Donate To Help Us Fight Back: Mobilization Rhetoric in Political Fundraising

Seo-young Silvia Kim[*†], Jan Zilinsky[‡], and Brian Brew[‡]

[†]American University
[‡]Technical University of Munich
[§]University of North Carolina, Chapel Hill

Abstract

How do campaigns differentially target donors and voters? We show that fundraising messages are an important class of electoral persuasion that reveals how campaigns perceive and target their "financial electorate." Because candidates' voters and donors can differ significantly, we theorize that rhetoric is chosen strategically for the target audience. Using data from the Facebook Ad Library for U.S. congressional candidates in the 2020 general election, we distinguish ads by persuasion targets. Then we use text analysis to test whether donor-targeting messages are, on average, more toxic, negative, and likely to reference a polarizing political figurehead (Donald Trump). While these expectations were largely borne out, there was significant variation by party and chamber. For example, Republican House candidates' appeals were more toxic than Democrats' and even more so when soliciting money. As the scramble for donations intensifies, these characteristics of appeals for cash may further polarize the electorate.

Electoral persuasion takes many forms, but the literature's primary interest has been in persuasion that affects voters' turnout and vote choice (Brader 2005; Arceneaux 2007; Huber and Arceneaux 2007; Gerber et al. 2011; Bailey et al. 2016; Kalla and Broockman 2018). In contrast, persuasion rhetoric that targets *campaign donors* has not received much attention relative to its importance, although the financial resources required to run a campaign are increasing with every election, and campaign chests enable the persuasion and mobilization of voters.

While the sender of messages that target voters and donors may be the same, these messages' intended recipients are very different. Direct donation solicitations do not necessarily explicitly

[*]Kim is an Assistant Professor of Government at American University, and is the corresponding author (email: sskim.research@gmail.com). Zilinsky is a post-doctoral research fellow at the School of Social Science and Technology at the Technical University of Munich. Brew is a Ph.D. student at the the Department of Political Science, University of North Carolina, Chapel Hill. An earlier version of this paper was presented at American University and at the MPSA 2022 annual conference. We thank Andy Ballard, Marcus Johnson, Clifford W. Brown, and Maggie McDonald for comments.

target voters who can cast a ballot for the given candidate. The financial electorate,[1] as opposed to the voting electorate, falls under the "monetary surrogacy" representation (Mansbridge 2003). Donors can be non-constituents as long as they satisfy the legal constraints; for example, Gimpel et al. (2008) show that almost two-thirds of congressional campaign cash flows from out-district. What's more, unlike the one-person-one-vote principle, the impact of monetary contributions can vary widely depending on the donor. And in general, donors are more resourced and ideological compared to voters (Verba et al. 1995; Francia et al. 2003; Hill and Huber 2017; Carey et al. 2022).

Given these features of the political environment, we ask the following: how do campaigns differentially target donors and voters? This paper fills the gap in the literature using data from the Facebook Ad Library for U.S. congressional candidates in the 2020 general election. This dataset provides a transparent and comprehensive set of online advertisements that target more broadly than direct mail or email appeals. Our primary motive is to show that persuasion that targets donors systematically differs from those targeting voters. Using text analysis, we find that donor-targeting messages are typically more toxic, negative, and likely to reference a polarizing political figurehead (Donald Trump), with some variations by party and chamber.

We argue that as the scramble for donations intensifies, a flood of negative sentiment, toxicity, and polarizing tactics in donor-targeting ads—especially online, where they are less regulated—may further polarize the electorate. These are instances of elite political communication that influence public agenda and mass political behavior (Zaller 1992; Lewandowsky et al. 2020), and the differences we analyze have important implications for mass behavior and the discursive environment on digital platforms.

1 Strategic Campaign Messages for Donations

American elections have always been comparatively expensive, but in recent years, the rate of increase in spending has been particularly striking. Congressional campaigns spent more than $8.7 billion in the 2020 elections (Evers-Hillstrom 2021) which is twice that of 2016 at $4.1 billion. In addition, the Center for Responsive Politics estimated that 4.7 million Americans donated to federal campaigns in 2020, compared to just 1.7 million in 2016.[2] All in all, campaigns are desperately trying to reach and persuade individual donors to keep their apparatus running.

However, there has been relatively little research on how politicians try to solicit money from potential donors although candidates' messages are often the primary driver of donations (Magleby et al. 2018). Some work has examined how politicians differentiate across donors. Hassell (2011), for example, shows that campaigns recognize that viable financial electorates differ between the primary and general stages, and they tailor their appeals accordingly. In addition, Hassell and Monson (2014) and Gaynor and Gimpel (2021) show that direct appeals to donors often specif-

[1] Hill and Huber (2017) uses the term 'donorate.'

[2] https://www.opensecrets.org/elections-overview/donor-demographics

ically target previous or frequent donors, rather than casting a broad net. Fowler et al. (2021) examine differences between political ads on TV versus online, finding that the former is comparatively less negative but more partisan. The only paper that, to our knowledge, mentions how donor- and voter-targeting rhetoric may differ is Hassell and Oeltjenbruns (2016), which shows that negative rhetoric is more prevalent in campaign emails with donation requests, controlling for characteristics such as election dynamics, incumbency, predicted electability. But how about in cases where campaigns target more broadly, as opposed to messages that are concentrated on known supporters?

Strong emotions such as anger are used strategically by campaigns and can mobilize, sustain partisan solidarity, and trigger information-seeking behavior (Marcus et al. 2000; Brader 2006; Druckman and McDermott 2008; MacKuen et al. 2010; Valentino et al. 2011; Webster 2020; Webster and Albertson 2022). We also know that donors are ideologically sophisticated and more extreme (Barber et al. 2017; Bafumi and Herron 2010; Hill and Huber 2017) even when compared to partisans. Combined with the fact that stronger mobilization of donors may lead to higher amounts donated, we expect campaigns to draft highly partisan messages that are both highly negative and polarizing so that they grab attention and stimulate the financial electorate.

Given these factors, we test three hypotheses regarding the distinction between persuasion rhetoric targeting donors and voters. Based on the understanding that politicians are strategic players who may employ negative and polarizing rhetoric for greater engagement (Hassell and Oeltjenbruns 2016; Ballard et al. 2022), we hypothesize that appeals for donations will (1) contain higher levels of negative sentiments such as anger, (2) contain higher levels of toxicity, and (3) will refer more to Donald Trump, who was a key polarizing political figurehead at the time. We also investigate themes and topics associated with appeals for donors and voters.

2 Data and Methods

The dataset consists of advertisements that politicians fielded on Facebook (now Meta). We gathered this data using the Facebook Ad Library, a public archive of all ads run on the platform. This dataset provides a wealth of opportunity in that it is not necessarily targeted toward existing donors, unlike those ads fielded through emails. It also provides transparent and tractable metadata, such as target state-level geography/demographics, date, reach, and amount spent. Moreover, U.S. residents are inadvertently exposed to Facebook ads—which are injected into their timelines—in contrast to emails for which individuals may have voluntarily signed up.

We downloaded all ads for candidates in the U.S. congressional general elections in the 2020 cycle.[3] Although citizens and activists may also initiate persuasion (Mutz et al. 1996) by setting up fundraisers independently from the campaign personnel, we restrict the data to official campaign

[3] We define the 2020 cycle as all ads between January 1, 2019, and December 31, 2020. We leave out Senate candidates not up for election.

messages. We also drop independent and third-party candidates from the analysis, limiting our sample to Democrats and Republicans.

Because politicians field the same ad content across different target geographies, demographics, and times, we deduplicate the ads by the candidate and unique content and summarise targets by averaging across the same ads. This leaves us with 26,113 Senate ads and 43,866 House ads. The number of unique ads for each candidate varies dramatically across individuals, with some running no ads whatsoever and others running hundreds of unique ads.[4]

Text-as-data methods can be applied to uncover patterns in text data, quantifying aspects of speech and scaling actors on theoretically relevant dimensions. They have a variety of applications in social sciences (Grimmer and Stewart 2013; Gentzkow et al. 2019; Grimmer et al. 2022). In this paper, we use sentiment analysis and the detection of toxicity using the Google Perspective API (Rieder and Skop 2021),[5]

We first classified ads that target donors (which we will call donor-targeting ads) and voters (voter-targeting ads) by (1) detecting keywords such as "chip in," "pitch in," or "donate," as well as (2) identifying whether the ad explicitly embeds links from fundraising platforms such as ActBlue and WinRed (Kim 2022; Kim and Li 2022).[6] Donor-targeting ads contain text such as "Send a message to Mike Pence by helping us flip a seat from red to blue in his own backyard. Chip in today to help Christina Hale win in November." On the other hand, voter-targeting ads contain text such as "My opponent Marcy Kaptur has been in office for 37 years. Public Office is supposed to be about service not a career. It is time for #TermLimits!"

To verify that voter-targeting ads do target the voting electorate, we first analyze the proportion of such ads that target in-state donors and prospective voters. Figure 1 shows the proportion of ads targeting in-state Facebook users. As should be the case, across both chambers and both parties, voter-targeting ads have a much greater probability of targeting in-state users, indicating that when candidates do target ads to out-of-state Americans, they are much more likely to be actively courting donations. On average, donor-targeting ads target in-state users at 62.2%, while voter-targeting ads target them at 91.3%. Given that the eligible electorate might be temporarily living out-district, we do not further restrict by geographic targets.

3 Toxicity, Negative Sentiment, and References to Trump

Figure 2 displays differences in toxicity across advertisement types. We see that for House candidates, toxicity is higher in donor-targeting ads for both parties (5.7% and 6.0% increase re-

[4]For example, Alexandria Ocasio-Cortez (D, NY-14)'s campaign featured 510 unique ads which were cumulatively run 48,171 times.

[5]Perspective API is a model trained to classify online speech, focused on detection of toxic and abusive comments. Toxicity is defined as "rude, disrespectful, and or unreasonable," and the classifier is trained on evaluations from human subjects.

[6]For full steps of the rule-based classification, see Supporting Information.

