campaigns, negativity

The use of digital advertising in political campaigns has skyrocketed in recent years. For instance, during the 2020 presidential general election in the United States, an estimated $435 million was spent on digital advertising, which is about a quarter of all ad spending in that race (Ridout et al., 2021a). Another estimate, based on Federal Election Commission filings, puts the percentage of

[1]Washington State University, USA
[2]Duke Kunshan University, China
[3]Wesleyan Media Project, USA
[4]Bowdoin College, USA
[5]Wesleyan University, USA

Corresponding Author:
Travis N. Ridout, School of Politics, Philosophy and Public Affairs, Washington State University, 816 Johnson Tower, Troy Lane, Pullman, WA 99164-4880, USA.
Email: tnridout@wsu.edu

digital spending even higher, at 33% of total media expenditures for the Biden campaign and 40% for the Trump campaign (Gulati & Williams, 2021). These percentages are sharply higher than just 8 years earlier in 2012 when both the Obama and Romney campaigns spent under 15% of their media expenditures on digital (Gulati & Williams, 2021).

Clearly, digital advertising is an increasingly important way by which campaigns disseminate their messages to voters, and yet there remains much that we do not know about how paid online advertising is used in today's political campaigns. For one, as Kruschinski et al., (2022) argue, there is "a notable paucity of empirical studies investigating the use and content strategy of paid messages" (p. 2). Key exceptions are a study of Facebook ads in the 2019 Canadian federal elections (Bennett & Gordon, 2021) and a study of Facebook ad content in a variety of federal and local races in the United States in 2018 (Fowler et al., 2021b). Second, those studies we do have are generally focused on a single platform, using that platform to stand in for digital advertising as a whole. Matassi and Boczkowski (2023), writing about the study of digital platforms more broadly, note that "most scholarship has tended to concentrate on patterns related to a single platform at a time—usually Facebook or Twitter" (p. 3). This is potentially a problem for coming to a full understanding of digital political advertising because, as Bode and Vraga (2018) suggest, research that examines a single platform "fails to recognize the diverse, fragmented, and complex modern media environment, of which a single media platform is only a small part" (p. 1). The upshot, then, is the need to engage in broader analysis of the content features of digital advertising and to do so across digital and social media platforms.

In recognition of this problem, there have been some recent efforts to study digital campaigning across platforms, though none of this research has looked at paid political advertising specifically. This research has often adopted a convergence–divergence framework, focusing on the extent to which messaging is similar across different platforms. In short, while there is good theory to support the idea of *platform divergence*—that campaigns craft different and distinct messages for different digital and social media platforms (e.g., Bossetta, 2018; Kreiss et al., 2018)—the limited empirical research that exists tends to support the idea of *platform convergence*—that campaigns are putting out similar messages on varied platforms (e.g., Bossetta & Schmøkel, 2023; Boulianne & Larsson, 2023).

Our research is unique in that it is the first to focus on digital messaging through the lens of online political advertising, which has become increasingly important in the United States and around the globe (Dommett, 2019). Second, we focus on three content features of that advertising that are especially important for understanding its impact: the use of partisan language, the tone of the advertising, and the goals pursued by the ads.

Whether ad content across platforms converges or diverges matters. First, it is important to know whether different segments of the electorate, defined by their social media use, are receiving different campaign messages. If campaign ads are hampering a shared public understanding of what the campaign was about, then this could create public disagreement about what issues or concerns elected officials should focus on after the election, making governing difficult. Second, this study has implications for research on digital politics, as it speaks to how far we can generalize from single-platform studies, which are still the norm. Do conclusions about candidate behavior drawn from a study of Facebook ads teach us about online advertising as a whole or only about advertising on Facebook? Finally, given that many theories of digital campaigning, such as those based on different audiences (Kreiss et al., 2018) or digital architectures (Bossetta, 2018), expect to find cross-platform differences, a finding that differences are minimal could serve as a call for scholars to rethink the application of those theories to paid political messages.

In this research, we examine variation in political ads across four platforms (Facebook, Instagram, YouTube, and Google search[1]), focusing on the use of partisan language, the tone of advertising, and the specific goals pursued. Our study contributes to the study of digital campaigning in several important ways. First, we examine four different platforms, whereas existing cross-platform research analyzes three or two. In a couple of cases (Bossetta & Schmøkel, 2023; Farkas & Bene, 2021), the comparison is between Instagram and Facebook, two platforms owned by the same company (Meta) and whose dashboards easily allow for cross-posting. These studies, then, stack the deck in favor of finding platform convergence. Second, while past research has examined campaign posts, we examine paid advertising on these platforms, which given its cost, presumably is reflective of careful campaign decision-making. Third, by using machine learning methods, we can process a large number of ads—just over 145,000—allowing us to speak with more confidence about platform differences.

All told, we find that platform convergence or divergence depends on the specific platforms and content features (language, tone, and goals) examined. The best evidence for platform divergence comes from YouTube, which features different content and goals than the other platforms. But the distinctiveness of YouTube is not just a function of its video format, as video and non-video were quite similar on Instagram and Facebook.

Comparing Political Messaging Across Platforms

Although comparisons across online platforms are not common in the study of campaign communications—and we were unable to find any that looked at paid campaign advertising—there are a few notable studies.

For instance, Bossetta and Schmøkel (2023) compare images in Facebook and Instagram posts from several U.S. presidential candidates in 2020. They find that about half of the images they examined were posted on both platforms, and there is very little variation in the emotions conveyed by the images of candidates' faces. All in all, this study provides strong support for the platform convergence hypothesis, but in some ways, the deck is stacked in favor of this outcome given that Facebook and Instagram are both owned by Meta, where the dashboard that campaigns use easily allows for cross-posting.

Another study, from Boulianne and Larsson (2023), examines leader posts during the 2019 federal election in Canada, focusing on Facebook, Instagram, and Twitter. The study finds that some types of posts resonated more on certain platforms—in the sense of driving user engagement metrics—but when it comes to the specific content posted, the authors conclude: "most of the content was cross-posted without changes to reflect the platform and/or what resonates on specific platforms" (p. 134). These findings, then, strike us as more consistent with the idea of platform convergence than divergence.

Differences between platforms are slightly more pronounced in a study by Stromer-Galley and colleagues (2021), which among other things compares how campaigns used Facebook and Twitter in the 2016 and 2020 U.S. presidential elections. During the general election phase of the campaign, Facebook was used more for attacks on opponents, discussion of issues (as opposed to candidate image), and "ceremonial" posts (as opposed to calls for action). These differences, however, were almost entirely non-existent during the primary phase of the election.

Farkas and Bene (2021) also find some differences across Facebook and Instagram in the images used by Hungarian politicians. The authors suggest that the "feel" of images is somewhat different across the two platforms. While photos used on both platforms emphasize candidates as individuals (not as party leaders), Instagram tends to show more of candidates' private lives as opposed to their work lives. In addition, there are differences in how campaigns use hyperlinks across platforms, with Twitter posts using more hyperlinks and thus linking to a broader range of sources and content (Jackson et al., 2020).

In sum, the evidence speaking to platform convergence versus divergence has not reached a consensus. Generally, research compares only two platforms, examines only a few elements of the content of campaign communications, focuses on organic rather than paid content, and is based on a small number of platforms. In the current study, we provide considerably more data speaking to platform convergence verses divergence and do so through examining paid political advertising. We examine three important aspects of content in over 145,000 ads placed by almost 200 candidates on four different platforms.

Platform Divergence

There are at least three broad explanations for platform divergence: audience differences, differences in digital architecture, and different platform genres. First, different platforms have different audiences. For instance, some platforms skew much younger than others. Among Americans aged 18–29, 65% report using Snapchat, 48% report using TikTok, and 30% report using Instagram. But among Americans older than age 65, 2% report using Snapchat, 4% report using TikTok, and 13% report using Instagram (Pew Research Center, 2021). Importantly, in addition to real differences in audience composition, there are perceptions of differences in who uses various digital and social media platforms. Practitioners perceive Facebook to have the broadest audience representing a span of demographic characteristics (Kreiss et al., 2018). By contrast, campaigns perceive Twitter as having a more politically engaged audience, including many journalists (Kreiss, 2016; Stier et al., 2018). And audience data suggest that Instagram and Snapchat are best for reaching young audiences. Thus, given real and perceived differences in who is using each platform, we might expect campaigns to customize their messaging for each. Indeed, Kelm (2020) finds a link between German politicians' perceptions of audience expectations and the ways in which those politicians use Facebook.

In addition, different platforms have different "digital architectures," which refers to "the technical protocols that enable, constrain, and shape user behavior in a virtual space" (Bossetta, 2018, p. 473). Bossetta notes the importance of a platform's network structure, its functionality, its algorithmic filtering (the visibility, prioritization, selection, and sequencing of posts), and datafication (the digital traces left by users that allow campaigns to target, match, and perform analytics on users). A platform's architecture creates different affordances, which are "the perceived actual or imagined properties of social media, emerging through the relation of technological, social, and contextual, that enable and constrain specific uses of the platforms" (Ronzhyn et al., 2022, p. 14). Thus, if campaigns perceive that some platforms are better at enabling some campaign goals than others, they may use these platforms in different ways.

