Measuring Partisanship and Representation in Online Congressional Communications

Michael R. Kistner
University of Houston

Michael Heseltine
University of Amsterdam

Robert Alvarez
University of Houston

Maya Fitch
University of Houston

Lucas Lothamer
University of Houston

Elizabeth N. Simas
Texas A&M University

June 2025

Abstract: Social media and the internet have created new ways for representatives to communicate. How have members of Congress responded to these opportunities? We introduce a multi-platform dataset of congressional communications extending back to the onset of the social media era. Using computational language processing, we classify approximately 5 million tweets, 2.5 million Facebook posts, and 184,000 email newsletters authored by members of Congress between 2009-2022 based on intended purpose, and scale the partisanship of each message along a continuous left-right dimension. After validation, we demonstrate how our data can be used to study partisanship and representation in the contemporary Congress. Importantly, our data show congressional rhetoric has become more partisan and negative as social media usage has increased. We identify one potential mechanism contributing to this trend: partisanship and negativity receive inflated levels of positive engagement on social media, relative both to other message types and reception by offline audiences.

Keywords: Representation, polarization, social media, political communication, Congress

Word Count: 10,927

Center for Effective Lawmaking Working Paper 2025-02

Introduction

Communication is a critical component of representation (Burke 1774; Cohen 1989; Habermas 1968; Mansbridge 2003; Mill 1861; Pitkin 1967). Effective representation requires a dynamic, deliberative sharing of information between the elected official and those they represent in office. Although communication has always played a central role in representative government, developments in both society and technology shape this role in important ways.

In this paper, we introduce new data and measures that shed light on how rhetoric by elected representatives has evolved in the era of social media, email, and other online communication. Our focus is Congress, perhaps the most studied representative assembly. In the 1970s, scholars such as Mayhew (1974) and Fenno (1978) put a spotlight on congressional communications, arguing that legislators strategically promote themselves to various constituencies for electoral and other reasons. A key insight was that many of the specific messages representatives send can be classified as accomplishing one of a few basic tasks, such as promoting the legislator's personal brand, claiming credit for legislative accomplishments, or taking a stance on issues voters cared about.

Fifty years later, there has been both continuity and change in congressional communication. Members today continue to advertise, claim credit, and take positions. But the way in which they do so differs considerably from earlier decades. Unlike when Mayhew and Fenno were writing, the advent of online communication and social media mean that legislators can reach a vast audience both inside and outside of their district boundaries nearly instantaneously (Gainous and Wagner 2013; Russell 2021*b*). Traditional gatekeepers such as parties and legacy media organizations no longer maintain an oligopoly on communication, enabling politicians to connect with the public directly using chosen messages (Jungherr and Schroeder 2022; Schroeder 2018). Changes in the technological environment have been accompanied by changes in the political environment. Partisanship and polarization both inside and outside of Congress have grown considerably (McCarty 2019), possibly leading to more negativity and incivility in congressional rhetoric (Ballard et al. 2022, 2023; Costa 2021; Yu, Wojcieszak, and Casas 2024).

These developments have motivated a new wave of research investigating how legislators use social media and other online forms of communication. Much of this research either focuses on how legislators use new technology to accomplish the same core representational purposes as before (e.g., Hunt and Miler 2025; Russell 2021*b*) or evaluates the type of language legislators use on these platforms, considering the extent to which political elites deploy polarizing or partisan rhetoric on social media (e.g., Ballard et al. 2022, 2023; Yu, Wojcieszak, and Casas 2024).

This paper advances the study of representation and communication in three ways. First, we introduce new data on congressional communications spanning multiple platforms over a 14 year period (2009 to 2022), covering almost the entire time period that social media has been widely used.[1] The dataset includes over 5 million tweets, 2.5 million Facebook posts, and 184,000 email newsletters. Unlike data used in most existing research, which typically cover one to two years and focus on only a single mode of communication, the scope and span of our new dataset makes possible more rigorous research designs not previously feasible, possibilities we highlight in several applications.

Second, we use computational language processing tools to classify both the representational purpose and partisan language within each message. A set of transformer-based classification models specifically pre-trained on social media are used to classify each tweet into six different categories based on the intended purpose of each message. Though not exhaustive, 74% of all tweets and Facebook posts and 77% of newsletter sentences are classified into at least one of these categories. Similarly, a textual scaling model is used to place each message on a partisan spectrum from far left to moderate to far right based on the language used. These measures together cover some of the most foundational concepts to the study of Congress.

[1]Both Facebook and Twitter/X first became available for public usage in 2006. For brevity, throughout this paper we refer to the latter platform as Twitter, as the name change to X occurred after our dataset's timespan.

Third, we demonstrate the possibilities our data unlock via several applications. In these applications, we chart the development of congressional rhetoric across the social media period. We find that rhetoric has become more partisan, particularly among congressional Democrats, as well as more negative. These trends are found across platforms.

Further analysis highlights one potential explanation for these trends: negativity, position-taking, and partisanship receive considerably more positive engagement on social media than other posts (in some cases, over doubling the expected number of retweets or shares). In the new public arena where attention rather than space is the primary constraint (Jungherr 2014), this incentivizes politicians to embrace partisanship and negativity, despite evidence that constituents dislike such rhetoric (Costa 2021). The mismatch between the rhetoric that is rewarded on social media versus the rhetoric rewarded at the voting booth, we term the *social media feedback mirage*. Our analysis of the social media feedback, though initial, reveals that the underlying incentive structure of social media can have important consequences for discourse by political elites, meriting further study.

As a final application, we take advantage of the dense and precisely-timed nature of our data to draw precise causal inferences about congressional behavior. Specifically, we consider whether primary elections lead to short-term changes in the partisanship of member rhetoric using a series of difference-in-differences analyses that compare messaging by House members before and after their primary election. The analyses reveal small but sharply-estimated effect sizes. Most notably, Democratic members use more partisan rhetoric in the run-up to their primary election.

The paper proceeds as follows. First, we detail our data collection process and the computational methods used to measure partisanship and classify representational content in congressional communications. After validating these measures, we conduct the applications described above, examining temporal trends in rhetoric, exploring differences in response to communication on- versus offline, and evaluate how primary elections affect messaging strategies. Throughout, we demonstrate how our new dataset enables more comprehensive investigations of elite political communication than previously possible.