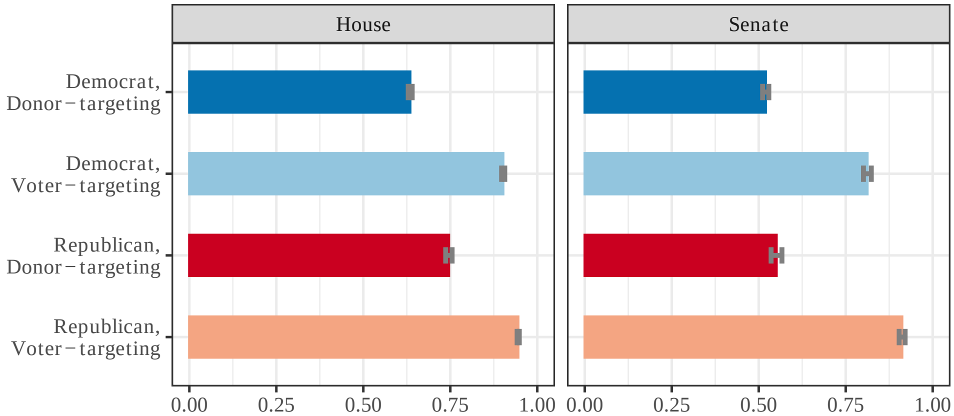

Figure 1: Proportion of In-state Targeting by Type of Facebook Ads

spectively for Democrats and Republicans compared to voter-targeting ads), but not for Senate candidates. In the U.S. House, we observe a clear hierarchy by party and target: toxicity is high in the order of Republican donor-targeting ads (13.5%), Republican voter-targeting ads (12.7%), Democratic donor-targeting ads (11.3%), and Democratic voter-targeting ads (10.6%). Republican ads are overall more toxic than Democratic ads for Senate candidates as well. Another way to put it is that within donor- or voter-targeting House ads, Republican ads are respectively 19.9% and 19.6% more toxic than Democratic ads. For Senate, the numbers are 11.6% and 21.8%.

OLS regressions that predict toxicity while conditioning on characteristics such as incumbency or campaign dynamics confirm that, on average, Republican-sponsored ads contain higher levels of toxic language (Table 1(a)). In addition, a within-candidate model reveals that donor-targeting ads are more toxic (Table 1(b)).[7][8]

An analysis of the presence of emotionally charged words, summarized in Figure 3, shows that in both chambers/parties, words associated with anger are more prevalent in donor-targeting ads compared to voter-targeting ads, with the 95% confidence interval also displayed. The evidence is more mixed for disgust and fear, which are two other prominent negative emotions. For House candidates, they are more likely to be used in donor-targeting ads, but such differences are not observed among Senate candidates.

[7] Note that 118 candidates (16.4%) only had either voter-targeting or donor-targeting ads, so they did not impact the model results.

[8] To check if particular candidates were driving the results, we also ran leave-one-out fixed effects models at the candidate level. For most of the 720 candidates, statistical significance still held if the candidate was excluded from the dataset, except for four: Raphael Warnock (Democratic Senator, Georgia), Jaime Harrison (Democrat, unsuccessful Senate challenger to Lindsey Graham, South Carolina), Rishi Kumar (Democrat, unsuccessful House challenger to Anna Eshoo (Democrat) for CA-18), Ammar Campa-Najjar (Democrat, defeated in an open-seat race by Darrell Issa (Republican) for CA-50). This seems to be driven by the large number of unique ads that these candidates run, relative to some Republican candidates, who had higher levels of relative toxicity in donor-targeting ads but ran smaller number of ads. However, a within-candidate regression with candidate/type average-summarized model was still borderline significant ($p < 0.1$)

Dependent Variable:	Toxicity
Republican	0.0211***
	(0.0032)
Donor-targeting	0.0055**
	(0.0027)
Senate	0.0047
	(0.0029)
Incumbent	-0.0027
	(0.0026)
Open seat	0.0029
	(0.0063)
Electoral safety	5.94×10^{-5}
	(0.0001)
Male	-0.0038
	(0.0031)
First ad delivery date	1.83×10^{-5} ***
	(6.29×10^{-6})
Republican × Donor-targeting	-0.0013
	(0.0038)
(Intercept/state dummies excluded for brevity)	
Fit statistics	
Observations	61,007
R^2	0.02776
Adjusted R^2	0.02684

Clustered (candidate) standard-errors in parentheses

(a) Clustered S.E.

Dependent Variable:	Toxicity
Donor-targeting	0.0043**
	(0.0020)
First ad delivery date	-1.71×10^{-5} ***
	(6.48×10^{-6})
Fixed-effects	
Candidate	Yes
Fit statistics	
Observations	61,007
R^2	0.09881
Within R^2	0.00114

Clustered (candidate) standard-errors in parentheses

(b) Candidate Fixed Effects

Table 1: Predicting Toxicity via Simple Linear Regression

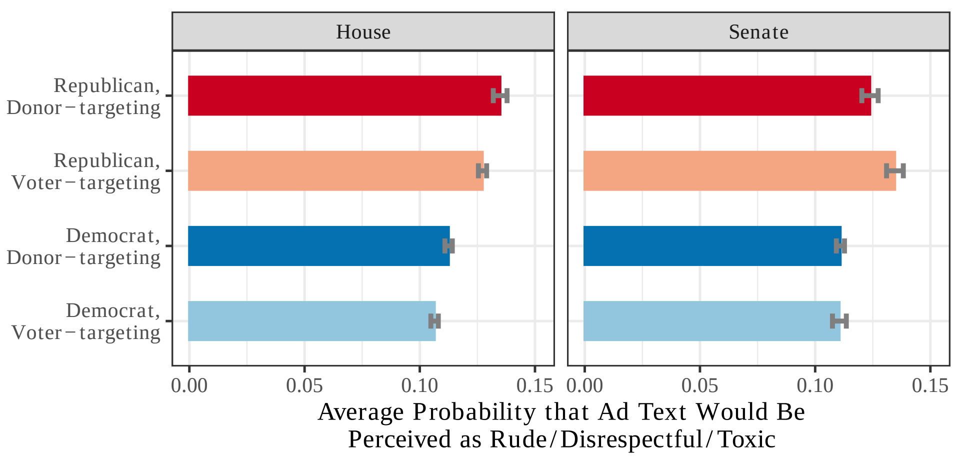

Figure 2: Level of Language Toxicity by Facebook Ad Type

Next, we examine the prevalence for references to Donald Trump. Trump, the incumbent president at the time when these ads were fielded, was the most prominent polarizing political figurehead. Is there a systematic difference in which types of ads mention Trump? Figure 4 shows that for both chambers and parties, donor-targeting ads are significantly more likely to reference Trump than voter-targeting ads—roughly one out of five ads that target donors mention Trump. There is also some variation by chamber and party: for Senate ads, there is a clear hierarchy in

(a) Words associated with anger

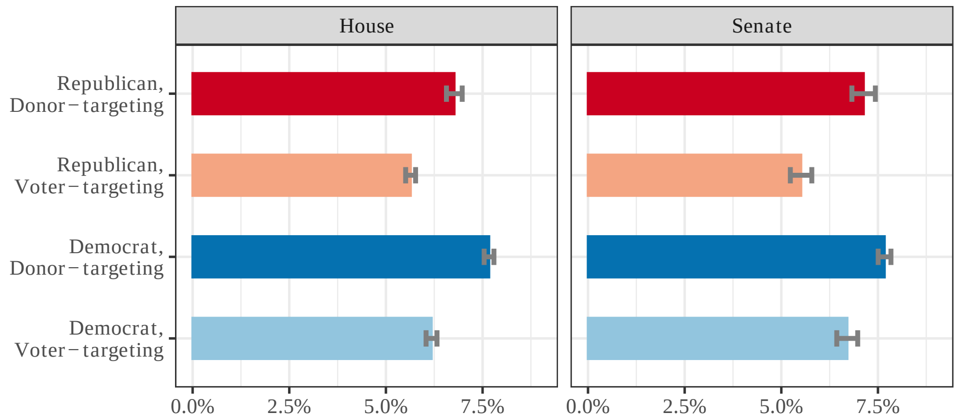

(b) Words associated with disgust

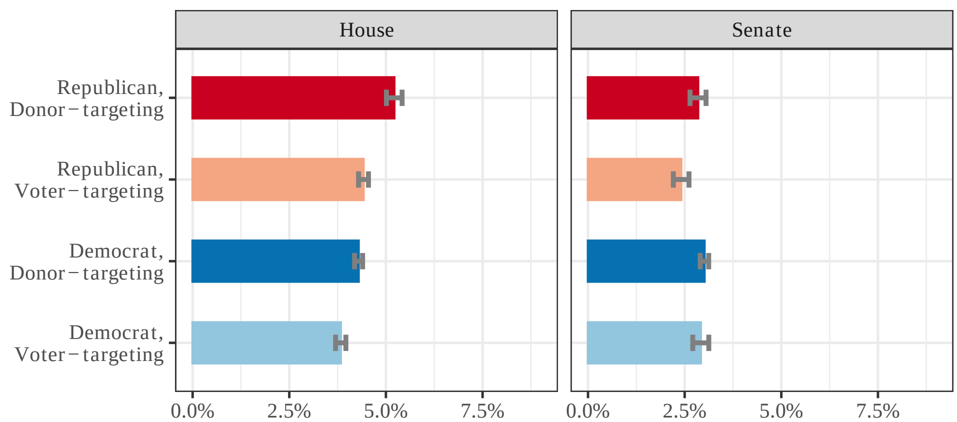

(c) Words associated with fear

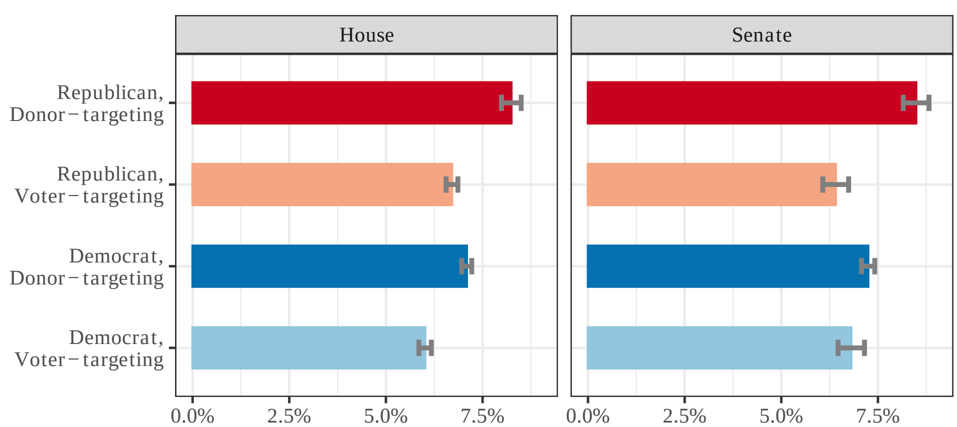

Figure 3: Presence of Words Associated with Negative Emotions by Types of Facebook Ads