Kreiss et al., 2018 also point to another difference across platforms that may affect how they are used: its genre. Genres are "shared social conventions that structure discourse and enable people to communicate" (p. 13). Importantly, audiences expect to find different types of content on different platforms. Different platforms *feel* different. As Matassi and Boczkowski (2023) explain, "People often sense that certain ways of communication and self-presentation are socially acceptable on some platforms and not on others" (p. 67). For example, people turn to Twitter when they want news and politics, and people turn to Instagram and Snapchat when they want good visual storytelling. So while a candid photo of a candidate making dinner in his home kitchen may feel like it belongs on Instagram, it may feel out of place on another platform. This all suggests that campaign content is not one-size-fits-all.

Digital Ad Content

As discussed, existing cross-platform research has analyzed several different aspects of campaign messaging, including the emotions conveyed by images, glimpses of candidates' private lives, and a focus on issues. But when it comes to studying political advertising, we argue that three features are especially important.

The first is the use of partisan language. Partisanship is a key predictor of the vote in the United States (Green et al., 2004), and the Internet is fertile ground for polarized language (Simchon et al., 2022). Some speculate that when people are exposed to more partisan language, they could become even more polarized (Fowler et al., 2021b). When it comes to political advertising on *television* in the United States, partisan cues are often de-emphasized or even absent (Neiheisel & Niebler, 2013), presumably because the intended audience with television ads is the independent, persuadable voter. However, the language employed in campaigns' Facebook ads is more partisan than the language used in their television ads (Fowler et al., 2021b).

But should we expect differences across platforms in the use of partisan language? Audience theories do not necessarily lead to this expectation, for even though audiences differ in age and other characteristics across the platforms we examine, the fact that ads can be targeted to highly partisan or non-partisan audiences makes the platforms equally useful for ads that employ a range of partisan language. But explanations of content divergence that are grounded in platform affordances might suggest that we would find variation in the use of partisan language. Consider, for instance, YouTube, which is built for video. Because of this, one might expect to find the sort of political advertising that one finds on television, which, as noted, tends to avoid strong partisan appeals. Explanations based on platform genres might also suggest platform divergence. For instance, ad sponsors on Instagram, whose audiences "expect photos and storytelling, not official campaign statements" (Kreiss et al., 2018, p. 16) might be less likely to use partisan language than other platforms. Given that we do not have strong expectations about how much cross-platform divergence we should see in the use of partisan language, we propose a research question.

RQ1: How does the use of partisan language in paid political advertising vary across digital platforms?

The second ad feature we examine is the tone of advertising. The tone of advertising is important because negativity is thought to have a different psychological impact than positivity. Negativity tends to be given more weight in impression formation (Soroka, 2014) and can mobilize political action (Martin, 2004). Campaign negativity has also been shown to reduce voter turnout in certain circumstances, such as when it is directed at a candidate one has already decided to support (Krupnikov, 2011).

As with the use of partisan language, we argue that differences in audience may not matter all that much given the capacity to target particular audiences, but a focus on affordances might suggest some more variation in tone across platforms. YouTube, for instance, specializes in video, which can be quite effective in creating (negative) emotional appeals. Google search, by contrast, with its focus on text, might be better at conveying information, such as voting locations, which does not require the use of negativity. Genres may also matter, with Instagram's focus on images and personalization, rather than hard-core politics, more conducive to positive messaging. Again, given our lack of strong expectations, we propose a research question:

RQ2: How does tone in paid political advertising vary across digital platforms?

The third ad characteristic that we focus on are the goals of the advertisement. Although it is assumed that the goal of television advertising is singular, namely, persuasion (Fowler et al., 2021a), scholars have noted several different goals pursued by online campaign advertisements. Ballard et al. (2016) note that one quarter of the online ads in the 2012 U.S. presidential election had a goal of fundraising, and 20% were aimed at recruitment. Thus, the distribution of goals of digital ads can provide insight into campaign strategy and how it can change over time (see Ridout et al., 2021b). Different audiences on different platforms may result in campaigns pursuing different goals, such as campaigns delivering more mobilization messages to the younger voters on Instagram. Different platform affordances, the result of digital architectures, could also result in the pursuit of different goals on different platforms, such as more ads that attempt to persuade on YouTube given the platform's focus on video. Thus, we propose a third research question:

RQ3: How do the goals of paid political advertising vary across digital platforms?

Data and Methods

We limit our focus to paid advertising by candidate sponsors in U.S. House and Senate races that were placed on four different platforms: Facebook, Instagram, YouTube, and Google search.[2] We examine ads that were delivered during the heart of the fall campaign, between September 1, 2020, and November 3, 2020, which was Election Day. We also limit our analysis to those 196 candidates who placed ads on all four platforms.

The Facebook and Google ad collections rely on the same technical processes but are different in what is used as the starting point for identifying the relevant ads. The Facebook Ad Library API (which also includes Instagram) provides the text for several attributes of an ad, such as the ad creative body, link title, link caption, page name, and paid-for-by disclaimer. To identify the relevant ads, we started with a collection of very broad keywords that included the name of a candidate, the candidate's home state, and paraphrased references to the race (e.g., Senate, Tennessee, TN). Using the search endpoint of the API, we obtained individual ads mentioning candidates to identify the names of pages running election-relevant ads. We then used a different endpoint of the Facebook Ad Library API to retrieve all ads that were launched from these pages and added them into our collection. Our final step was to match the page names and disclaimers to a master list of candidate names from the Federal Election Commission (FEC).

For Google, we rely on the Google Transparency Report, a library for ads run on Google platforms, including Google search and YouTube. However, the text fields of the ads in Google's library are not searchable the same way as in Facebook's, which meant we had to restrict the ads to those that were relevant to us in a different way. We first identified all U.S. ads that ran during the election period. We then matched the names of advertisers to the names in the FEC's collection. That allowed us to reduce the number of ads that we would need to inspect. We then visited the ads' web pages within the Transparency Report using the Selenium web scraper (for which we had Google's permission) and retrieved their text and images. To download video ads, which are hosted on YouTube, we relied on the widely used youtube-dl package. Then we used automatic speech recognition in order to extract the transcripts of these videos.

The distribution of the number of text characters that we examined in each ad varied across platforms, as Table 1 shows. Facebook-only ads contained, on average, the most characters (659), and Google text ads had the fewest (247). On average, YouTube ads contained 396 characters, but there was a wide standard deviation.

Our next step was to append content information to the dataset, specifically partisan language, tone, and goals.

Partisan Language

Our initial goal was to measure the extent to which the language used in the ad was partisan, that is, invoked particular words and phrases that are commonly associated with either the Democratic or Republican party. Our intuition, following the work of Fowler et al. (2021b), is that ads whose

Table 1. Summary Statistics on Number of Characters by Ad Platform.

Platform	Mean	SD	Kurtosis
Facebook only	658.78	527.68	2.02
Facebook and Instagram	578.21	339.59	6.37
Instagram only	522.40	303.13	1.84
Google text	246.55	47.68	3.15
Google video (YouTube)	396.33	759.99	1309.23

language is easy to classify are using partisan language, compared to ads more difficult to classify, which are using more neutral language or language employed by both Democrats and Republicans. Thus, the probabilities produced by our party classifier—the probability that the ad sponsor is Republican and the probability that the ad sponsor is Democratic—can be a useful measure of the partisanship of the ad's language.

To create our party classifier, we relied on a set of 850k Meta ads and 102k Google ads from 2020 for which we knew the party because we had matched them to FEC data, which contain the party of each candidate. Because party labels were available for both Meta and Google ads, we included data from both to create a combined model that would perform well on either platform. We retained 10 Democratic and 10 Republican advertisers who had between 500 and 1000 creatives for the test set, using the rest for the training set.[3] For both sets, we masked the names of politicians, so that they would not bias our classification of partisanship (this also helps performance, as a model that retains the names has an F1 score of 0.81 compared to 0.84 for a model that removes them). Furthermore, we deduplicated ads that were exact matches of one another.[4] At this point, we were left with 2236 ads for the test set, and 102,483 ads for training. Given that almost two thirds of the latter were ads by Democrats, we retained all Republican ads (38,465) and sampled an equal number of Democratic ads. This ensured that the classifier wouldn't prioritize performance on Democratic ads. The ads were tokenized and represented as term frequency-inverse document frequency (TFIDF). We experimented with multiple classifiers (Logistic Regression, Multinomial Naïve Bayes, SVM, and Random Forest) and ultimately opted for the Logistic Regression due to its simplicity and relatively smooth output probability distribution (as opposed to a lopsided distribution where most ads are either $p > .99$ or $p < .01$). We used an L2 penalty (often called ridge regression in political science) in order to gain smoother probabilities. The regularization strength parameter was tuned until the recall for both Democrats and Republicans was the same (0.84). This guaranteed similar output probability distributions, which is essential for ensuring that a very Democratic ad and a very Republican ad are classified as being similarly partisan.