Congressional Rhetoric and Online Communication

The literature on congressional representation has long recognized that members play an active and strategic role in communicating information for various purposes, particularly electoral success (Fenno 1978; Mayhew 1974; Yiannakis 1982). For instance, Fenno (1978) emphasizes how political representatives shape their public image to align with crucial district elements, crafting a distinct "homestyle" to cultivate the trust of their constituents. Mayhew (1974) highlights three activities in particular that members do to enhance their re-election chances: position-taking (establishing stances on political issues); credit-claiming (taking responsibility for work done to pass legislation and secure resources for their districts); and advertising (bolstering name recognition, highlighting appearances at events and mentions in the media). More recent work confirms that members of Congress continue to use similar messaging strategies in modern times (Grimmer 2013; Grimmer, Westwood, and Messing 2015; Russell 2021*b*).

But while the underlying purposes of communication may seem similar, the means of communication have changed dramatically. In the contemporary era, the rise of the internet and social media has transformed communication, allowing representatives to reach a wider audience than ever before. Members of Congress are no longer dependent on a "franking" privilege to directly reach those they represent, and even lowly rank-and-file legislators now have the ability to broadcast views far beyond the boundaries of their districts. Moreover, messages no longer have to pass through newspaper reporters or television anchors before reaching the public (Gainous and Wagner 2013; Russell 2021*b*). The internet and social media allow representatives to reach constituents in a direct and unmediated way. At the same time, the proliferation of information requires information to break through the noise; with the demise of traditional news sources such as local papers (Hayes and Lawless 2015, 2017; Peterson 2021), there is no single front page or editorial section equivalent that guarantees messages will be seen by a large, captive audience (Jungherr and Schroeder 2022; Schroeder 2018). Social media has also enabled two-way communication between the mass public and political elites, allowing each group to influence the discourse of the other (Barberá et al. 2019; Warner 2023).

Researchers have begun studying how political elites use these new forms of media, although the field is still nascent. As recently as 2017, scholars were referring to online tools of campaign communication – "smartphones, Facebook, blogs, and the like" as "niche communication(s)" (Frankel and Hillygus 2017). Still, considerable progress has been made in understanding how these tools are used. We view these works as creating at least two major strands of research.

One of these major strands focuses on explaining what factors shape differential usage of social media and online communication, both in terms of the volume of social media usage as well as which types of messages (e.g., position-taking versus credit claiming) members choose to prioritize (Albert 2020; Cormack 2016*a*,*b*; Evans and Clark 2016; Evans, Cordova, and Sipole 2014; Hemphill, Russell, and Schöpke-Gonzalez 2021; Hunt and Miler 2025; Russell 2018*a*,*b*, 2021*a*,*b*; Scherpereel, Wohlgemuth, and Lievens 2018; Smith and Russell 2022; Straus et al. 2016; Tillery 2021). This work typically considers the political, institutional, demographic, sociological, and other variables that lead politicians to communicate in different ways.

A second major strand of research focuses on how rhetoric has evolved in response not just to the rise of social media, but also to the growing political polarization in American society. While many have studied the polarizing effects of social media usage on the mass public (for a review, see Tucker et al. 2018), others have focused more specifically on how political elites use social media in polarizing ways, via the language they use and how they discuss political issues online. Research on congressional rhetoric has considered the extent to which lawmakers deploy polarizing or extreme rhetoric (Ballard et al. 2023; Cowburn and Sältzer 2024; Heseltine 2023, 2024; Kaslovsky and Kistner 2025), negativity (Macdonald, Russell, and Hua 2023; Macdonald, Hua, and Russell 2024; Russell 2018*a*; Yu, Wojcieszak, and Casas 2024), and uncivil language (Ballard et al.

2022). Once again, a common theme in this work is the importance of electoral incentives in shaping member behaviors. Both social media users and donors (categories with some overlap) have been found to reward this type of language with engagement (Ballard et al. 2023; Macdonald, Russell, and Hua 2023) and dollars (Fu and Howell 2020; Yu, Wojcieszak, and Casas 2024).

While this research has improved our understanding of how representatives communicate, most of this work has been hamstrung by four key limitations. The first limitation, common to almost all of the above-cited research, is a focus on short time periods. Due in part to the difficulties in collecting and cleaning communication data, most research uses at most a few years of data, which can lead to inferential issues.[2] For instance, studies using data from a single session (Hemphill, Russell, and Schöpke-Gonzalez 2021; Yu, Wojcieszak, and Casas 2024) have found communication differences between Democrats and Republicans, which are attributed to one party being in the majority and the other in the minority. But Democrats and Republicans differ from each other in many ways, making it impossible to separate partisan differences from majoritarian differences when analyzing just a single session's worth of data.

Besides avoiding problems such as these, longer time spans are desirable for another reason. It's unclear whether congressional communication has stayed largely constant or evolved as social media usage and technology have changed. Particularly given the speed of developments in online communication, a major concern is the temporal validity of findings in this area (Munger 2023). Assessing how stable conclusions are over longer periods of time requires data covering longer periods of time.

A second issue with most existing research is examining communication on a single platform.[3] The audiences members speak to when posting on a social media platform like Twitter/X

[2]Section A of the Supplemental Materials displays a table of published political science articles over the past twelve years that analyze text from online Congressional communication data (Twitter/X, Facebook, or e-newsletters). Out of these 30 articles, 9 used one year of data, 8 used two years, 4 used three years, 3 used four to six years, and 6 used more than six years of data.

[3]Again referencing Section A of the Supplemental Materials, only 4 of 30 published articles studying online congressional communication examined more than one platform.

– where messages are seen by a heterogeneous mix of political enthusiasts, journalists, interest group members, fellow politicians, and more – look very different from the recipients of email newsletters, which are targeted more directly towards constituents.[4] Recently published research demonstrates that these different audiences matter. Members vary in terms of how much they post on Facebook versus Twitter/X (Blum, Cormack, and Shoub 2023). Furthermore, the partisanship of member speech varies by venue, as some members appear more partisan when measured using communication in one form versus another (Green et al. 2024). Other hypotheses researchers are interested in testing may be platform-specific, and demonstrating similarities or differences across platforms can provide deeper insight.