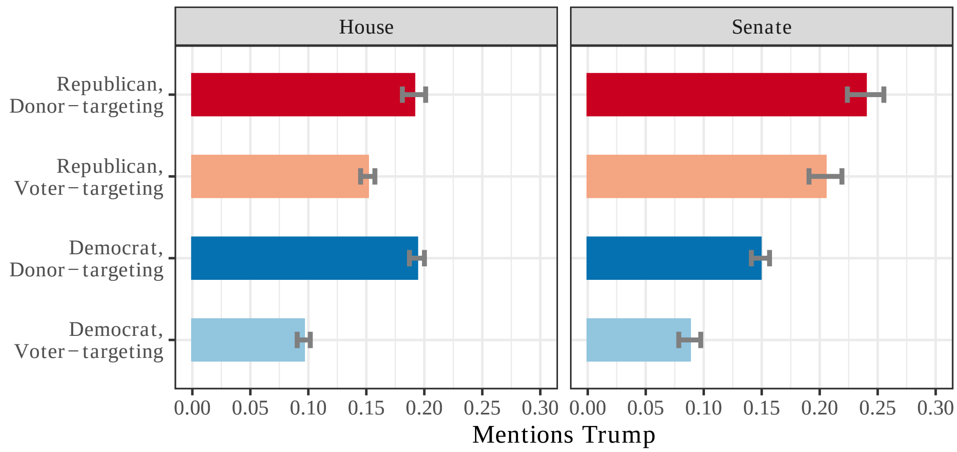

Figure 4: Mention of Trump by Facebook Ad Type

the order of Republican donor-targeting ads, Republican voter-targeting ads, Democratic donor-targeting ads, and Democratic voter-targeting ads. Within House donor-targeting ads there is not much difference between the two parties, but Republican candidates exhibit a greater propensity to feature Trump in their voter-targeting ads compared to Democratic candidates.

4 Conclusion

This paper has established important insights on campaign-driven persuasion rhetoric that targets voters and donors. We have shown that donor-targeting messages are, on average, more negative, toxic, and likely to reference a polarizing political figurehead—in this case, Donald Trump. Moreover, compared to voter-targeting messages, candidates rely less on substantive policy issues when they target potential donors. In addition, we document significant variation by party and chamber; specifically, Republican candidates' language was more toxic than Democrats'. This is consistent with the academic assessment of GOP's recent turn towards extremism (Skocpol 2020).

The next natural question is the following: does this matter? Do differences in rhetoric for the financial electorate create differences in political behavior? Are specific types of appeal more or less effective among different groups of donors? For example, Haenschen (2022) shows experimentally that Facebook ads did not have strong turnout effects.

While we do not have a definite answer, we believe that ads, as another type of elite political communication, are likely to have long-term repercussions. Electoral persuasion effects can be moderated by audience (Suhay et al. 2020), and the pouring solicitations for money can serve as another kind of "partisan media" that can further polarize citizens and make governing difficult (Levendusky 2013). This is especially true because online ads can target more broadly and reach more than just habitual donors. In addition, since small-donor-based strategies have proven to be somewhat viable (Alvarez et al. 2020), it is likely that campaigns will increasingly douse the average American voter with polarizing appeals for donations. Note that small dollars online were already flowing more towards polarizing candidates (Karpf 2013).

To be sure, harsh language may invigorate democratic debate or promote engagement (Schudson 1997; Sydnor 2019), and name-calling and insults are generally viewed as a milder form of incivility (Sobieraj and Berry 2011) compared to outright hate speech (Siegel 2020). However, disrespectful discourse may silence or demobilize citizens, or accelerate democratic backsliding (Kalmoe 2014; Jamieson et al. 2017; Finkel et al. 2020). Although here we do not identify causal effects of political ads with toxic language, existing works suggest that our findings are consistent with the view that U.S. political elites are chipping away at the quality of democracy.

These findings may hold concerning implications for democratic backsliding in the U.S.. Recent analyses have shown that surprisingly many Americans are willing to trade democratic principles for conflicting considerations such as partisan loyalties (Graham and Svolik 2020). Recent events, such as the attack on the Capitol on January 6, 2021, have made it clear that American democracy is not as secure as was previously believed. The content of the candidates' advertisements serves as a reminder that rancor and hostility are becoming increasingly normalized as part and parcel of political competition in the United States—and particularly because online ads are less regulated than traditional TV ads (Fowler et al. 2021).

References

Alvarez, R. Michael, Jonathan N. Katz, and Seo-young Silvia Kim (2020). Hidden Donors: The Censoring Problem in U.S. Federal Campaign Finance Data. *Election Law Journal: Rules, Politics, and Policy 19*(1), 1–18.

Arceneaux, Kevin (2007). I'm Asking for Your Support: The Effects of Personally Delivered Campaign Messages on Voting Decisions and Opinion Formation. *Quarterly Journal of Political Science 2*, 43.

Bafumi, Joseph and Michael C. Herron (2010). Leapfrog Representation and Extremism: A Study of American Voters and Their Members in Congress. *American Political Science Review 104*(3), 519–542.

Bailey, Michael A., Daniel J. Hopkins, and Todd Rogers (2016). Unresponsive and Unpersuaded: The Unintended Consequences of a Voter Persuasion Effort. *Political Behavior 38*(3), 713–746.

Ballard, Andrew O., Ryan DeTamble, Spencer Dorsey, Michael Heseltine, and Marcus Johnson (2022). Dynamics of Polarizing Rhetoric in Congressional Tweets. *Legislative Studies Quarterly 0*(0).

Barber, Michael J., Brandice Canes-Wrone, and Sharece Thrower (2017). Ideologically Sophisticated Donors: Which Candidates Do Individual Contributors Finance? *American Journal of Political Science 61*(2), 271–288.

Brader, Ted (2005). Striking a Responsive Chord: How Political Ads Motivate and Persuade Voters by Appealing to Emotions. *American Journal of Political Science 49*(2), 388–405.

Brader, Ted (2006). *Campaigning for Hearts and Minds: How Emotional Appeals in Political Ads Work*. University of Chicago Press.

Carey, John, Katherine Clayton, Gretchen Helmke, Brendan Nyhan, Mitchell Sanders, and Susan Stokes (2022). Who will defend democracy? Evaluating tradeoffs in candidate support among partisan donors and voters. *Journal of Elections, Public Opinion and Parties 32*(1), 230–245.

Druckman, James N. and Rose McDermott (2008). Emotion and the Framing of Risky Choice. *Political Behavior 30*(3), 297–321.

Evers-Hillstrom, Karl (2021, February). Most expensive ever: 2020 election cost $14.4 billion. *OpenSecrets*.

Finkel, Eli J., Christopher A. Bail, Mina Cikara, Peter H. Ditto, Shanto Iyengar, Samara Klar, Lilliana Mason, Mary C. McGrath, Brendan Nyhan, David G. Rand, Linda J. Skitka, Joshua A. Tucker, Jay J. Van Bavel, Cynthia S. Wang, and James N. Druckman (2020). Political sectarianism in America. *Science 370*(6516), 533–536.

Fowler, Erika Franklin, Michael Franz, and Travis N. Ridout (2021, November). *Political Advertising in the United States* (2 ed.). New York: Routledge.

Fowler, Erika Franklin, Michael M. Franz, Gregory J. Martin, Zachary Peskowitz, and Travis N. Ridout (2021). Political Advertising Online and Offline. *American Political Science Review 115*(1), 130–149.

Francia, Peter L., John C. Green, Paul S. Herrnson, Clyde Wilcox, and Lynda W. Powell (2003). *The Financiers of Congressional Elections: Investors, Ideologues, and Intimates*. Columbia University Press.

Gaynor, SoRelle Wyckoff and James G. Gimpel (2021). Small Donor Contributions in Response to Email Outreach by a Political Campaign. *Journal of Political Marketing 0*(0), 1–25.

Gentzkow, Matthew, Bryan Kelly, and Matt Taddy (2019). Text as Data. *Journal of Economic Literature 57*(3), 535–574.

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

Gimpel, James G., Frances E. Lee, and Shanna Pearson-Merkowitz (2008). The Check Is in the

Mail: Interdistrict Funding Flows in Congressional Elections. *American Journal of Political Science 52*(2), 373–394.

Graham, Matthew H. and Milan W. Svolik (2020). Democracy in America? Partisanship, Polarization, and the Robustness of Support for Democracy in the United States. *American Political Science Review 114*(2), 392–409.

Grimmer, Justin, Margaret E. Roberts, and Brandon M. Stewart (2022). *Text as Data: A New Framework for Machine Learning and the Social Sciences*. Princeton University Press.

Grimmer, Justin and Brandon M. Stewart (2013). Text as Data: The Promise and Pitfalls of Automatic Content Analysis Methods for Political Texts. *Political Analysis 21*(3), 267–297.

Haenschen, Katherine (2022, March). The Conditional Effects of Microtargeted Facebook Advertisements on Voter Turnout. *Political Behavior*.

Hassell, Hans J. G. (2011). Looking Beyond the Voting Constituency: A Study of Campaign Donation Solicitations in the 2008 Presidential Primary and General Election. *Journal of Political Marketing 10*(1-2), 27–42.

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

Hassell, Hans J. G. and Kelly R. Oeltjenbruns (2016). When to Attack: The Trajectory of Congressional Campaign Negativity. *American Politics Research 44*(2), 222–246.

Hill, Seth J. and Gregory A. Huber (2017). Representativeness and Motivations of the Contemporary Donorate: Results from Merged Survey and Administrative Records. *Political Behavior 39*(1), 3–29.