The classifier trained in this way achieves 0.84 overall F1 score (the harmonic mean of precision and recall) on the held-out test set. Class F1 scores are 0.84 for Democratic ads and 0.85 for Republican ads. In short, our classifier does a good job of predicting the party of the candidate sponsoring the ad. We then apply this classifier to both the Facebook and the Google data to produce both party predictions and, more importantly, the probabilities that each ad was sponsored by a Republican and a Democrat.

We then turned these party probabilities into a measure of the partisanship of the ad's language. We calculated the absolute value of the probability that the ad was sponsored by a Democrat and subtracted from that the probability that the ad was sponsored by a Republican, resulting in a theoretical range from 0 to 1. Thus, when these probabilities are similar, the resulting score is low, suggesting that the party classification is "difficult" based on the ad's language and thus that the campaign is using non-partisan language. By contrast, when there is a large difference in the two probabilities, then it is much easier to classify the ad, suggesting that it uses more partisan language.

Tone

To determine the tone of an ad, we rely on a standard and widely used reference-based measure in the literature on political advertising (Geer, 2008) that requires information on which candidates are referenced (mentioned or pictured) in the ad. Ads that contain no reference to either candidate within the ad body or reference only the sponsoring candidate are defined as positive (since advertising is created by candidates to help their campaigns), those that mention only an opponent are negative, and those that mention both the sponsoring candidate and an opponent are contrast ads. Geer's definition does not take into account the sentiment of the ad, which also removes subjectivity from coding. For instance, an ad that featured the sponsoring candidate talking about rising crime, a bad economy and corruption in government would be coded positive if it did not mention any other candidates.

To identify the candidates referenced, we relied on an entity-linking model consisting of named-entity recognition, a lookup table containing all federal candidates, and candidate and ad embeddings. Information from the text fields, automatic speech recognition (ASR), and text onscreen extracted using OCR were used to determine mentions in text, audio, and imagery. To determine which candidates were pictured, we relied on a pretrained face recognition deep learning model (AWS Rekognition) to search all federal candidate faces in ad images and videos. Importantly, the ad disclaimer or sponsorship field was not used in determining which candidates were referenced, for if it were included, then it would be impossible to have a negative ad under Geer's (2008) definition.

After identifying the referenced candidates, we were able to classify each as positive, contrast, or negative, following Geer (2008). But in a very few instances (fewer than one percent of ads) as described above, an ad mentioned a candidate only in the disclaimer or sponsorship information; no candidate was referenced in the main part of the ad. For instance, an ad might ask viewers to "chip in" and feature a graphic of a post-it note with a fundraising deadline without naming or picturing the candidate in the body of the ad. We coded such ads as positive, given our presumption that candidates run ads to advance their own election efforts. In these instances, the viewer would still be able to learn the sponsoring candidate through the disclaimer information. Our tone predictions were then turned into a negativity scale that ranged from 1 to 3, with 1 denoting a positive ad, 2 denoting a contrast ad, and 3 denoting a negative ad.

Because the tone measure is constructed solely based on who is referenced in the ad, we validated our results by asking human coders to identify all candidates mentioned and pictured in the ads.[5] We then compared the machine to the human coding results for a sample of 3671 ads. The approach achieved an F1 score of 0.87, suggesting that agreement between our machine classifier and the hand-coded ads was high.

Goals

Several categorizations of online political ad goals have been proposed. Ballard et al. (2016) focus on four goals: get-out-the-vote, fundraising, persuasion, and campaign recruitment. Zhang et al. (2017) uses slightly different categories, including "calls to action," information, persuasion, advocacy, and attack. Ridout et al. (2021b) coding initially identified 16 different goals, but they aggregate those into five: acquisition of audience information, fundraising, mobilization, persuasion, and other.

Our own coding scheme was developed using these past categorizations. Coders were asked, "Which of the following elements are included in the ad?" and given the opportunity to choose one or more of the following: request for donations, request for viewer to contact a legislator, selling of merchandise, GOTV efforts, advertisement of an event/rally, asking the viewer to fill out a survey/poll, solicitation of viewer's email/name/information (other than a poll), and inviting viewers to learn more. In order to identify "persuasion" ads, coders were asked, "Is the primary purpose of the ad to promote or attack a candidate or persuade the viewer to vote for a specific candidate?" Response choice were "yes" or "no, the primary goal is something else." Our coders answered these questions about 3671 Facebook ads to build our training dataset.

In this analysis, we focus on four of the most common goals: persuasion, fundraising (the request for donation category), mobilization (the GOTV category), and encouraging viewers to learn something (the invite to learn more category), which is similar to the information category used by Zhang et al. (2017).[6] For instance, one ad from Kim Olson, an agriculture commissioner candidate in Texas, asked viewers to check out her blog for her thoughts on renewable energy.

To detect the goals of an ad, we trained and applied a set of supervised models. More specifically, we trained the model using 3671 ads that were hand-coded as described above, split into an 80/20 train/test set. Because an ad can be coded to have multiple goals, we chose to estimate multiple binomial models rather than a single multinomial one. We experimented with both Random Forest and DistilBERT models but ultimately opted for the Random Forest given that performance was similar and the Random Forest—through the use of cross-validation and a logistic calibration function—provides a wider range of output probabilities (whereas the predicted probabilities of the DistilBERT model are effectively all 0.01 or 0.99). F1 scores were 0.99 for donate, 0.93 for mobilization, 0.84 for learn, and 0.85 for persuasion.

In the end, the usable dataset to which we applied our model consisted of 145,914 unique ads delivered in House and Senate races. These included 32,990 ads that were placed only on Facebook, 2774 ads placed only on Instagram, and 92,771 ads placed on both platforms. There were also 12,166 ads placed only on YouTube and 5213 ads placed only on Google search.

Results

We ended up with 152 U.S. House and 44 U.S. Senate campaigns that had paid advertising on Facebook, Instagram, YouTube, and Google search, allowing us to make within-candidate comparisons on these 196 campaigns across all four platforms. Table 2 shows average scores across campaigns on all six dependent variables by platform. The top four rows included all 196 cases, where the next two rows show averages for ads that were seen only on Facebook and only on Instagram within those 26 campaigns that posted both "Facebook only" and "Instagram only" ads. The final two rows show a comparison between Facebook video ads and YouTube ads within those 191 campaigns that purchased both types of advertisements. Letters (a, b, c, d) indicate a statistically significant difference at $p < .05$ between the examined platform and each other platform on the basis of a t test.

Partisan Language

The "partisan language" column in Table 2 displays mean values on the partisan language score for each platform. Figure 1 shows these same data visually through density plots across the 196 campaigns for each of the four platforms. The mean on the partisan language scale is lowest on YouTube at 0.17 and highest on Facebook at 0.23. The score is 0.18 on Google search and 0.22 on Instagram. Thus, viewers will be exposed to the least partisan language on YouTube and the most partisan language on Facebook and Instagram.

The magnitude of the difference in the use of partisan language between Facebook and Instagram is not large, which is not a surprise given that many of the same ads are placed on both Meta platforms. If one analyzes just those 26 campaigns that place some ads only on Facebook and other ads only on Instagram, there is more partisan language on Facebook (mean = 0.27) than on Instagram (mean = 0.19).

In sum, YouTube ads and Google search are consistently less partisan in their language than are the other two platforms, lending support to the idea of platform divergence.

Table 2. Average Content Feature Scores Across Platform.

	N	Partisan language		Negativity		Donate	
All Facebook	196	0.233	bcd	1.35	c	0.319	cd
All Instagram	196	0.223	acd	1.35	c	0.301	cd
Google search	196	0.176	ab	1.25		0.221	abd
YouTube	196	0.174	ab	1.40	c	0.008	abc
Facebook only	26	0.271	*	1.49	ns	0.429	*
Instagram only	26	0.193		1.44		0.199	
Facebook video	191	0.238	*	1.36	ns	0.324	*
YouTube	191	0.175		1.40		0.008	
	N	**Persuasion**		**Mobilization**		**Learn**	
All Facebook	196	0.544	d	0.037	d	0.069	cd
All Instagram	196	0.561	d	0.037	d	0.075	cd
Google search	196	0.555	d	0.031	d	0.191	abd
YouTube	196	0.971	abc	0.004	abc	0.017	abc
Facebook only	26	0.456	*	0.031	ns	0.032	ns
Instagram only	26	0.652		0.046		0.076	
Facebook video	191	0.567	*	0.031	*	0.069	*
YouTube	191	0.970		0.004		0.017	

Note. *$p < .05$, a: significantly different from Facebook at $p < .05$, b: significantly different from Instagram at $p < .05$, c: significantly different from Google search at $p < .05$, d: significantly different from YouTube at $p < .05$.

Tone

When it comes to the tone of advertising, there were some differences across platforms. Negativity ranged from 1.25 on Google search to 1.40 on YouTube, as Figure 2 shows. The negativity value was identical on Facebook and Instagram at 1.35. The language used in Google search ads was significantly less negative than the language on each of the other three platforms, though the others were not significantly different among themselves at $p < .05$. Within those 26 campaigns that placed both Facebook-only and Instagram-only ads, the Facebook-only ads were slightly more negative, 1.49 versus 1.44, but the difference was not statistically significant. Interestingly, while YouTube used the least partisan language, it was the most likely to contain an attack on an opponent.