A third issue is that research has largely studied *either* the representational content *or* the partisan tone of congressional speech. While the concepts are distinct, they are not mutually exclusive. Position-taking is often inherently partisan, and elites frequently use partisan rhetoric and one or more representational strategies in the same message. Being able to analyze both content and tone simultaneously allows researchers to more precisely isolate what component of messaging is having the effects they find, without worrying about confounding the impact of one dimension for the other. For these reasons, having readily accessible and easily comparable measures of both concepts (representational purpose and partisanship) enables more robust scholarship than possible when studying either in isolation.

The fourth issue is that the current system, where research teams individually download, clean, and prepare different versions of similar datasets, is inherently wasteful and slows the pace of scientific progress. Having a central repository of easily accessible, ready-to-use data allows researchers to spend their time developing and testing theories of political communication, not repeating time-consuming data work that has been done many times over. To build on this point, having a single commonly-used dataset ensures that similar cleaning and sample inclusion

[4]On that latter point, offices sometimes require individuals sign up for e-newsletters using a zip code, to confirm constituency residency.

decisions have been made. Idiosyncratic data processing decisions – that may not be immediately obvious to readers – can be eliminated as possible explanations when differing results emerge, making comparison of results more transparent.

For these reasons, the study of communication and representation stands to benefit enormously from a single publicly available dataset with multimodal and longitudinal data, classified by representational purpose and scaled according to the partisan positioning expressed via the message.

Measuring Partisanship and Representation

To address these limitations and advance future research, we create the *Scaled and Classified Congressional Communication (SCCC) dataset*, a new multimodal dataset spanning the years 2009 to 2022. The dataset includes posts on the two most-used social media platforms by members of Congress, Twitter/X and Facebook, as well as email newsletters, a common form of communication members use to address (primarily) constituents. In addition to the texts of these communications, we possess auxiliary variables such as social media engagement metrics (likes, retweets, shares, etc.), as well as variables measuring partisanship and representative purpose.

Our measurement schemata are shown below in Figure 1, along with example messages that correspond to each category or scale position. The representational categories come directly from Mayhew (1974), although these or similar categories have been studied frequently by others (e.g., Grimmer 2013; Russell 2021*b*; Yiannakis 1982). These categories are *Advertising* (“any effort to disseminate one’s name among constituents in such a fashion as to create a favorable image"), *Credit Claiming* (“generat[ing] a belief...that one is personally responsible for causing the government, or some unit thereof, to do something that the actor...considers desirable"), and *Position Taking* (“a judgmental statement on anything likely to be of interest to political actors"). We further subdivide the credit claiming category to encompass two different forms of credit claiming, *Credit Claiming for Constituency Work* (a message taking responsibility for particularized benefits provided to constituents, the district, or the state) and *Credit Claiming for Policy Work* (a message taking responsibility for non-constituency-specific policy accomplishments).

The partisanship categories mirror those studied by Russell (2018*b*); specifically, we identify both *Negative Partisanship* (an attack on the policies and politicians of the opposing party) and *Bipartisanship* (advocating the value of bipartisan collaboration) in messages.[5] In addition, we scale the *Partisan Orientation* of each message on a scale that ranges from -1 (most Democrat-leaning) to 0 (nonpartisan) to 1 (most Republican-leaning).

The categories are neither mutually exclusive nor collectively exhaustive. A single tweet could, for instance, advertise ("As the representative for TX-23, "...), position-take ("securing the border is one of my top priorities."), credit claim ("This is why I'm sponsoring legislation..."), and make a negative partisan attack ("The Biden border crisis must be stopped!"). It can also be none of the above, wishing (for example) followers a Merry Christmas or Happy Holidays.

While we choose these categories in part based on congruence with existing research, these categories also each represent critical components of the information voters require in order to hold legislators accountable. A useful theoretical framework to think about representation is the *principal-agent model of political accountability* (Ashworth 2012; Besley 2006), in which voters select politicians with aligned preferences whom they believe will do the best jobs of advancing their interests. As a consequence, politicians motivated by re-election take actions that correspond to voter preferences and benefit constituents, in order to maximize their chances of remaining in office. This relationship requires that voters know who their representatives are (which *advertising* facilitates), how they vote for or against specific policies (*position-taking*), and

[5]Russell (2018*b*) also studies a third partisanship category, *Positive Partisanship*, messages that "signal favoritism or support for one's own party [or one's] party's candidates" (p. 703). We omit this category in our measurement for two reasons. First, the percentage of messages that were positive partisan (as classified by our research assistants in the manually classified sample) was quite low, lower than other categories. Second, the accuracy of our supervised machine learning classifications was considerably lower for this category than all others, likely due to the small sample size.