Huber, Gregory A. and Kevin Arceneaux (2007). Identifying the Persuasive Effects of Presidential Advertising. *American Journal of Political Science 51*(4), 957–977.

Jamieson, K. H., A. Volinsky, I. Weitz, and K. Kenski (2017). The political uses and abuses of civility and incivility. In K. Kenski and K. H. Jamieson (Eds.), *The Oxford Handbook of Political Communication*, pp. 205–218. New York: Oxford University Press.

Kalla, Joshua L. and David E. Broockman (2018). The Minimal Persuasive Effects of Campaign Contact in General Elections: Evidence from 49 Field Experiments. *American Political Science Review 112*(1), 148–166.

Kalmoe, Nathan P. (2014). Fueling the Fire: Violent Metaphors, Trait Aggression, and Support for Political Violence. *Political Communication 31*(4), 545–563.

Karpf, David (2013). The Internet and American Political Campaigns. *The Forum 11*(3), 413–428.

Kim, Seo-young Silvia (2022). How Much To Ask? Platforms, Parties, and Suggested Amounts in Political Fundraising.

Kim, Seo-young Silvia and Zhao Li (2022). Keep Winning with WinRed? Digital Fundraising Platform as the Party's Public Good.

Levendusky, Matthew (2013). *How Partisan Media Polarize America*. University of Chicago Press.

Lewandowsky, Stephan, Michael Jetter, and Ullrich K. H. Ecker (2020). Using the president's tweets to understand political diversion in the age of social media. *Nature Communications 11*(1), 5764.

MacKuen, Michael, Jennifer Wolak, Luke Keele, and George E. Marcus (2010). Civic Engagements: Resolute Partisanship or Reflective Deliberation. *American Journal of Political Science 54*(2), 440–458.

Magleby, David B., Jay Goodliffe, and Joseph A. Olsen (2018). *Who Donates in Campaigns?: The Importance of Message, Messenger, Medium, and Structure*. Cambridge University Press.

Mansbridge, Jane (2003). Rethinking Representation. *American Political Science Review 97*(4), 515–528.

Marcus, George E., W. Russell Neuman, and Michael MacKuen (2000). *Affective Intelligence and Political Judgment*. University of Chicago Press.

Mutz, Diana Carole, Paul M. Sniderman, and Richard A. Brody (1996). *Political Persuasion and Attitude Change*. University of Michigan Press.

Rieder, Bernhard and Yarden Skop (2021). The fabrics of machine moderation: Studying the technical, normative, and organizational structure of Perspective API. *Big Data & Society 8*(2).

Schudson, Michael (1997). Why conversation is not the soul of democracy. *Critical Studies in Mass Communication 14*(4), 297–309.

Siegel, Alexandra (2020). Online hate speech. In J. A. Tucker and N. Persily (Eds.), *Social Media and Democracy*, pp. 56–88.

Skocpol, Theda (2020). The Elite and Popular Roots of Contemporary Republican Extremism. Upending American Politics, pp. 3 – 28. Oxford University Press.

Sobieraj, S. and J. Berry (2011). From incivility to outrage: Political discourse in blogs, talk radio, and cable news. *Political Communication 1*(28), 19–41.

Suhay, Elizabeth, Bernard Grofman, and Alexander H. Trechsel (2020). A Framework for the Study of Electoral Persuasion. In E. Suhay, B. Grofman, and A. H. Trechsel (Eds.), *The Oxford Handbook of Electoral Persuasion*, pp. xiv–25. Oxford University Press.

Sydnor, Emily (2019). *Disrespectful Democracy: The Psychology of Political Incivility*. Columbia University Press.

Valentino, Nicholas A., Ted Brader, Eric W. Groenendyk, Krysha Gregorowicz, and Vincent L. Hutchings (2011). Election Night's Alright for Fighting: The Role of Emotions in Political Participation. *The Journal of Politics 73*(1), 156–170.

Verba, Sidney, Kay Lehman Schlozman, and Henry E. Brady (1995). *Voice and Equality: Civic Voluntarism in American Politics*. Harvard University Press.

Webster, Steven W. (2020). *American Rage: How Anger Shapes Our Politics*. Cambridge University Press.

Webster, Steven W. and Bethany Albertson (2022). Emotion and Politics: Noncognitive Psychological Biases in Public Opinion. *Annual Review of Political Science 25*(1).

Zaller, John R. (1992). *The Nature and Origins of Mass Opinion*. Cambridge University Press.

Online Appendices

Appendix A Key Variables

Below are the key variables that are featured within our analysis:

- Anger, Fear, and Disgust: each of these sentiment variables was coded using the NRC's Emotions Lexicon.
- First Ad Delivery Date: the date when a given advertisement was first delivered on Facebook.
- Chamber: a dichotomous variable indicating whether a given candidate was running for the House or the Senate.
- Gender: an indicator of the candidate's gender identity, primarily coded using the `Kmisc` package's `gender_mutate_df` function. For more on the logistics of coding gender as a variable, see Blevins and Mullen (2015).
- Incumbency: a categorical variable, with 'Incumbent' indicating an incumbent candidate, 'Challenger' indicating a candidate who was challenging the incumbent, and 'Open' indicating that a candidate was running in a race without an incumbent on the ballot.
- Open Seat: a dichotomous variable, indicating whether or not there was an incumbent running in the race a candidate participated in.
- Party: a dichotomous variable indicating whether a candidate was a Democrat or a Republican; Independent candidates were not included in our analysis.
- Targeting: a dichotomous variable. Advertisements that contained links to fundraising pages were coded as donor-targeting, as were those advertisements that contained explicit fundraising appeals (e.g. "Chip in," "Donate today").
- Toxicity: A score from the Google Perspective API, the values of which fall between 0 to 1. The closer the score is to 1, the more likely it is that the advertisement is perceived as toxic. By way of example, an advertisement from incumbent Dan Bishop (R, NC-9) which stated "These crazy liberal clowns and what they believe are not funny, they are downright scary," received a toxicity score of 0.9, whereas an advertisement from incumbent Gil Cisneros (D, CA-39) which stated "Gil Cisneros is helping us and our economy deal with the COVID-19 pandemic – working to assist small businesses and people who have lost their jobs or wages," received a toxicity score of 0.02.
- Trump: a dichotomous variable, indicating whether or not a given advertisement mentioned Donald Trump by name.

Appendix B Building the Facebook Ads Data

The texts of advertisements, and key variables associated with them, were scraped from the Facebook Ad Library API using the `Radlibrary` package, developed by Fraser and Shank, in R. `Radlibrary` allows researchers with a Facebook API key to query the API and scrape data directly from R. This enables scholars to gather a host of variables on the advertisements, including but not limited to their text and any links they contained, how much money was spent upon them, when they were aired, what links they contained, where they were delivered, and a range of impressions of each advertisement.

First, we gathered a list of all the Ad Library ID numbers for candidates who ran as their party's general-election nominees in the 2020 House and Senate elections. These lists for House and Senate candidates served as the basis of the data-collection process. As `Radlibrary` has a limit of 5,000 advertisements or ten pages per query, and as some candidates had more than 5,000 advertisements in the 2020 campaign cycle, we queried each campaign page individually.

The Ad Library API provides different types of data for different queries. Ad-focused queries furnish information about a particular advertisement. The most important among these for our purposes were the advertisement's text, when it was aired, whether it contained any links, a lower and upper limit to the range of funds that may have been spent upon the advertisement, and a lower and upper limit to the range of possible impressions of the advertisement (the number of times the advertisement was viewed). Demographic-focused queries provide broad insight into the gender and racial groups that were targeted by an individual advertisement. Region-focused queries reveal the geographic regions where advertisements were distributed. These were almost exclusively U.S. states in our analysis, though a handful of ads were also aired in foreign countries.

After using `Radlibrary` to scrape the Ad, Demographic, and Region data for all those candidates who had campaign Facebook pages, we set about cleaning the data and adding in other relevant variables. Incumbency status and partisanship were both added in from outside sources. The advertisement texts were used to generate other variables, such as the presence of mentions of Donald Trump and whether the ads were primarily targeted towards voters or donors. The texts were also the data used to gain measures of advertisement toxicity using the Google Perspective API. Once we had the data cleaned and all the variables assembled, we rearranged it into several different datasets that were conducive to different forms of analysis. For example, one dataset was used primarily for the purposes of constructing a keyword-assisted topic model (see Appendix C for more on this process), while others were used to run regressions and other supplementary analyses and inspections (see Appendices D and E for some of these).

Appendix C Estimation of Keyword-assisted Topic Models

To gain further insight into the ways in ads for donors and voters differ, we employed a keyword-assisted topic model (keyATM, proposed by Eshima, Imai and Sasaki, 2021), a method which allows researchers to condition the model to search for topics containing pre-identified keywords, avoiding topic instability (a core drawback of topic models, see: Wilkerson and Casas, 2017). We identified nine topics on substantive policy areas and rhetorical subjects, and then ran keyATM on the corpus of Facebook ads to explore how candidates presented themselves to different targets.

A keyword-assisted topic model was estimated using the `keyATM` package in R (Eshima, Imai and Sasaki, 2021). Compared to unsupervised methods, one benefit of the Eshima, Imai and Sasaki (2021) model is that post-hoc topic labeling is avoided (and miscellaneous topics which organically emerge remain unlabelled, and there is thus no post-hoc topic assignment).

Before standard pre-processing steps (lower-casing, and stemming the words used in the Facebook ads) were applied, we generated nearly 30 bi-grams or tri-grams which we expected would occur in political speech and signal the theme of a social media post (e.g. `green new deal`, `right to choose`, `national security`, `student debt`). The full list of manually chosen items used to generate these n-grams is available on the paper's GitHub repository.