Donation

Ads that seek donations are most common on Facebook and Instagram. Thirty-two percent of ads on Facebook in the average campaign have a donation goal, compared to 30% on Instagram. On average, 22% of Google search ads have a donation goal. YouTube is again distinct, as it is much less likely than the other platforms to feature ads that focus on fundraising. Across campaigns, less than 1% of YouTube ads, on average, have a donation goal.

One interesting finding regarding donation ads is that requests for donations are much more common in Facebook-only ads than in Instagram-only ads—43% on Facebook versus 20% on Instagram (Table 2). This is a sign of platform divergence, that campaigns are using Facebook for different purposes than they are using Instagram. To campaigns, Instagram may not “feel” as appropriate for fundraising as Facebook (Figure 3).

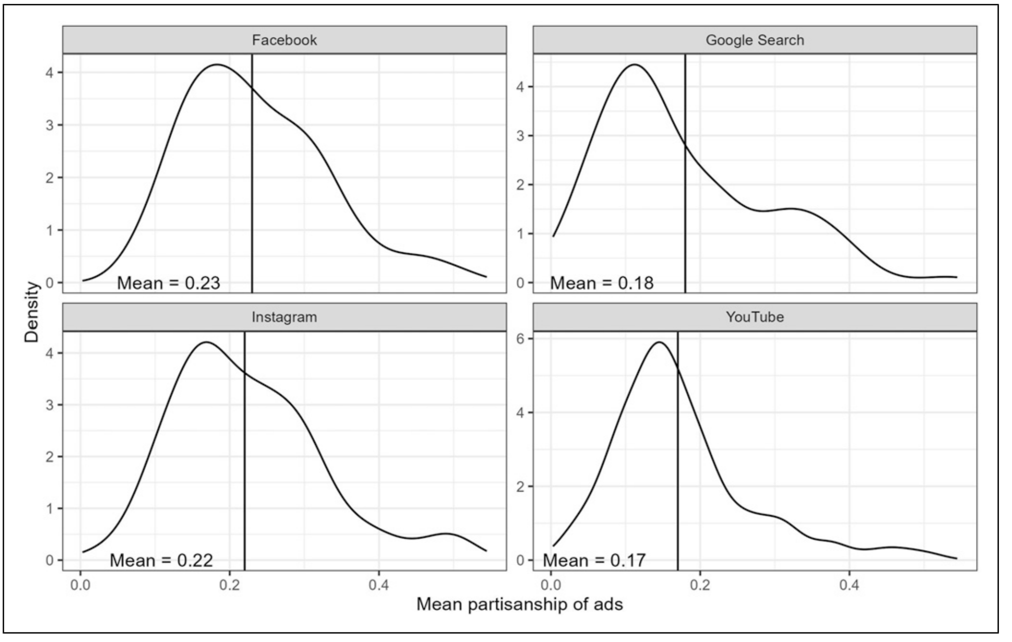

Figure 1. Distribution of Average Partisan Language by Platform.

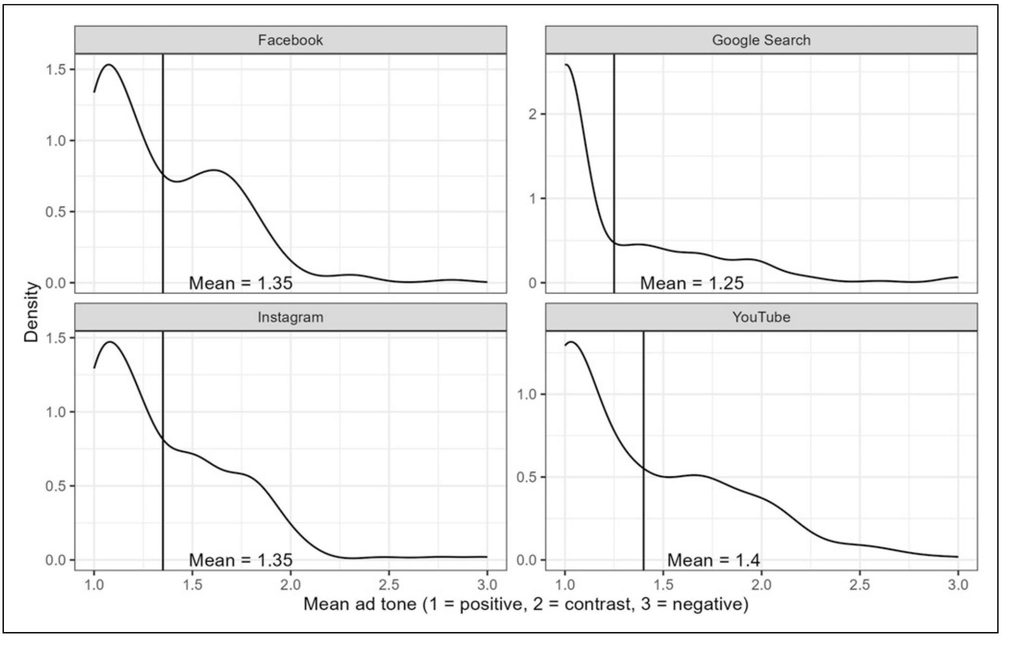

Figure 2. Distribution of Average Tone by Platform.

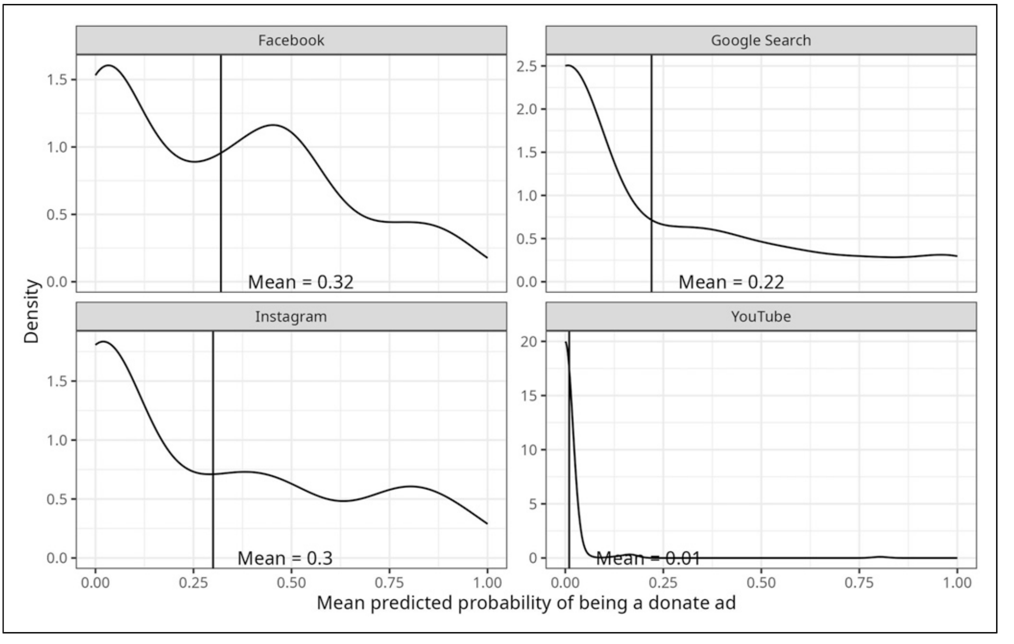

Figure 3. Distribution of Donation Ads by Platform.

Persuasion

One of the primary goals of campaigns in digital advertising is persuasion, but the extent to which that is pursued varies quite a bit across platforms, as Figure 4 shows. On Facebook, 54% of the average campaign's ads were devoted to persuasion, with many campaigns having a small percentage of persuasion ads. The percentage of persuasion ads was slightly higher on Instagram at 56%. On Google search, 55% of ads employed persuasion. Persuasion was, by far, most common on YouTube, a platform that emphasizes video. On YouTube, persuasion was an ad goal 97% of the time in the average campaign. Interestingly, if one examines the 26 campaigns that placed both Facebook-only and Instagram-only ads, there was a marked difference in the use of persuasion ads on the two platforms. The average proportion of ads with a persuasive message was 0.46 on Facebook and 0.65 on Instagram, a difference that was statistically significant at $p < .05$.

Note that, because our data cover advertising placed between the first of September and Election Day, the percentage of persuasion ads is likely higher than if our data had covered the entire campaign, as evidence suggests that persuasion ads ramp up closer to Election Day (Ridout et al., 2021b).

Mobilization

As Figure 5 shows, ads that primarily seek to mobilize are relatively rare on all platforms, but there is some variation. Such ads are most common on Instagram and Facebook (about 4% of ads, on average) but are non-existent on YouTube. Indeed, the average proportion of mobilization ads is significantly lower ($p < .05$) on YouTube than on each of the other platforms. But there is no significant difference in the use of mobilization ads between Instagram-only and Facebook-only ads (Table 2).