FIGURE 1: Classification and Scaling Schemata With Example Text

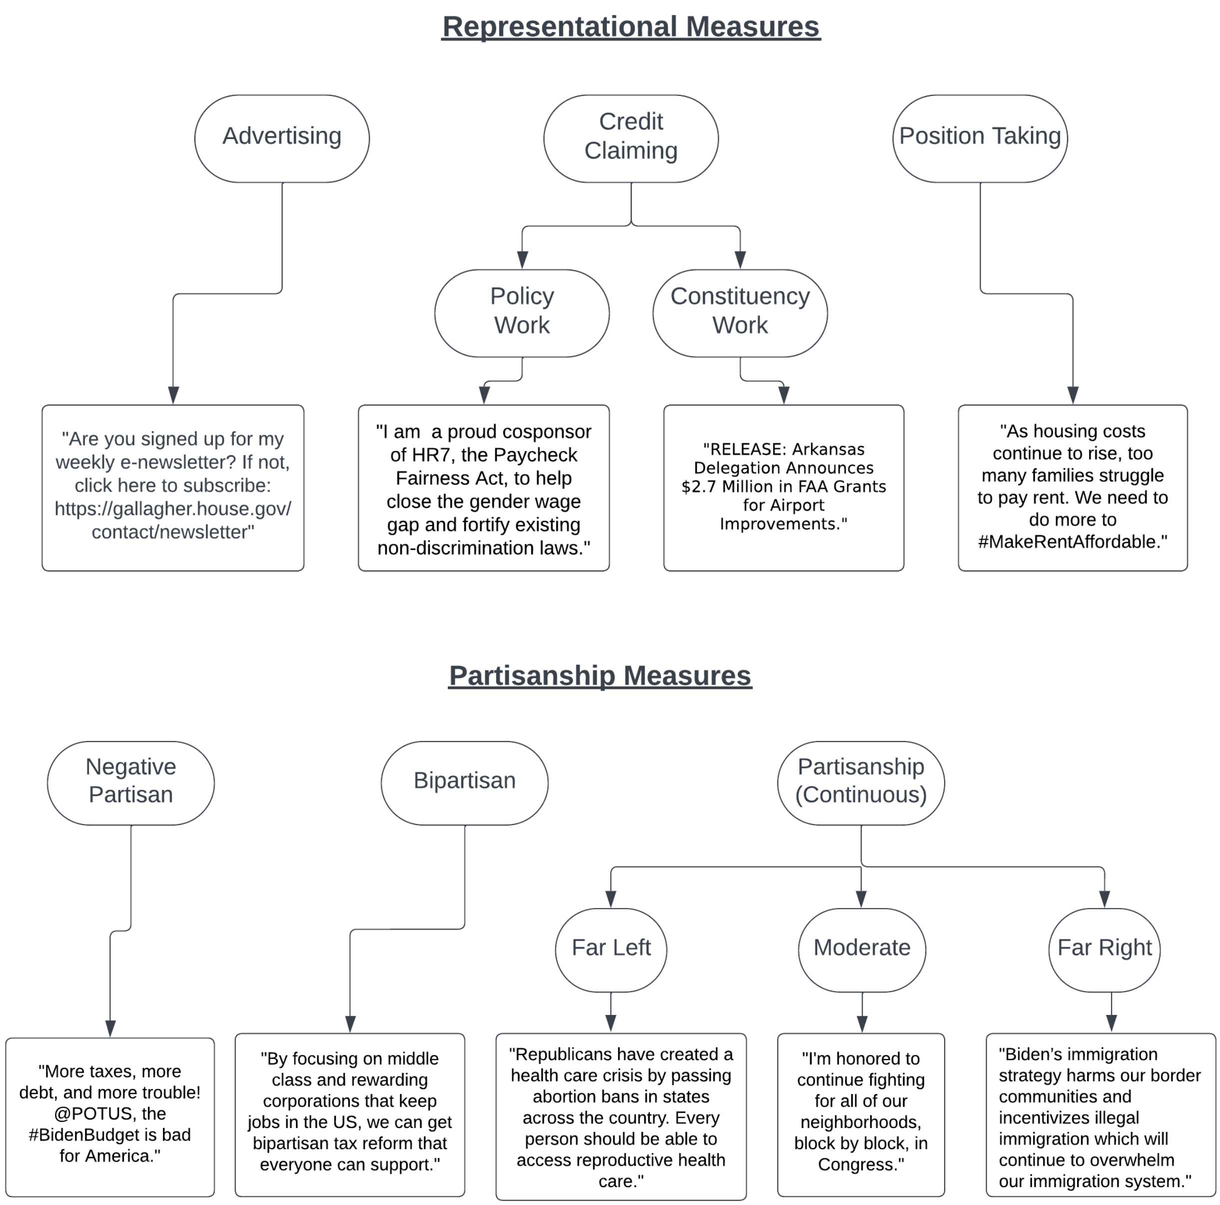

Note: The figure displays the six separate classification categories and the continuous partisanship dimension our data contains. The top half of the figure displays the representation categories (Advertising, Credit Claiming for Policy Work, Credit Claiming for Constituency Work, and Position Taking), while the bottom half shows the partisanship categories (Negative Partisanship, Bipartisanship) and the continuous partisan dimension (Partisan Score). Example tweets classified into each category, as well as three tweets located at different points of the Partisan Score dimension, are shown to illustrate typical messages.

the actions they take to aid constituents (*credit-claiming*). *Bipartisan* and *negative partisan* messages send signals to voters about whether representatives will work with or oppose members of the other party, also critical pieces of information in evaluating how legislators do their job in a polarized age.

Data Collection

The communications data come from multiple sources. For tweets and Facebook posts, data were downloaded directly on a rolling basis beginning in 2018.[6] Tweets were first downloaded using the V2 API endpoint and then later using the Academic Twitter API. For Facebook, posts were downloaded using the CrowdTangle Platform. For newsletters, data come from the publicly available www.DCInbox.com repository of email newsletters collected and cleaned by Cormack (2017). For newsletters, the data are available dating back to 2010; in the case of Twitter/X and Facebook, the data are available dating back to 2009. The dataset currently spans through 2022, although we aim to make periodic future updates of the dataset to broaden the timespan and enable study of contemporary congressional communication.

Table 1 displays the total number of tweets, Facebook posts, and email newsletters included in the dataset, listed by biennial legislative session. In total, the data consist of 7,827,972 unique communications from 1,025 US senators and representatives. As the table shows, online communication appears to have grown considerably, although some of the social media differences in the pre-2018 years may be due to the deletion of accounts. For Twitter and Facebook, the table thus displays lower bounds on the total number of users as well as tweets and posts from this time period. On the other hand, the number of tweets and posts from the median member has gone

[6] Our Twitter/X and Facebook data includes official, campaign, and personal accounts publicly associated with members of Congress. All account names/handles were hand-collected by one of the authors, with periodic updates to include new Members and check for newly created accounts. Content from posts and accounts deleted before the date of collection were, by definition, not retrievable. In the case of Facebook accounts, not all accounts were created as official accounts. As Crowdtangle only facilitated the collection of data from specifically-created flagged account types and not personal pages, retrieval of data from some accounts was not possible, an issue which mostly impacted accounts from earlier time periods.

up considerably over this time period (over a fourfold increase for each social media platform), suggesting that the considerable growth in social media usage over this time period is not merely an artifact of missing accounts.

Classification Procedure

To classify messages into the six binary categories displayed in Figure 1, we begin with approximately 22,500 messages manually classified by political science PhD-level researchers. These messages are then used to train a transformer-based classification model pre-trained on social media data. Below, we describe the classification process in more detail.