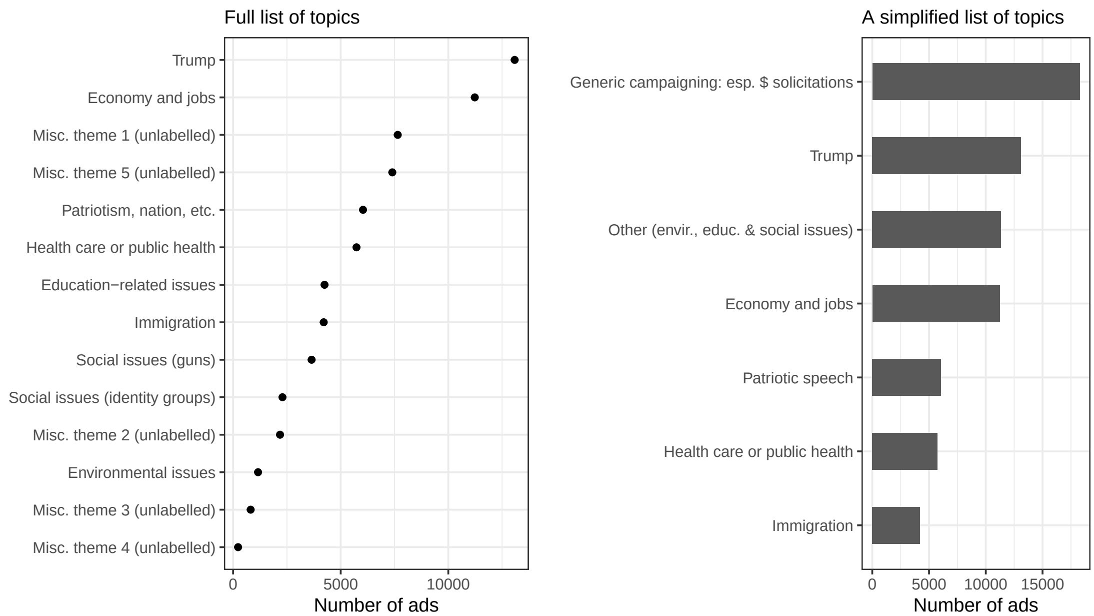

Figure 1: Topics of ads fielded by candidates for Congress in the 2020 electoral cycle estimated via the keyword-assisted topic model (Eshima, Imai and Sasaki, 2021).

Next, we pre-specificed keywords for 9 expected themes, namely:

- Economy-related ads (including those related to pocket-book issues – such as jobs,or wages – or the brader national economy)
- Posts centered around President Trump
- Posts with generic patriotic pronouncements
- Health-related messages (e.g. health care system, Covid, etc.)
- Immigration
- Environment and the climate change
- Education (e.g. school choice, student loans, etc.)
- (At least) Two topics related to cultural issues. We expected that one set of ads would relate to guns / safety and related issues; the second set encompasses the remaining cultural and social issues.

The top keywords for each of these topics are listed in Table 1, and the full list of seed words initially supplied by us is available on GitHub. We also specified that 5 miscellaneous topics (with no keywords specified ex-ante) were to be identified; we sometimes refer to theses ads falling into these 5 topics as "generic campaigning" and we verified that most of them are simple requests for monetary contributions. The term document matrix was not trimmed, except for the removal of stop-words.

The distribution of ads across the estimated topics is displayed in Figure 1.

Beyond the explicit mentions of the president, other themes and topics discussed by candidates' 2020 Facebook ads, are displayed in Figure 2, which presents the results from keyATM models.

Pre-defined topic	Keyword (stemmed)	Proportion (% of all words)	Word Count
Economy and jobs	work	0.502	11048
Economy and jobs	import	0.142	3117
Economy and jobs	job	0.141	3107
Economy and jobs	tax	0.121	2668
Economy and jobs	econom	0.053	1171
Health care or public health	health_care (bi-gram)	0.153	3376
Health care or public health	health	0.078	1720
Health care or public health	pandem	0.067	1475
Health care or public health	covid-19	0.062	1374
Health care or public health	coronavirus	0.037	817
Social issues (guns)	violenc	0.039	865
Social issues (guns)	nra	0.015	340
Social issues (guns)	shoot	0.005	119
Social issues (guns)	rifl	0.003	66
Social issues (guns)	shooter	0.001	20
Social issues (remaining)	reproduct	0.035	760
Social issues (remaining)	abort	0.030	662
Social issues (remaining)	parenthood	0.027	587
Social issues (remaining)	pro-choic	0.003	75
Social issues (remaining)	pregnanc	0.002	40
Immigration	border	0.062	1366
Immigration	immigr	0.049	1089
Immigration	dreamer	0.006	142
Immigration	migrant	0.002	55
Immigration	daca	0.002	50
Environmental issues	green_new_deal	0.035	764
Environmental issues	climat	0.030	669
Environmental issues	environ	0.026	577
Environmental issues	environment	0.015	325
Environmental issues	oil	0.012	256
Trump	trump	0.627	13804
Trump	presid	0.455	10016
Patriotism, nation, etc.	nation	0.209	4596
Patriotism, nation, etc.	america	0.179	3934
Patriotism, nation, etc.	patriot	0.014	311
Patriotism, nation, etc.	usa	0.003	62
Education-related issues	school	0.078	1713
Education-related issues	educ	0.070	1547
Education-related issues	student	0.042	915
Education-related issues	univers	0.034	746
Education-related issues	student_debt	0.006	134

Table 1: Top words associated with each of the 9 pre-selected topics. The full list of topics is displayed in the left panel of Figure 1.

Strikingly, nearly half of Republican financial ads at both chambers invoke Donald Trump or the executive, speaking to the then-president's prominence within inter-party politics during the 2020 election. Republican candidates clearly believed that leaning on Trump would encourage donors to contribute. Another trend present in both parties but more for Democrats is the following: financial ads were less likely to contain a substantive discussion of major wedge issues such as economy, healthcare, and immigration than non-financial ads. In part, this indicates that many financial ads principally consist of requests for donations, but it also suggests that many candidates place a higher priority on presenting substantive positions to prospective voters compared to those appeals that are directed towards the financial electorate.

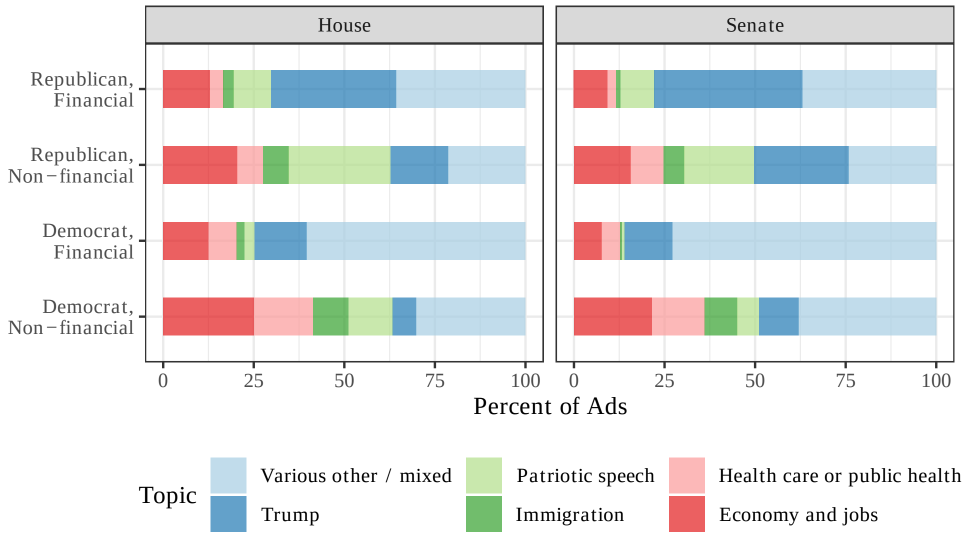

Figure 2: Themes and Topics in Ads Estimated with a Keyword-Assisted Topic Model by Facebook Ad Type

Appendix D Time-Series Plots of Facebook Data

Figure 3 shows the number of unique donor-targeting and voter-targeting advertisements that were created each week in 2019 and 2020 at the House and Senate levels. The marked increase in the number of unique Senate advertisements from late 2019 through early 2020 is largely a function of the fact that many senators ran for the Democratic nomination in 2020, and used their campaign Facebook pages to advertise their presidential primary bids.

Figure 4 shows the number of unique ads by party and financial status over time.

Figure 5 shows the total number of impressions for advertisements created in each month of 2019 and 2020. Figure 6 shows the total funds spent on Facebook advertisements in each month of 2019 and 2020.

In Figures 7 and 8, toxic ads are those to which the Google Perspective API assigned a score greater than 0.2, following the example of Muddiman, McGregor and Stroud (2019). These two figures contain every House advertisement, and every advertisement run by a 2020 Senate candidate (those senators who were not up for reelection in 2020 were excluded from these data). As both these figures make very clear, there were many more non-toxic ads, by this threshold, than there were toxic ads. 12.682% of House advertisements and 13.085% of Senate