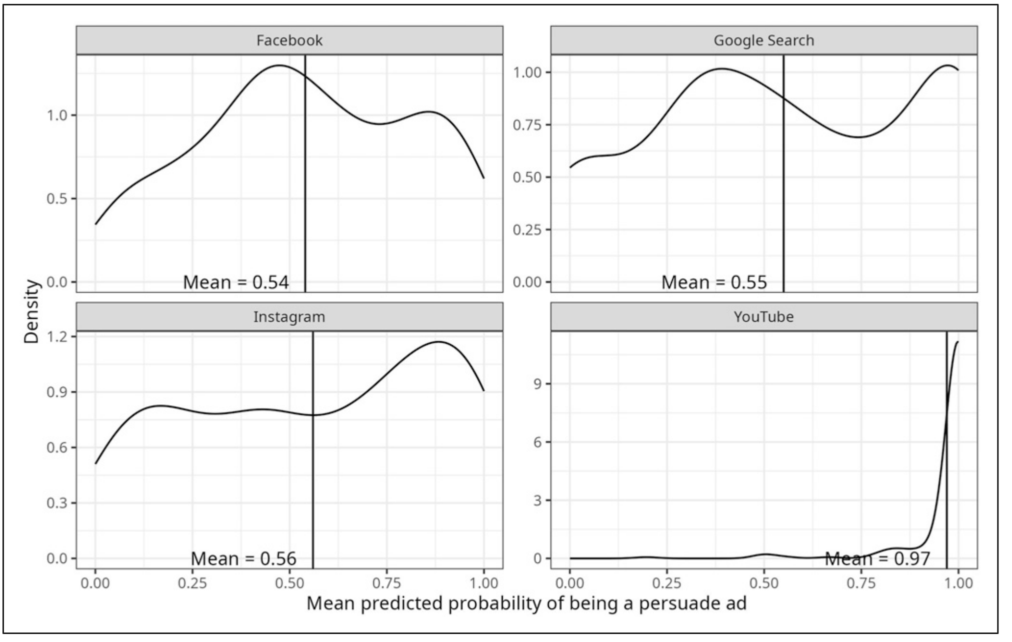

Figure 4. Distribution of Persuasion Ads by Platform.

Learning

Finally, we consider the extent to which ads encourage voters to learn something. As Figure 6 shows, these ads are relatively common on Google search (the proportion of "learn" ads placed by the average campaign was 0.19), but they are much rarer on the other three platforms. The proportion of ads that primarily encourage learning on Instagram is 0.08 and is 0.07 on Facebook. That proportion drops to 0.02 for YouTube. When it comes to this type of ad, there was not a significant difference in their use in Facebook-only ads compared to Instagram-only ads.

Platform or Mode of Communication?

Our focus has been on differences across platforms, but platforms may serve as a proxy for the mode of communication. We can get some leverage on this question—the extent to which content differences depend on the platform itself versus the audio and visual format of an ad—by taking a close look at ads on Facebook. Among the approximately 125,000 Facebook ads in our dataset, 51% contain only an image, 38% contain a video, and 11% contain both a video and an image. Thus, we were able to make comparisons between Facebook video ads (the 38%) and YouTube ads, which are entirely video. Differences in content that remain reflect platform differences as the use of video is controlled for.

In short, our analysis, reported in Table 2, suggests that some substantial differences exist between YouTube and Facebook, even when we examine only video ads. Facebook video ads use significantly more partisan language (a substantive difference of 0.06 on the 1-point scale), are much more likely to seek a donation, are much less likely to use persuasive appeals, and are significantly more likely to encourage people to learn something and to mobilize. The one area that did not differ between YouTube and Facebook video was the degree to which ads were negative.

To provide additional support for the impact of video versus non-video format on content, we estimated several mixed-effects models at the ad level, one for each content variable. (We did not

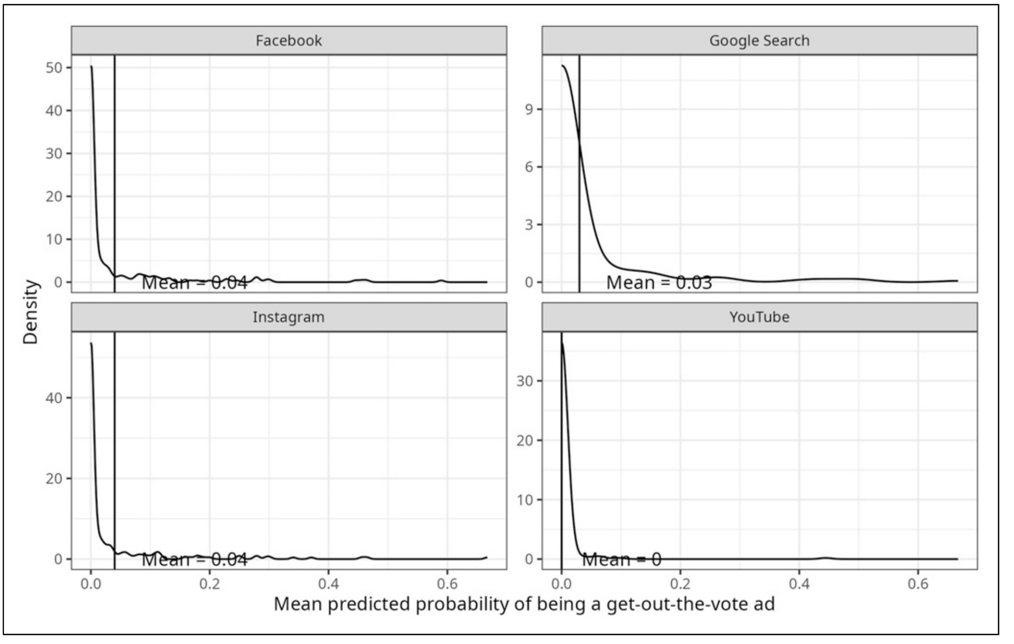

Figure 5. Distribution of Mobilization Ads by Platform.

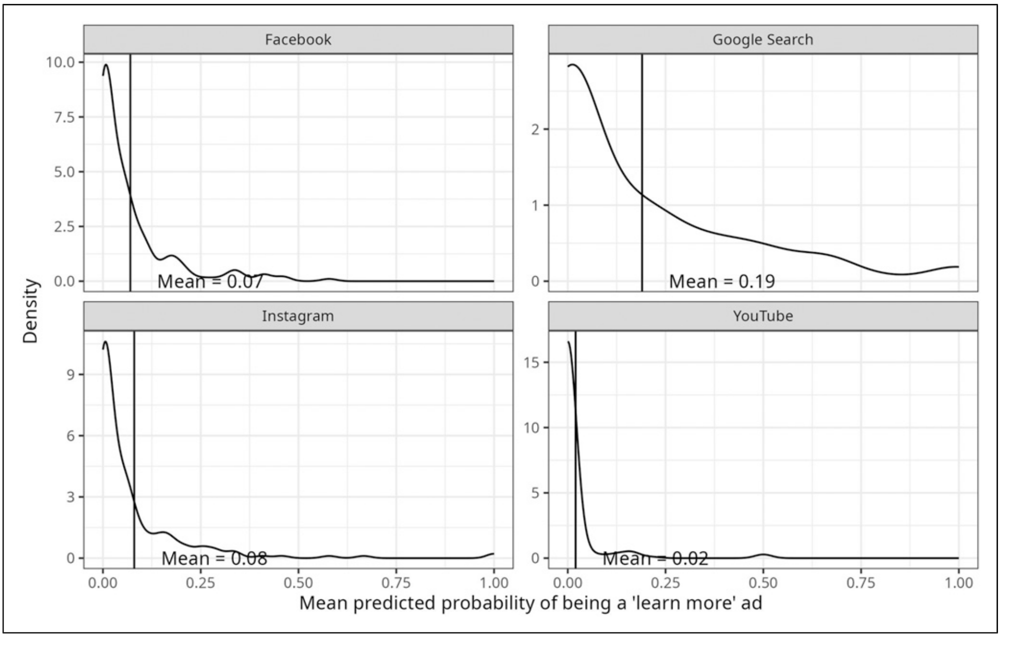

Figure 6. Distribution of Learning Ads by Platform.

estimate models for ads with mobilization or learning goals given the small number of ads falling into each of these categories.) Each model contains an indicator of the company (Meta v. Google), an indicator of whether the ad was a video ad or not, the interaction of the company with the video indicator, and an indicator of whether the race was a Senate race (as opposed to a House race). By structuring the independent variables in this way, we can make distinctions between Google search ads, Google video ads (YouTube), Meta (Facebook and Instagram) video ads, and Meta non-video ads.

We estimated linear mixed-effects models for the negativity and partisanship dependent variables; for the donation and persuasion models, which have a dichotomous dependent variable, we used a mixed-effects model with a binomial link function. In addition, for each of these models, we weighted by logged candidate spending on each creative, thus giving more weight to those ads that were placed more often. Model estimates are found in the Appendix.

We used the model estimates to create predicted values of each content feature. Figure 7 shows predicted partisanship across ad types. Google video (YouTube) ads were the least partisan, with a predicted score of about 0.14 on the 0 to 1 scale. Google non-video (search) ads were significantly more partisan, with a predicted score of 0.18. Facebook and Instagram ads, whether video or not, were the most partisan, with predicted scores of above 0.24.