First, members of the research team classified a stratified random sample of tweets posted by members of Congress across the entire time span (2009 to 2022), sampled to ensure an approximately equal number of tweets for each chamber-year combination. 3,500 of these tweets were read separately by three members of the research team, who then made a binary decision for each category. In all categories, the Krippendorf's alpha was 0.90 or above, indicating high levels of intercoder reliability.[7] An additional 14,000 tweets were read and classified by a single member of the research team. Doing so resulted in a dataset of 17,500 manually classified tweets. After withholding 1,375 tweets for out-of-sample validation data, the remaining tweets were used as training data.[8] A variety of classification algorithms were tested, with out-of-sample balanced accuracy and F1 score used to choose among alternatives. Ultimately, the BERTweet language model (Nguyen, Vu, and Tuan Nguyen 2020), a RoBERTa variant pre-trained specifically on English-language tweets, proved to have the best performance.[9] Besides being pre-trained

[7] This level of intercoder reliability is above the norm for most published political science research. We attribute the high levels of agreement to both the quality of the coders as well as the rigorous preparation process, which involved a week of training, practice, and group discussion on decision-making and protocol before final classifications began. In section B of the Supplemental Materials, we include a copy of training materials provided to coders during these training sessions.

[8] The training data were split into 14,125 tweets used for training, and 2,000 tweets used for in-sample testing.

[9] An alternative to the supervised learning procedure described here would be to use a large language model to perform classifications, either using one-shot learning or fine-tuning a model. Such an approach is less than optimal

TABLE 1: Online Communication by Members of Congress (2009 - 2022)

Medium	Session	Years	Percent Using	# (Median Member)	# (Total)
Tweets	111	2009-2010	–	–	77,798
	112	2011-2012	82.4	473.0	365,725
	113	2013-2014	89.6	843.5	608,535
	114	2015-2016	91.4	968.0	712,953
	115	2017-2018	94.1	1234.0	907,365
	116	2019-2020	96.9	1636.0	1,149,654
	117	2021-2022	97.8	1723.0	1,157,227
Facebook Posts	111	2009-2010	–	–	51,023
	112	2011-2012	69.7	204.5	164,254
	113	2013-2014	81.0	276.0	204,511
	114	2015-2016	87.3	533.0	343,100
	115	2017-2018	91.0	747.0	474,591
	116	2019-2020	95.4	930.5	623,920
	117	2021-2022	97.1	1088.5	702,340
Newsletters	111	2009-2010	–	–	8,085
	112	2011-2012	89.0	27.0	21,711
	113	2013-2014	90.5	27.0	22,416
	114	2015-2016	92.6	30.0	23,962
	115	2017-2018	91.9	33.0	24,849
	116	2019-2020	92.4	43.0	30,558
	117	2021-2022	88.5	39.0	29,417

Note: The table displays the usage of tweets, Facebook posts, and email newsletters by members of Congress in our dataset. Percent Using and Number by Median Member are excluded for the 111th session, for which data spanning the full session do not exist.

on social media data, the BERTweet classifier also uses vector embeddings to read messages in their entirety with the original ordering preserved, as opposed to more simplistic bag-of-words approaches, making it unsurprising that the accuracy outperformed alternatives such as Naive Bayes and random forests models. The training data were used to tune the base BERTweet model for three epochs, with a learning rate of 2e-5.

Out-of-sample classification performance metrics (accuracy, balanced accuracy, and macro F1) for tweets, as well as out-of-sample hand-coded validation datasets for Facebook and email newsletters, are displayed in Table 2, while more detailed classification metrics are provided in section C of the Supplemental Materials.[10] The macro F1 for each category of tweets ranged from 0.83 to 0.92, values indicating strong performance. The BERTweets models were then applied to these unseen messages from other communication platforms to compare results. Classifier performance on these other communication platforms was lower than it was for the tweet data, but still within acceptable ranges, demonstrating the possibility of cross-platform application of pre-existing classification models. For these non-Twitter messages, macro F1 ranged from 0.62 to 0.92 depending on the category. For comparison's sake, the bottom of Table 2 displays the accuracy, balanced accuracy, and F1 score (as reported) for six recently published research articles classifying social media messages by members of Congress (Ballard et al. 2022, 2023; Yu, Wojcieszak, and Casas 2024), U.S. state legislators (Butler, Kousser, and Oklobdzija 2023; Payson

here, for multiple reasons. First, using a proprietary LLM for such a task would be prohibitively expensive given the approximately 10 million messages that would need to be classified. Second, the classifications themselves would be subject to change based on updates to the LLM, which occur regularly, producing instability and limiting replicability. Despite this, we assessed the relative classification success on a small subset (2,000 tweets) of our data using a fine-tuned GPT-4o mini model. Differences in classification were minimal. As also shown by Heseltine and von Hohenberg (2024), the difference in downstream results based on LLM or transformer-based classification is negligible, at a fraction of the expense or computational resource requirements.

[10] For Twitter, we use the 1,375 hold-out set. For Facebook, we use 2,462 coded sentence bigrams, and for the newsletters, we use 2,500 coded newsletter sentence bigrams. For evaluating out-of-sample performance, we use sentence bigrams because they are approximately the same length as tweets, on which the model was pre-trained and trained. Because approximately half (45.5%) of Facebook posts are two sentences or fewer and two-thirds (69%) are three sentences or fewer, when applying our models to make final classifications, the Facebook models are applied to entire posts. Because newsletters are considerably longer, final classifications are made on sentence bigrams. Unless otherwise noted, analyses for newsletters are conducted at the sentence bigram level.

et al. 2022), or both (Fowler et al. 2021).