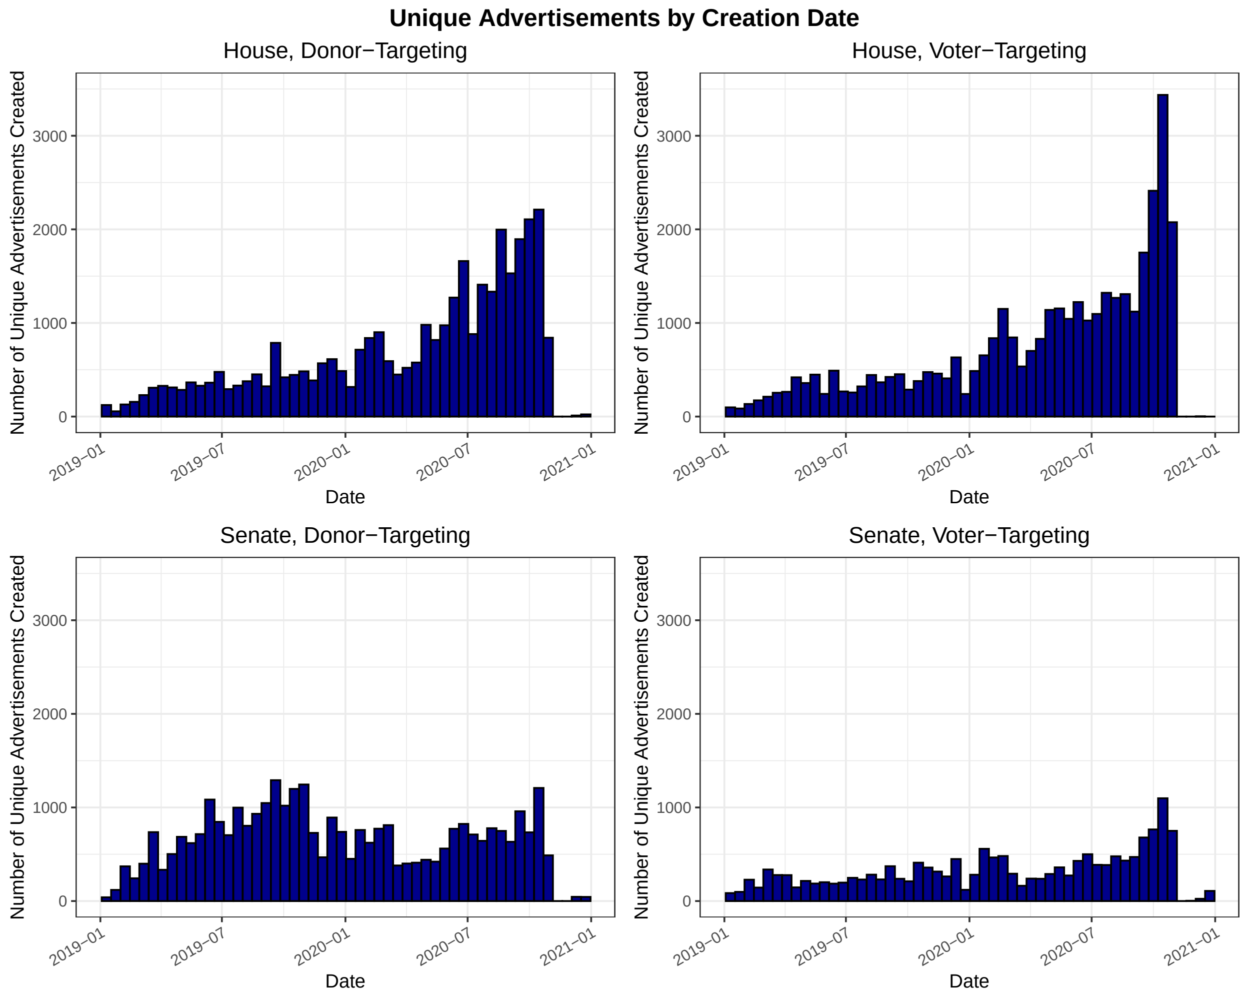

Figure 3: Unique Financial and Non-Financial House and Senate Advertisements by Creation Date

advertisements run by candidates in the 2020 cycle were toxic. However, in the aggregate, these advertisements received relatively more attention from campaigns and audiences alike, although not dramatically more so than nontoxic advertisements. Based on the information from the Facebook Ad Library API, the 12.682% of House advertisements that were toxic received an estimated 17.597% of all House advertisement funding, and these ads received an estimated 16.749% of all House advertisement impressions. The same trend holds in the Senate: the 13.085% of advertisement that were toxic received an estimated 15.716% of the funding, and 14.936% of all impressions. One limitation that must be addressed is that the Facebook Ad Library API does not provide the exact amount of money devoted to an individual advertisement, nor does it furnish the exact number of times an advertisement was viewed. As mentioned in Appendix B, the API returns the lower and upper limits of an estimated range for both of these variables. Because the Facebook Ad Library's maximum estimate for the reach of a single advertisement is 999,999, some upper estimates for the number of impressions were blank in the data we collected; accordingly, the estimates in Figure 8 are the lower estimates.