Figure 8 shows predicted levels of negativity by platform and media type. Google video ads (YouTube) are the most negative at 1.45, much more negative than the non-video (search) ads on Google. But if one looks at ads on Meta, video format matters less for predicted levels of negativity, as there is no statistically significant difference between Instagram and Facebook video and non-video ads.

When it comes to donation ads, Figure 9 reveals that, for Google, YouTube ads are significantly less likely to feature a request for a donation than search ads. The probability of an ad being a donation ad is near 0 on YouTube versus 0.10 for a Google search ad. But when we turn to the two Meta platforms, differences across video and non-video ads again disappear.

Finally, we consider the probability that an ad is considered a persuasion ad in Figure 10. Again, there are distinct differences between Google search ads, with a probability of just above 0.5 of being a persuasion ad, and YouTube, with a probability approaching 1. But for the Meta platforms, the difference in the probability of being a persuasion ad does not differ by video versus non-video format.

In sum, our ad-level analyses, which are focused on differences between video and non-video ads, suggest that Google's video offering (YouTube) is quite distinct from its non-video offering (search advertising), but the video and non-video offerings on Meta are used in a quite similar fashion.

Discussion

Our research demonstrates that the content of political campaign messages differs not just between "old media" and digital and social media but across social media platforms as well, though the extent of the difference depends on the specific platforms and content features being compared. We found, first, that the use of partisan language is higher on Facebook and Instagram than it is on Google search and YouTube. One possibility for this finding is that ads on Google and YouTube are more likely to be delivered to a politically diverse audience that is similar to a television audience. Whereas Facebook and Instagram ads may be targeted to individuals based on their (presumed) political characteristics, Google and YouTube targeting likely depend much more on search behavior given the centrality of search to Google's business. Thus, a Google search for a candidate's name will likely result in delivery of a persuasion ad that is designed to appeal to the swing voter.

Second, we found that YouTube was the most negative and Google search the most positive platform, though the divergence was not enormous. Our speculation is that YouTube, because it is the platform that most closely resembles television, is often being used to attack opponents (even

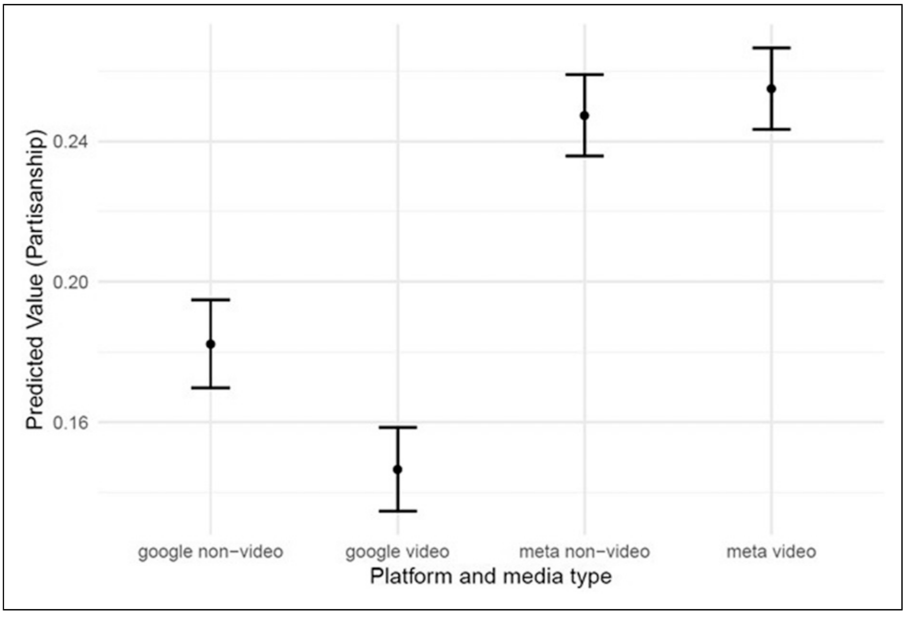

Figure 7. Predicted Value of Partisanship by Platform and Media Type.

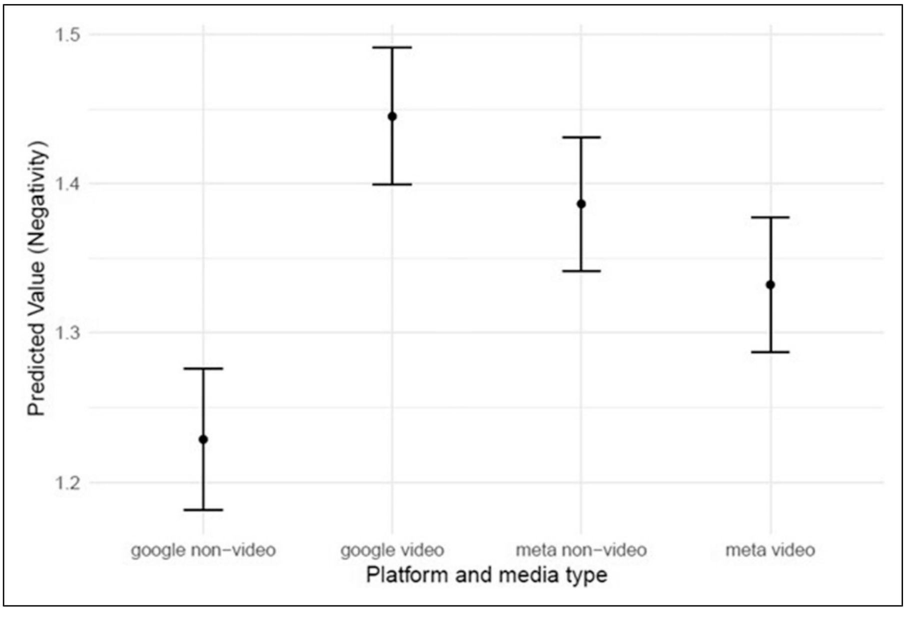

Figure 8. Predicted Value of Negativity by Platform and Media Type.

if those attacks are not explicitly partisan ones). Search, by contrast, may be less likely to contain an attack because it can be used to convey information in response to a user's search.

Third, we found that the primary goals of ads are also different across platforms, with YouTube being the most distinct. YouTube ads were much less likely to feature fundraising appeals, much

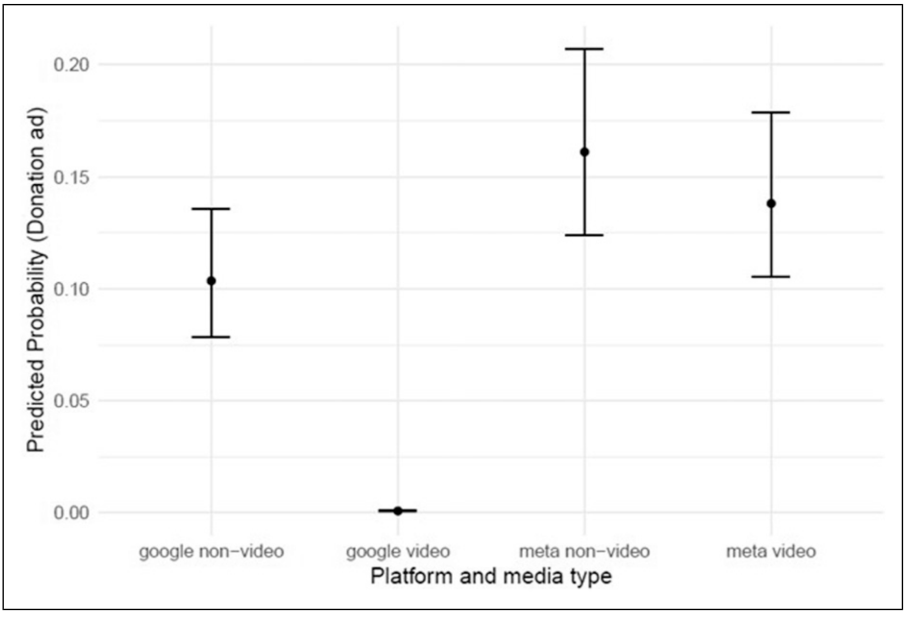

Figure 9. Predicted Probability of Donation Ads by Platform and Media Type.

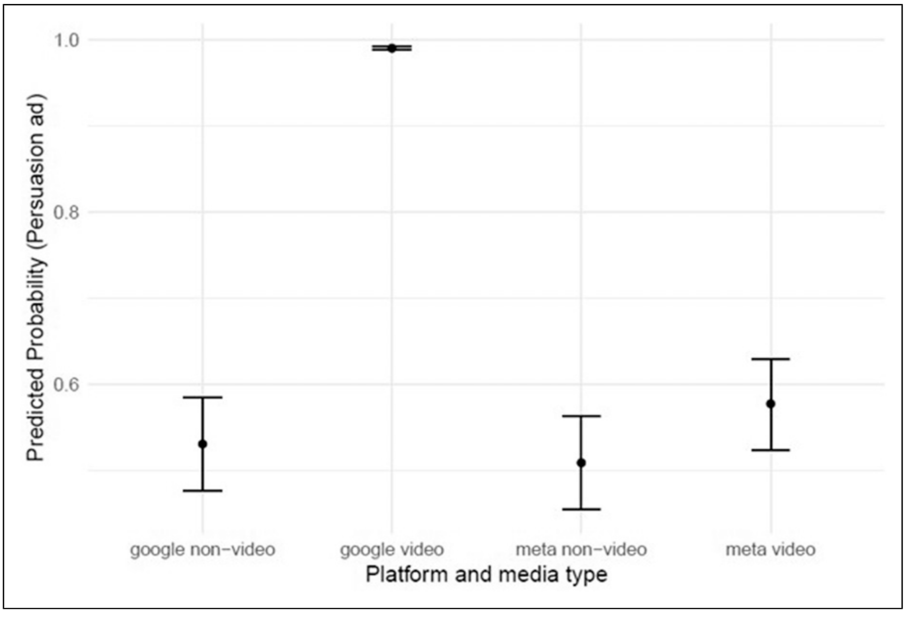

Figure 10. Predicted Probability of Persuasion Ads by Platform and Media Type.

less likely to feature mobilization, much less likely to encourage viewers to learn but were much more likely to be primarily about persuasion.