TABLE 2: Classification Accuracy Metrics by Category and Mode, With Comparisons

Out-of-Sample Accuracy Metrics

Platform	**Category**	**Accuracy**	**Balanced Accuracy**	**F1**
Twitter	Advertising	0.89	0.89	0.83
	Credit Claiming (Constituency)	0.94	0.75	0.83
	Credit Claiming (Policy)	0.94	0.83	0.84
	Position Taking	0.87	0.82	0.87
	Bipartisanship	0.99	0.93	0.92
	Negative Partisanship	0.95	0.91	0.91
	Range:	0.87 - 0.99	0.82 - 0.93	0.83 - 0.92
Facebook	Advertising	0.87	0.76	0.77
	Credit Claiming (Constituency)	0.91	0.68	0.72
	Credit Claiming (Policy)	0.94	0.85	0.84
	Position Taking	0.86	0.86	0.86
	Bipartisanship	0.99	0.90	0.91
	Negative Partisanship	0.95	0.92	0.86
	Range:	0.86 - 0.99	0.68 - 0.92	0.72 - 0.91
Newsletters	Advertising	0.87	0.75	0.71
	Credit Claiming (Constituency)	0.80	0.62	0.64
	Credit Claiming (Policy)	0.90	0.81	0.84
	Position Taking	0.77	0.79	0.77
	Bipartisanship	0.99	0.88	0.91
	Negative Partisanship	0.96	0.92	0.82
	Range:	0.77 - 0.99	0.62 - 0.92	0.64 - 0.91

Comparisons From Other Published Work

Article	**Platform**	**Accuracy**	**Balanced Accuracy**	**F1**
Fowler et. al. (2021)	Facebook	0.80 - 0.99	0.50 - 0.95	–
Payson et. al. (2022)	Tweets	0.55 - 0.59	–	0.35 - 0.92
Ballard et. al. (2022)	Tweets	0.63 - 0.97	–	0.64 - 0.97
Ballard et. al. (2023)	Tweets	–	–	0.75 - 0.94
Butler et. al. (2023)	Tweets	0.20 - 0.99	–	–
Yu et. al. (2024)	Tweets	0.66 - 0.74	0.67 - 0.85	0.50 - 0.77
	All:	0.20 - 0.99	0.50 - 0.95	0.35 - 0.97

Note: The top portion of the table displays the accuracy, balanced accuracy, and F1 score (out-of-sample) for each of the six representational categories for tweet, Facebook post, and newsletter sentence data. The bottom portion of the table displays the equivalent metrics reported in recently published research articles using supervised classification techniques to classify social media messages by legislators, to give context for classifier performance.

We validate the resulting classifications in two ways. First, we establish the face validity of our classifications by creating keyword plots, showing which words are most strongly associated with classification into each of the six categories. Figure 2 shows the top 20 words for each category. The words in each category accord with expectations. For example, messages that credit claim for constituency work discuss words associated with visiting the district ("town", "hall", "office"), aiding constituents ("assistance", "help"), and securing money for the district ("funding", "grant"). Similarly, messages that attack the other party feature the names of top leaders of each party ("trump", "biden", "pelosi") or divisive issues ("obamacare", "border").

Second, we establish the convergent validity of our classifications by comparing the communication styles members use to members' legislative styles, as introduced by Bernhard, Sewell, and Sulkin (2017) and discussed further in Bernhard and Sulkin (2018). These authors use a cluster analysis approach applied to behavioral data of Congressional actions (bill introductions and cosponsorships, party-line voting frequency, quantity of district offices and staff, etc.) to categorize members of the U.S. House into five distinct groupings: District Advocates, Policy Specialists, Party Builders, Party Soldiers, and Ambitious Entrepreneurs. While their data only extends until 2008, 1,079 members in our data served at least one session in the Bernhard and Sulkin data.

In Section D of the Supplemental Materials, we show what percent of these member types' tweets, Facebook posts, and newsletter bigrams fall into the different categories. As can be seen in Figure D.1, there are clear differences in communication by members with different legislative styles. For example, District Advocates have high levels of credit claiming relative to other members, particularly claiming credit for constituency work. Similarly, Party Soldiers have low levels of bipartisanship in their communications.