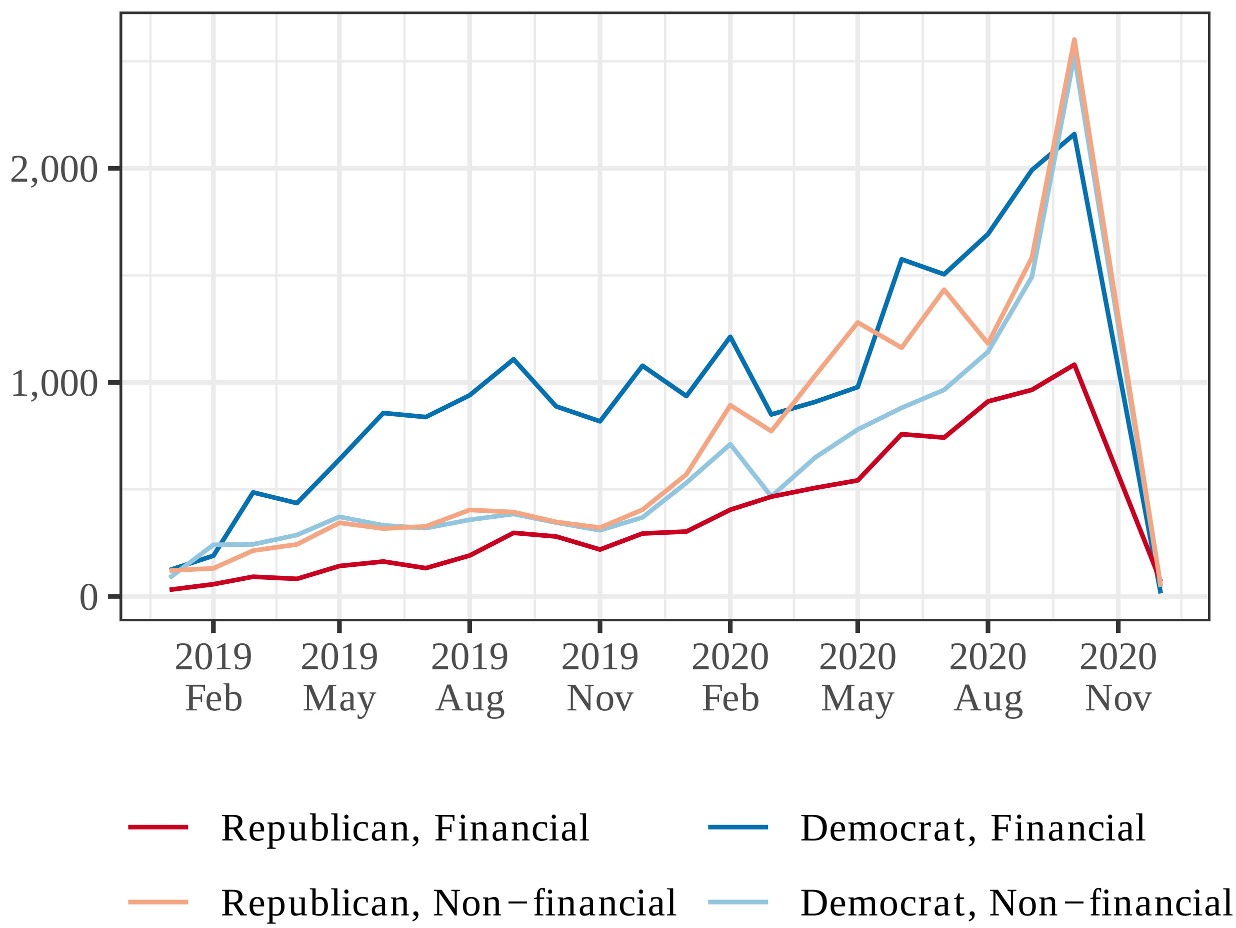

Figure 4: Number of Unique Ads Over Time by Type of Facebook Ads

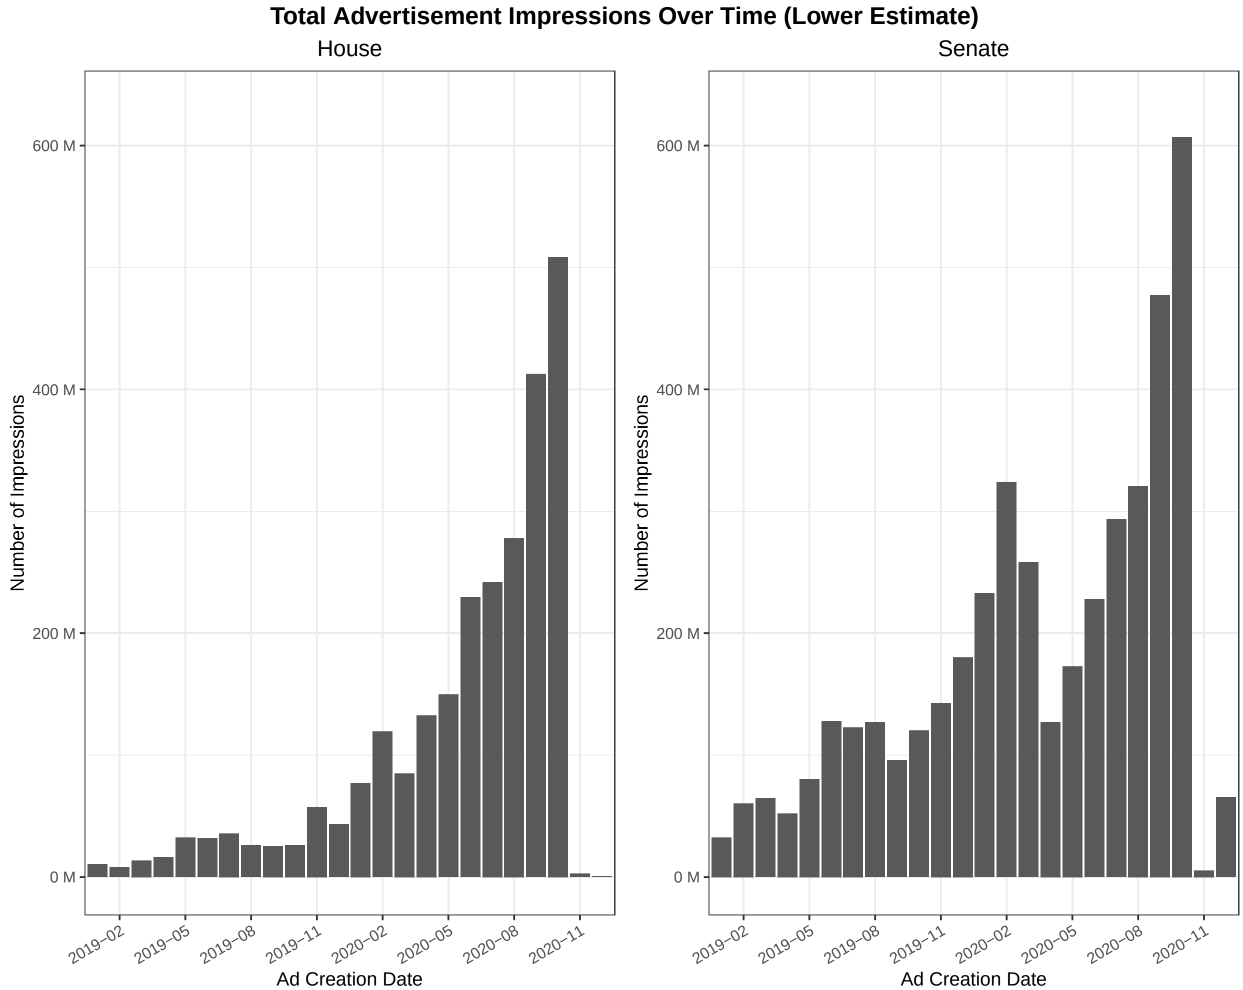

Figure 5: Total Impressions by Ad Creation Date

Appendix E Toxicity of Trump-Mentioning Advertisements

Figure 9 presents the toxicity of advertisements that do and do not mention Donald Trump by party and chamber. In the aggregate, for both parties in both chambers, advertisements that mention Donald Trump are more likely to be perceived as toxic than advertisements that do not mention him. To test whether these aggregate tendencies are statistically significant, we employed a series of Welch two-sample t-tests. While t-tests are only viable for normally distributed data in small samples, the law of large numbers and the central limit theorem make them a viable difference-in-means test for sufficiently large datasets, a condition which our hundreds of thousands of advertisements satisfy. The key results from these t-tests are as follows:

- House Republicans, Trump-mentioning versus non-Trump-mentioning:
  - T-statistic: 11.21
  - p-value: $p < 2.2 * 10^{-16}$
  - Mean toxicity of Trump-mentioning ads: 0.149

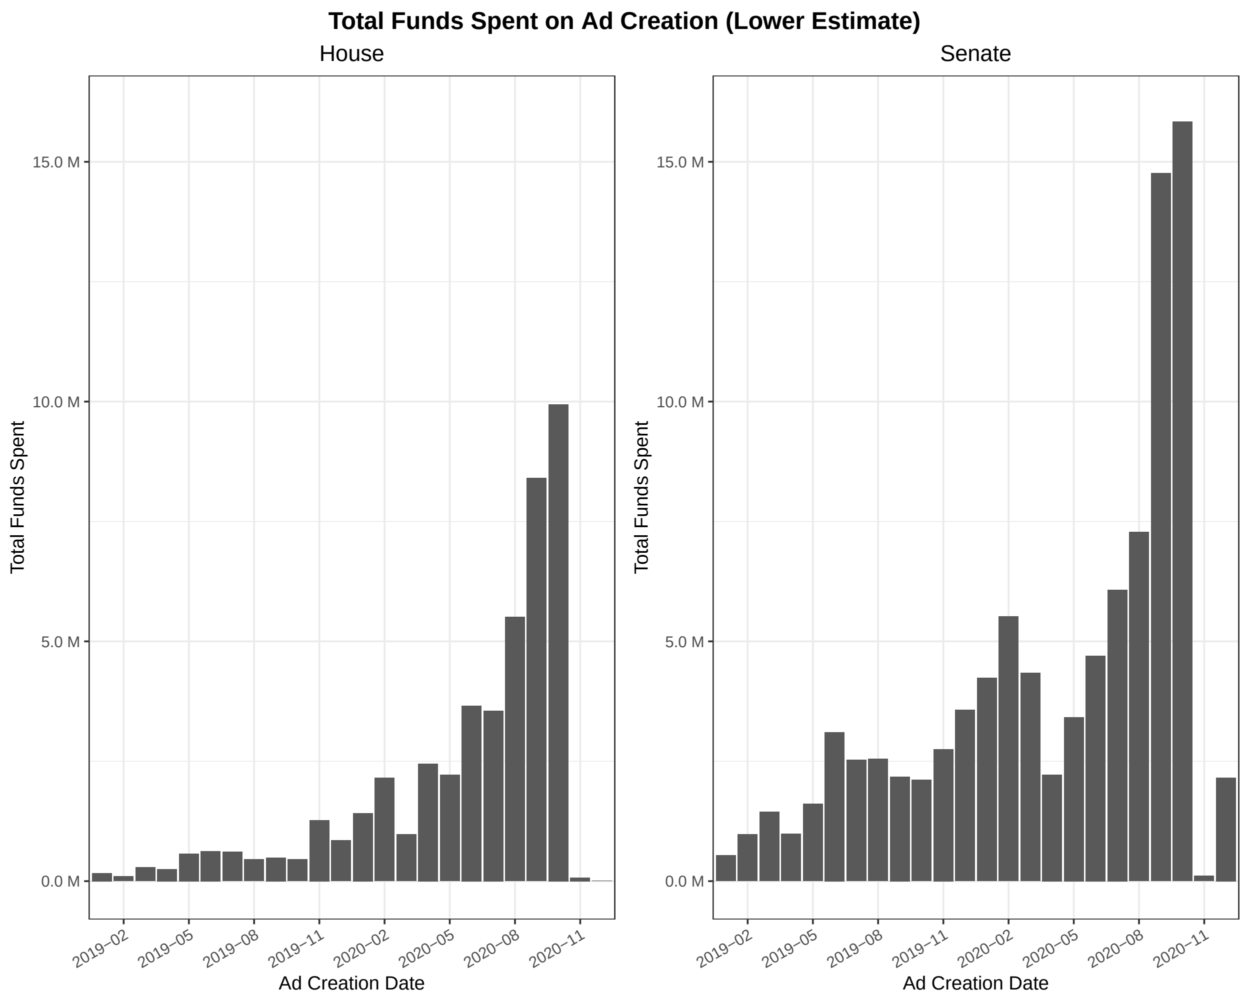

Figure 6: Total Funds Spent on Ad Creation Over Time

  - Mean toxicity of non-Trump-mentioning ads: 0.126

- House Democrats, Trump-mentioning versus non-Trump-mentioning:
  - T-statistic: 30.963
  - p-value: $p < 2.2 * 10^{-16}$
  - Mean toxicity of Trump-mentioning ads: 0.158
  - Mean toxicity of non-Trump-mentioning ads: 0.102
- Senate Republicans, Trump-mentioning versus non-Trump-mentioning:
  - T-statistic: 5.3841
  - p-value: $p = 7.98 * 10^{-8}$
  - Mean toxicity of Trump-mentioning ads: 0.143
  - Mean toxicity of non-Trump-mentioning ads: 0.128