We wondered if YouTube's focus on persuasion might stem from its unique ability to deliver video. Two pieces of evidence speak to this. First, when we compared YouTube with

Facebook ads that used a video format at the campaign level, we found that differences remained. Thus, it is not just YouTube's use of video that makes it the preferred platform for persuasion. Perhaps it is YouTube's reputation for being all about video or characteristics of the YouTube audience that make its use distinct. Second, when we compared video to non-video ads on the Meta platforms in our ad-level analysis, we found no significant differences. Again, this suggests that platform is more important than video capability in explaining content divergence.

All told, our findings lend more weight to a platform divergence perspective than a platform convergence perspective, as on every feature examined there were at least some statistically significant differences in content across platforms—and in many instances, the differences were large enough to be substantively meaningful as well. We are not able to explain this divergence definitively using one of the theories outlined earlier (audiences, digital architectures, or genres), as we do not have sufficient variation across the platforms examined here. But given that differences remained even after controlling for video capability, the digital architecture theory cannot suffice in full. And it is not the case that ad characteristics vary linearly with the number of text characters that we examined, as Google search, with the fewest characters, was sometimes more similar to YouTube and sometimes more similar to Facebook and Instagram.

It is also worth pointing out that convergence is a more plausible story when talking about Instagram and Facebook, two platforms owned by Meta. One reason for this is that many of the ads in our sample (about 72%) were posted on both platforms. Still, when we examined those ads only posted on Facebook or only posted Instagram, there was less convergence between the two platforms, with Facebook ads using slightly less partisan language, using many more requests for donations, and being less likely to pursue persuasion.

As campaigns transition away from traditional means of campaign communication, such as television advertising and billboards, and toward more individualized campaign messaging made accessible through digital and social media, one concern is that they may alter the content of the campaigns to which voters are exposed—and in a way that is deleterious to voters or democracy. For instance, evidence shows that the move to digital campaigning may strengthen partisan polarization, as television advertising uses less partisan language than does advertising on Facebook (Fowler et al., 2021a). Moreover, some are concerned that the move to targeted digital ads may result in a campaign environment in which different people receive different messages—and thus different understandings of what the campaign is about and what the parties' priorities are (Borgesius et al., 2018). This can make it difficult for the winners to govern, as different constituencies bring with them different ideas about what the winning politicians promised to do in office. Finally, Matthes et al. (2022) find that perceived online political micro-targeting can decrease trust in democracy.

Our research suggests that these concerns may exist not just when comparing old campaign means of communication, for example, comparing television to digital media, but when comparing across digital and social media platforms. People who spend hours each day on Facebook, for instance, may view the campaign as much more partisan than those whose social media diet is concentrated on watching YouTube videos.

A few caveats are in order. First, we did not examine all platforms that sell political advertising—the most important omission being Snapchat (although the amount of spending on that platform is much smaller than the others we do examine)—nor did we examine non-paid content online, such as that found on TikTok. Moreover, by focusing on the last 2 months of the campaign, we are focusing on a period that tends to be more negative and more about persuasion than the campaign as a whole, and thus our conclusions about the content of advertising do not necessarily hold for earlier in the campaign or during primary races.

Third, our data come from the United States and thus may not generalize to other countries. Although campaigns in most countries have access to advertising on these same platforms, differences in campaign finance, election rules, and privacy regulations may affect how campaigns use various platforms. For example, in a country with compulsory voting and publicly financed campaigns, persuasion may be the primary goal regardless of the platform. And in places with strict privacy regulations, such as in the European Union, the ability to take full advantage of each platform by targeting different messages to different audiences based on the characteristics of those individuals may not be available. Thus, a similar message may be sent to all voters, regardless of the platform. At the very least, though, our data from the United States show that when campaigns are relatively unencumbered by regulation—and have access to ample supplies of money—the capacity to use different platforms in different ways is eminently possible.

Of course, a single study will not be the final word on the question of platform convergence versus divergence. But we believe our study remains an important contribution to the larger debate given our focus on paid political advertising, our examination of four platforms (as opposed to two or three), our large sample of campaigns (196), the large number of ads examined (145,000), and our focus on six different features of campaign advertisements (the use of partisan language, ad tone, and four different campaign goals).

Appendix

Ad-Level Models Predicting Content Features

	Dependent variable:			
	Ad Negative	Partisanship	Goal Donate	Goal Persuade
	Linear mixed-effects (1)	Linear mixed-effects (2)	Generalized linear mixed-effects (3)	Generalized linear mixed-effects (4)
Meta	0.158***	0.052***	0.508***	−0.087***
	(0.009)	(0.004)	(0.019)	(0.018)
Video	0.216***	−0.140***	−4.891***	4.549***
	(0.010)	(0.004)	(0.039)	(0.029)
Senate	−0.009*	0.021***	−0.431***	0.329***
	(0.005)	(0.002)	(0.018)	(0.016)
Meta*Video	−0.271***	0.144***	4.710***	−4.273***
	(0.011)	(0.005)	(0.040)	(0.030)
Constant	1.233***	0.626***	−1.979***	−0.011
	(0.024)	(0.010)	(0.156)	(0.112)
Candidate random intercepts	Yes	Yes	Yes	Yes
Observations	104,077	104,039	104,077	104,077
Log Likelihood	−92,279.880	−6,124.018	−233,712.200	−226,015.100
Akaike Inf. Crit.	184,573.800	12,262.040	467,436.400	452,042.200
Bayesian Inf. Crit.	184,640.600	12,328.900	467,493.700	452,099.500

Note. *$p < .1$; **$p < .05$; ***$p < .01$.

Acknowledgement

We would like to acknowledge Sebastian Zimmeck and Saray Shai for their collaboration and respective roles in helping to secure additional funding, and we especially thank Breeze Floyd and Colleen Bogucki for all of their effort along with numerous students on our Delta Lab computational and human coding teams who make the work possible.

Declaration of Conflicting Interests

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

Funding

The author(s) disclosed receipt of the following financial support for the research, authorship, and/or publication of this article: This research was made possible, in part, through Cross-Platform Election Advertising Transparency Initiative (CREATIVE) funded in part through a grant from the National Science Foundation (Award Numbers: 2235006, 2235007, 2235008). We also acknowledge funding from the John S. and James L. Knight Foundation (GR-2019-60044), the Democracy Fund (R-201903-03831, R-201912-05114), and Wesleyan University for the Wesleyan Media Project's 2020 tracking. The views presented here along with any errors are solely those of the authors.

ORCID iDs

Travis N. Ridout https://orcid.org/0000-0003-3268-4358
Erika Franklin Fowler https://orcid.org/0000-0003-3141-572X

Notes

1. Google search is somewhat different than the other three, which are social media platforms. Google search, rather, is a search engine and advertising platform. For the sake of simplicity, we use the term "platform" to refer to all four.
2. Other types of Google ads, such as non-YouTube mobile video ads, are not included in our sample.
3. We perform the train/test split at the advertiser level, not the ad level. This avoids having similar ads by the same advertiser in both the training set and test set, which would make the model look better than it is. The downside of this is that advertisers vary greatly in the number of ads they have, which might lead to widely differing test set sizes based on which advertisers are selected. Therefore, we select the test set by sampling from advertisers with a medium number of ads (500–1000), reducing variance in its size. Furthermore, to ensure that model evaluation works equally well for Democrats and Republicans, the sample for the test set included 10 advertisers from each party.
4. We concatenated ad creative body (ad content for Google), ad creative link description, ad creative link title (ad title for Google), OCR, and ASR and omitted the other fields since those often contain the names of the advertisers, which often provides clues about the party and would therefore provide information to the model that it shouldn't have.
5. Intercoder reliability on this task was high both for the number of candidates identified (alpha = 0.80) along with which individual candidates were identified.
6. Intercoder reliability was assessed by undergraduate research assistants who double coded a random sample of 462 of the Facebook ads. The Krippendorff's alphas are 0.7 for persuasion, 0.9 for fundraising, 0.6 for mobilization, and 0.9 for learning.