FIGURE 2: Words Most Associated With Each Category

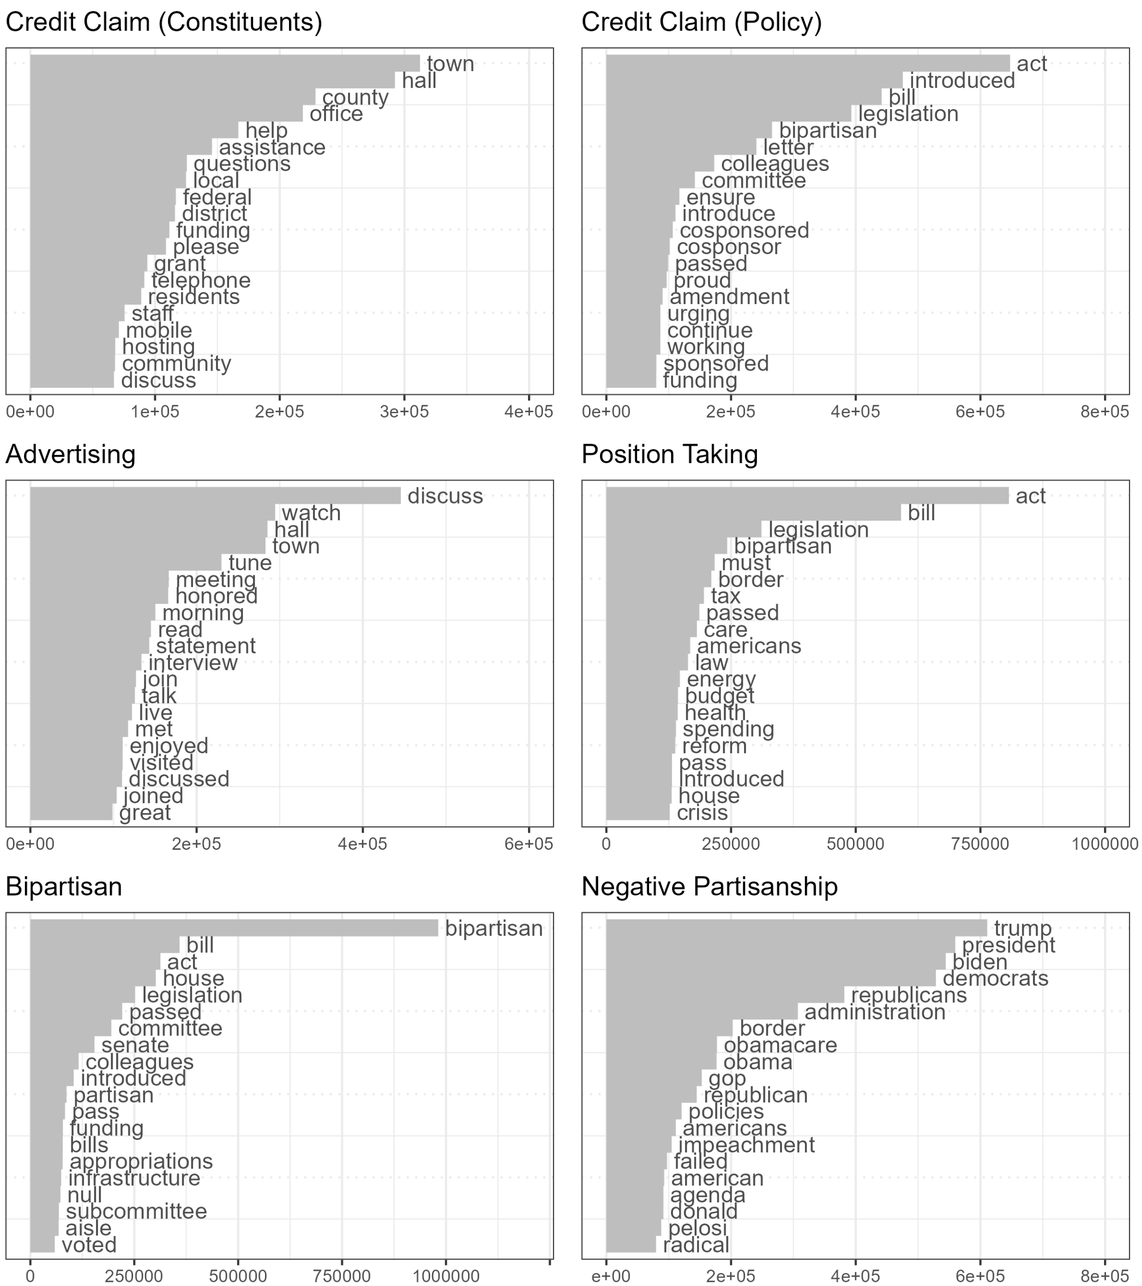

Note: The figure displays the top 20 words most associated with a message (tweets, Facebook posts, and newsletters combined) being classified into each of the six categories. The x-axis shows the likelihood ratio for each word.

Scaling Procedure

To scale speech as more or less partisan, we adopt the same general strategy as Gentzkow, Shapiro, and Taddy (2019), who measure partisanship in Congressional floor speeches. Following their lead, we define speech as partisan on the basis of how strongly it identifies the party of the speaker. Extreme partisan speech is used almost exclusively by members of one or the other party, while moderate speech is used by both. For instance, an individual who uses terms such as "border crisis" when discussing immigration is likely to be a Republican, while the utterance of "pathway to citizenship" means the speaker is likely to be a Democrat.[11] Rhetoric is partisan if a speaker's language is dominated by terms used primarily by one party or the other.

To capture this definition of partisanship, a class affinity scaling model (Perry and Benoit 2017) is fit using the words contained in the tweets, Facebook posts, and newsletters.[12] While similar in some respects to supervised text classification, this method is better described as a scaling procedure because it models speakers as having a continuous "affinity" towards classes (e.g., parties) rather than simply belonging to a binary class or not. In this model, affinity towards party (partisanship) is parameterized as $\pi_r \in [0, 1]$, the probability that for any given token of speech W_i the underlying orientation of the speaker is $U_i = r$, or Republican. The probability that the speaker's underlying orientation is Democratic ($U_i = d$) for a given token of speech is thus $1 - \pi_r$. The orientation of a speaker for each token $i = 1, \cdots, n$ determines the probability of a specific word w being used:

$$\Pr(W_i = w) = \pi_r \Pr(W_i = w | U_i = r) + (1 - \pi_r) \Pr(W_i = w | U_i = d).$$

[11] Our approach can more precisely be described as measuring the *partisanship* of speech as opposed to *ideology*, given that partisan speech may encompass not just the discussion of policy issues, as in the examples above, but also non-policy topics, such as individual politicians (e.g., "crooked Hillary"). In practice, there is considerable overlap between the ideology and partisanship of speech.

[12] Prior research, e.g., Kaslovsky and Kistner (2025), has used the class affinity scaling model to measure the partisanship of tweets, although not, to our knowledge, Facebook posts or newsletters.

The partisanship for any given tweet, Facebook post, or newsletter is the expected proportion of time the underlying orientation is r versus d, which can be calculated as $\pi_r = \mathbb{E}\left\{\frac{1}{n}\sum_{i=1}^{n} U_i = r\right\}$. For interpretability sake, we rescale the resulting probability so it ranges from -1 (most Democrat-leaning) to 0 (neutral partisanship) to 1 (most Republican-leaning), calculated as $2\pi_r - 1$. This rescaling we refer to as the *Text Partisanship Score*. This measure can be folded, i.e., $|2\pi_r - 1| \in [0, 1]$, so that higher values indicate more extreme partisan rhetoric regardless of the speaker's party. This measure we refer to as the *Text Partisan Extremity Score*.

In Supplemental Section E, we discuss the choice of a class affinity scaling algorithm versus alternatives such as Wordscores (Laver, Benoit, and Garry 2003) and Wordfish (Slapin and Proksch 2008) that are commonly used to scale political texts. While the performance, interpretability, and computational efficiency of the class affinity model make it our preferred choice for this task, reassuringly, as Figure E.1 shows, the text-level scalings of alternative algorithms are highly correlated (r = 0.74 to r = 0.99) with the measure produced by the class affinity model, providing confidence that conclusions should not not be highly sensitive to the choice of scaling algorithm. As before, we create keyword plots showing words most associated with each tercile of our scale (Democratic, bipartisan, Republican) that are displayed in Figure F.1 of Supplemental Section F.[13]

To validate this scaling, we aggregate to the member level by taking the average Text Partisanship Score across all texts (tweets, Facebook posts, and newsletters) in our dataset.[14] These aggregated member-level Text Partisanship Scores are then compared to two separate, commonly-used measures of Congressional ideology. Figure 3 compares a member's average Text Partisanship Score to the member's DW-NOMINATE score (Poole and Rosenthal 2000), which is estimated using roll call voting, and the member's campaign finance (CF) score (Bonica 2014), which is esti-

[13] Our measurement strategy differs from (but complements) approaches that scale politicians using features of social media besides the textual content, such as the network structure of followers (Barberá 2015) or the hyperlinks included in the message (Heseltine 2023). Unlike these approaches, our measurement strategy enables the scaling of *individual messages*, regardless of whether (for instance) a hyperlink is included or not.