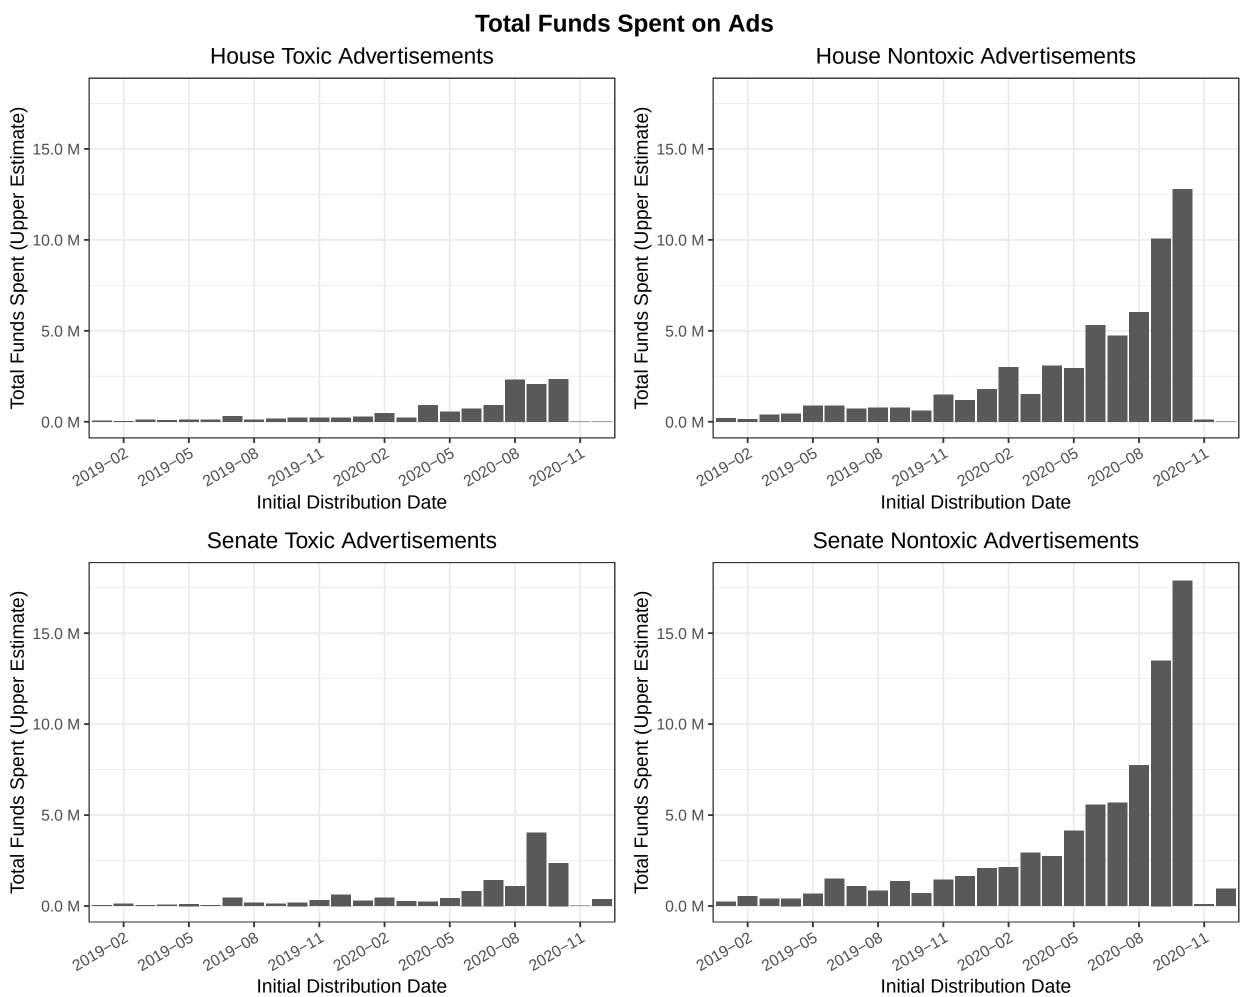

Figure 7: Total Funds Spent on Toxic and Non-Toxic Ads by Distribution Date

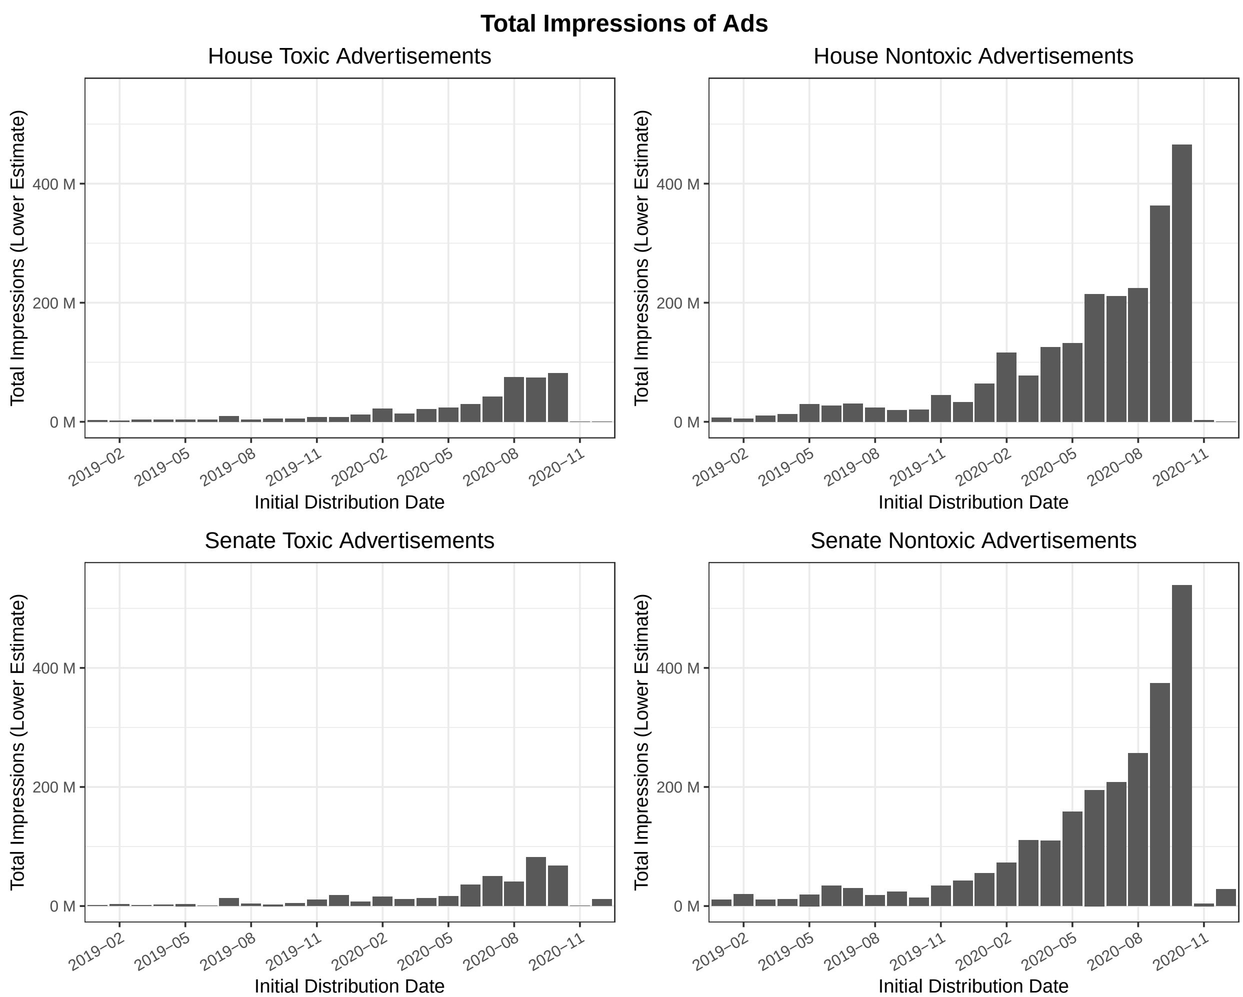

Figure 8: Total Impressions of Toxic and Non-Toxic Ads by Distribution Date

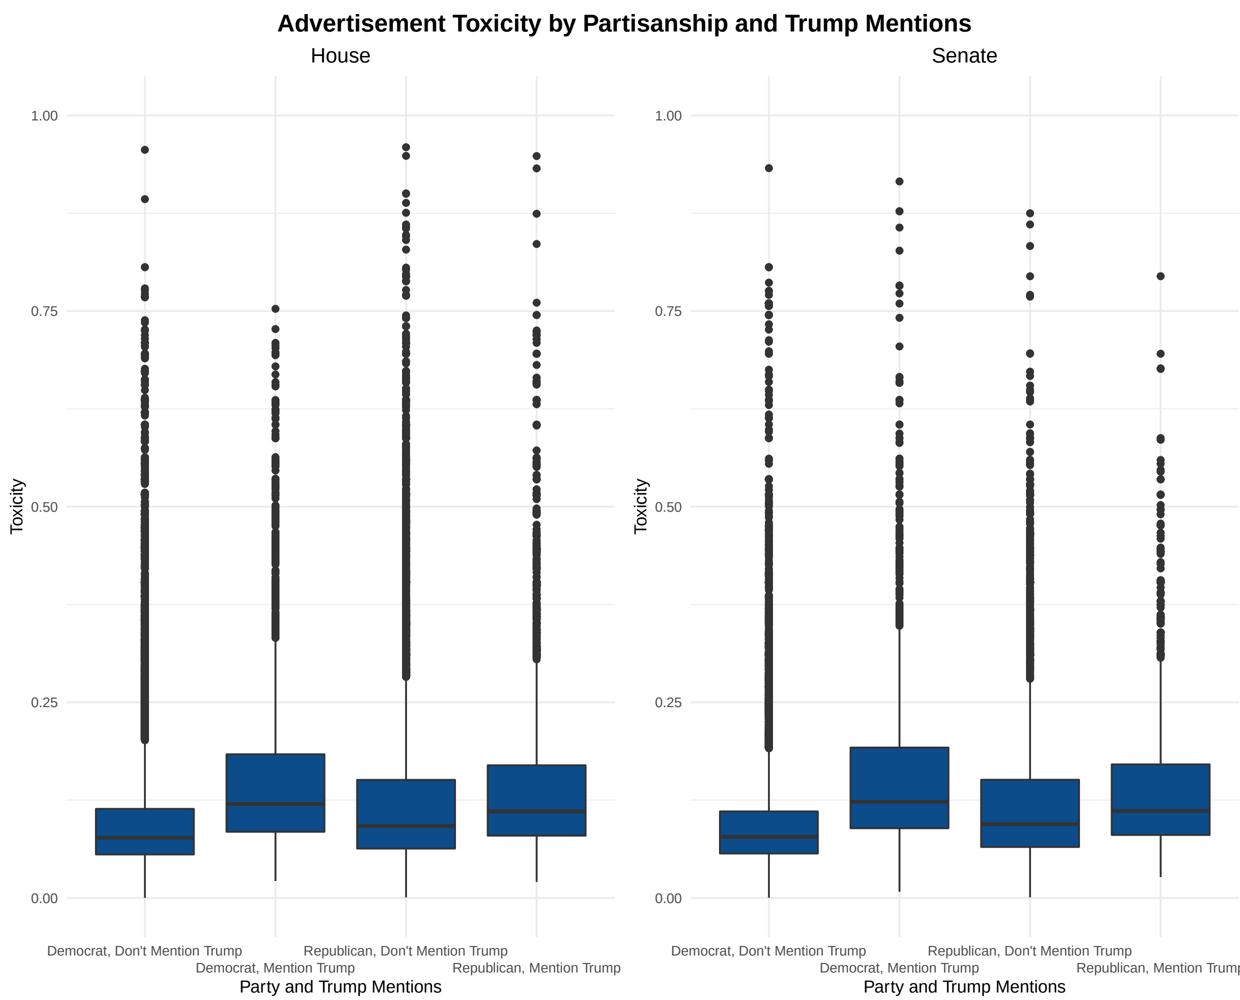

Figure 9: Toxicity of Trump-Mentioning Ads Across Chambers

- Senate Democrats, Trump-mentioning versus non-Trump-mentioning:
  - T-statistic: 27.87
  - p-value: $p < 2.2 * 10^{-16}$
  - Mean toxicity of Trump-mentioning ads: 0.165
  - Mean toxicity of non-Trump-mentioning ads: 0.097

The results of these t-tests reinforce the conclusion that advertisements that mentioned Donald Trump were wont to be more likely to be perceived as toxic than advertisements that did not mention him.

References

Blevins, Cameron and Lincoln Mullen. 2015. "Jane, John... Leslie? A Historical Method for Algorithmic Gender Prediction." *DHQ: Digital Humanities Quarterly* 9(3).

Eshima, Shusei, Kosuke Imai and Tomoya Sasaki. 2021. "Keyword Assisted Topic Models.".

Muddiman, Ashley, Shannon C. McGregor and Natalie Jomini Stroud. 2019. "(Re) Claiming Our Expertise: Parsing Large Text Corpora with Manually Validated and Organic Dictionaries." *Political Communication* 36(2):214–226.

Wilkerson, John and Andreu Casas. 2017. "Large-Scale Computerized Text Analysis in Political Science: Opportunities and Challenges." *Annual Review of Political Science* 20(1):529–544. tex.ids= wilkerson_large-scale_2017-1.
URL: *https://www.annualreviews.org/doi/10.1146/annurev-polisci-052615-025542*