References

Ballard, A. O., Hillygus, D. S., & Konitzer, T. (2016). Campaigning online: Web display ads in the 2012 presidential campaign. *PS: Political Science & Politics*, *49*(3), 414–419. https://doi.org/10.1017/S1049096516000780

Bennett, C. J., & Gordon, J. (2021). Understanding the "micro" in political micro-targeting: An analysis of Facebook digital advertising in the 2019 Federal Canadian election. *Canadian Journal of Communication*, *46*(3), 431–459. https://doi.org/10.22230/cjc.2021v46n3a3815

Bode, L., & Vraga, E. K. (2018). Studying politics across media. *Political Communication*, *35*(1), 1–7. https://doi.org/10.1080/10584609.2017.1334730

Borgesius, F. J. Z., Möller, J., Kruikemeier, S., Fathaigh, R. Ó., Irion, K., Dobber, T., Bodo, B., & De Vreese, C. (2018). Online political microtargeting: Promises and threats for democracy. *Utrecht Law Review*, *14*(1), 82–96. https://doi.org/10.18352/ulr.420

Bossetta, M. (2018). The digital architectures of social media: Comparing political campaigning on Facebook, Twitter, Instagram, and Snapchat in the 2016 US election. *Journalism & Mass Communication Quarterly*, *95*(2), 471–496. https://doi.org/10.1177/1077699018763307

Bossetta, M., & Schmøkel, R. (2023). Cross-platform emotions and audience engagement in social media political campaigning: Comparing candidates' Facebook and Instagram images in the 2020 US election. *Political Communication*, *40*(1), 48–68. https://doi.org/10.1080/10584609.2022.2128949

Boulianne, S., & Larsson, A. O. (2023). Engagement with candidate posts on Twitter, Instagram, and Facebook during the 2019 election. *New Media & Society*, *25*(1), 119–140. https://doi.org/10.1177/14614448211009504

Dommett, K. (2019). The rise of online political advertising. *Political Insight*, *10*(4), 12–15. https://doi.org/10.1177/2041905819891366

Farkas, X., & Bene, M. (2021). Images, politicians, and social media: Patterns and effects of politicians' image-based political communication strategies on social media. *The International Journal of Press/Politics*, *26*(1), 119–142. https://doi.org/10.1177/1940161220959553

Fowler, E. F., Franz, M., & Ridout, T. (2021a). *Political advertising in the United States*. Routledge.

Fowler, E. F., Franz, M. M., Martin, G. J., Peskowitz, Z., & Ridout, T. N. (2021b). Political advertising online and offline. *American Political Science Review*, *115*(1), 130–149. https://doi.org/10.1017/s0003055420000696

Geer, J. G. (2008). *In defense of negativity: Attack ads in presidential campaigns*. University of Chicago Press.

Green, D. P., Palmquist, B., & Schickler, E. (2004). *Partisan hearts and minds: Political parties and the social identities of voters*. Yale University Press.

Gulati, J., & Williams, C. B. (2021). Digital media expenditures in presidential campaigns, 2008–2020. In T. Towner & J. Baumgartner (Eds.), *The Internet and the 2020 campaign* (pp. 25–48). Rowman & Littlefield.

Jackson, S., Stromer-Galley, J., & Hemsley, J. (2020). Differentiated information flows: Social media curation practices in the 2016 US presidential elections. *International Journal of Communication*, *14*, 22.

Kelm, O. (2020). Why do politicians use Facebook and Twitter the way they do? The influence of perceived audience expectations. *Studies in Communication and Media*, *9*(1), 8–34. https://doi.org/10.5771/2192-4007-2020-1-8

Kreiss, D. (2016). Seizing the moment: The presidential campaigns' use of Twitter during the 2012 electoral cycle. *New Media & Society*, *18*(8), 1473–1490. https://doi.org/10.1177/1461444814562445

Kreiss, D., Lawrence, R. G., & McGregor, S. C. (2018). In their own words: Political practitioner accounts of candidates, audiences, affordances, genres, and timing in strategic social media use. *Political Communication*, *35*(1), 8–31. https://doi.org/10.1080/10584609.2017.1334727

Krupnikov, Y. (2011). When does negativity demobilize? Tracing the conditional effect of negative campaigning on voter turnout. *American Journal of Political Science*, *55*(4), 797–813. https://doi.org/10.1111/j.1540-5907.2011.00522.x

Kruschinski, S., Haßler, J., Jost, P., & Sülflow, M. (2022). Posting or advertising? How political parties adapt their messaging strategies to Facebook's organic and paid media affordances. *Journal of Political Marketing*, 1–21. https://doi.org/10.1080/15377857.2022.2110352

Martin, P. S. (2004). Inside the black box of negative campaign effects: Three reasons why negative campaigns mobilize. *Political Psychology*, *25*(4), 545–562. https://doi.org/10.1111/j.1467-9221.2004.00386.x

Matassi, M., & Boczkowski, P. J. (2023). *To know is to compare: Studying social media across nations, media, and platforms*. MIT Press.

Matthes, J., Hirsch, M., Stubenvoll, M., Binder, A., Kruikemeier, S., Lecheler, S., & Otto, L. (2022). Understanding the democratic role of perceived online political micro-targeting: Longitudinal effects on trust in democracy and political interest. *Journal of Information Technology & Politics*, *19*(4), 435–448. https://doi.org/10.1080/19331681.2021.2016542

Neiheisel, J. R., & Niebler, S. (2013). The use of party brand labels in congressional election campaigns. *Legislative Studies Quarterly*, *38*(3), 377–403. https://doi.org/10.1111/lsq.12019

Pew Research Center. (2021). *Social media fact sheet*. https://www.pewresearch.org/internet/fact-sheet/social-media/#panel-c14683cb-c4f4-41d0-a635-52c4eeae0245

Ridout, T. N., Fowler, E. F., & Franz, M. M. (2021a). Spending fast and furious: Political advertising in 2020. *The Forum*, *18*(4), 465–492. https://doi.org/10.1515/for-2020-2109

Ridout, T. N., Fowler, E. F., & Franz, M. M. (2021b). The influence of goals and timing: How campaigns deploy ads on Facebook. *Journal of Information Technology & Politics*, *18*(3), 293–309. https://doi.org/10.1080/19331681.2021.1874585

Ronzhyn, A., Cardenal, A. S., & Batlle Rubio, A. (2022). *Defining affordances in social media research: A literature review*. New Media & Society. https://doi.org/10.1177/14614448221135187

Simchon, A., Brady, W. J., & Van Bavel, J. J. (2022). Troll and divide: The language of online polarization. *PNAS Nexus*, *1*(1), pgac019. https://doi.org/10.1093/pnasnexus/pgac019

Soroka, S. N. (2014). *Negativity in democratic politics: Causes and consequences*. Cambridge University Press.

Stier, S., Bleier, A., Lietz, H., & Strohmaier, M. (2018). Election campaigning on social media: Politicians, audiences, and the mediation of political communication on Facebook and Twitter. In E. K. Vraga & L. Bode (Eds.), *Studying politics across media* (pp. 50–74). Routledge.

Stromer-Galley, J., Rossini, P., Hemsley, J., Bolden, S. E., & McKernan, B. (2021). Political messaging over time: A comparison of US presidential candidate Facebook posts and tweets in 2016 and 2020. *Social Media + Society*, *7*(4), Article 205630512110634. https://doi.org/10.1177/20563051211063465

Zhang, F., Tanupabrungsun, S., Hemsley, J., Robinson, J. L., Semaan, B., Bryant, L., Stromer-Galley, J., Boichak, O., & Hegde, Y. (2017). Strategic temporality on social media during the general election of the 2016 US presidential campaign. In Proceedings of the 8th international conference on social media & society (pp. 1–10). ACM.

Author Biographies

Travis N. Ridout is Thomas S. Foley Distinguished Professor in the School of Politics, Philosophy and Public Affairs at Washington state University. He also serves as co-director of the Wesleyan Media Project. His research focuses on political communication, including political campaigns and advertising.

Markus Neumann is an Assistant Professor of Political Science and Computational Social Science at Duke Kunshan University. He has a Ph.D. in Political Science and Social Data

Analytics from Penn State, and spent three years as a Postdoc at the Wesleyan Media Project. His research revolves around the application of machine learning and especially natural language processing methods to social science problems, in particular in political advertising.

Jielu Yao was a postdoctoral research fellow at the Wesleyan Media Project. Her research interests center on political methodology and American politics, including applying supervised, unsupervised, and deep learning techniques to nontraditional data such as text, image/video, and audio.

Laura Baum was formerly Associate Director of the Wesleyan Media Project at Wesleyan University.

Michael M. Franz is professor of Government and Legal Studies at Bowdoin College in Brunswick, Maine. He is also co-director of the Wesleyan Media Project. His research interests include political advertising, campaign finance, and interest groups.

Pavel Oleinikov is associate director of the Quantitative Analysis Center at Wesleyan University, an assistant professor of the practice, and an affiliate of the Wesleyan Media Project. He uses his background in social and natural sciences to advance the application of quantitative methods to data from the social world.

Erika Franklin Fowler is professor of Government at Wesleyan University and co-director of the Wesleyan Media Project. Her research examines the volume, content, and effect of media messaging in electoral and health policy contexts.