[14] To establish face validity, we also construct a similar plot to Figure 2 for terciles (liberal, moderate, conservative) of our partisanship scale. This keyword plot is displayed in Figure F.1 of the Supplemental Materials.

mated using campaign contributions to the member.[15] While we view partisanship and ideology as distinct conceptually, empirically the two are closely related. More ideologically extreme conservatives and liberals typically speak in more partisan ways, and vice-versa. DW-NOMINATE scores and CFscores thus each provide useful benchmarks for our text partisanship scores, representing two important domains (roll call voting and campaign finance) of American politics.

FIGURE 3: Comparing Text Partisanship Scores to Ideology Measures

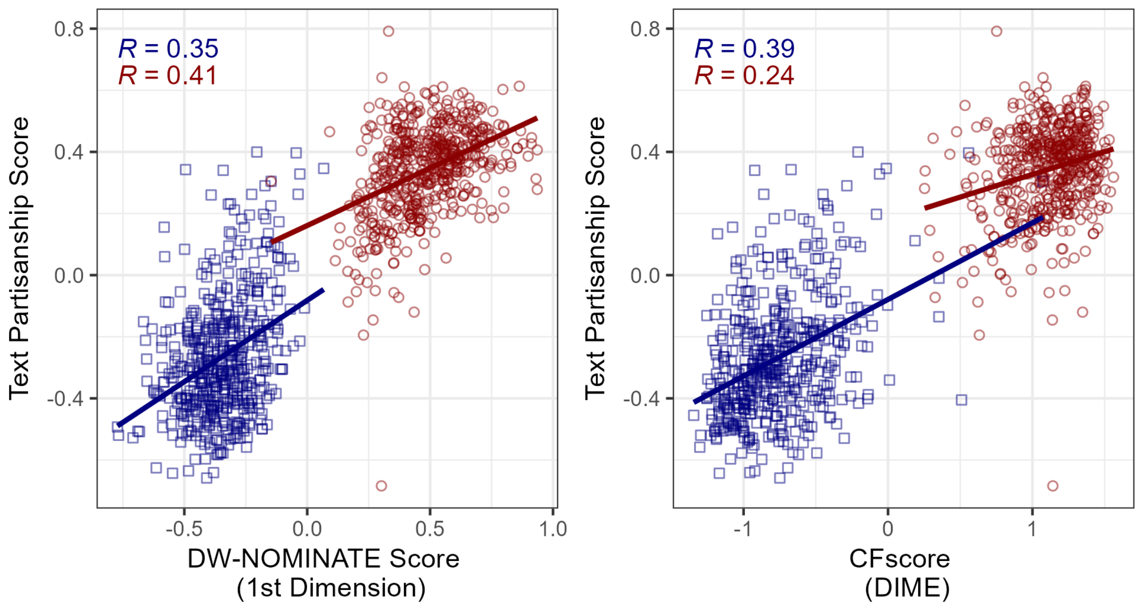

Note: The figure displays the relationship between a member's average text partisanship score to the member's first dimension DW-NOMINATE score (on the left) or CFscore (on the right). Democrats are denoted with blue squares, Republicans with red circles. Within-party regression lines and correlation coefficients (Pearson's r) are shown to assess relationship strength.

As Figure 3 shows, the average Text Partisanship Scores are predictive of members' estimated ideology (using either the roll call or campaign finance-based measure) even after accounting for party identification. The within-party correlation between text partisanship scores and DW-NOMINATE scores is 0.41 for Republicans and 0.35 for Democrats. The within-party correlation between text partisanship scores and CF scores is 0.24 for Republicans, and 0.39 for

[15]DW-NOMINATE score data is publicly available at www.voteview.com, while CFscores can be downloaded as www.data.stanford.edu/dime

Democrats.[16] These results indicate that the partisanship contained within the text of members' tweets, Facebook posts, and newsletters is clearly related to both the way a member votes and whom a member raises money from, but distinct from each.[17] Furthermore, although we compare member-level average scores in Figure 3, partisanship scores based on member communications are calculable using far shorter time spans than either roll call or contribution based measures. In our applications, we show how researchers can take advantage of this to examine, for instance, changes in partisan positioning before versus after a member's primary election.

The Evolution of Online Congressional Communication

How has congressional rhetoric evolved in the age of social media? Our data, which spans back to almost the beginning of widespread social media usage, is well equipped to answer this question.

We first explore how *partisanship* in congressional rhetoric has evolved across time. Existing research has clearly demonstrated that polarization, measured via roll call voting patterns and in other ways, has been increasing in Congress during this time period.[18] Do we see an analogous increase in partisan rhetoric?

To evaluate this question, we plot the average Partisan Extremity Score by calendar day for all three forms of communication and each of the two major political parties across our entire time period. A smoothed GAM regression line is fit to each of the time series, to flexibly capture

[16]For comparison, the within-party correlation between DW-NOMINATE scores and CF scores is 0.58 for Republicans and 0.20 for Democrats. These within-party correlations are also similar to those of other ideology measures commonly used in political science research (Tausanovitch and Warshaw 2017).

[17]Furthermore, in many cases we believe our scalings have more facial validity than other measures. For example, members of "The Squad" (www.usatoday.com/story/news/politics/elections/2024/11/05/aoc-wins-squad-reelection-results/76081070007/) are popularly described as some of the most liberal Democrats in Congress. Out of 447 House Democrats in our dataset, Squad members Alexandria Ocasio-Cortez, Ilhan Omar, Ayanna Pressley, Rashida Tlaib, Jamaal Bowman, and Cori Bush rank between 12th and 121st most left-leaning. By first dimension DW-NOMINATE scores, these members rank between the 254th to 350th most liberal House Democrats.

[18]For a summary of this research, see McCarty (2019).

changes across the period. The resulting plots are shown in Figure 4.

FIGURE 4: TRENDS IN PARTISAN RHETORIC (2009 - 2022)

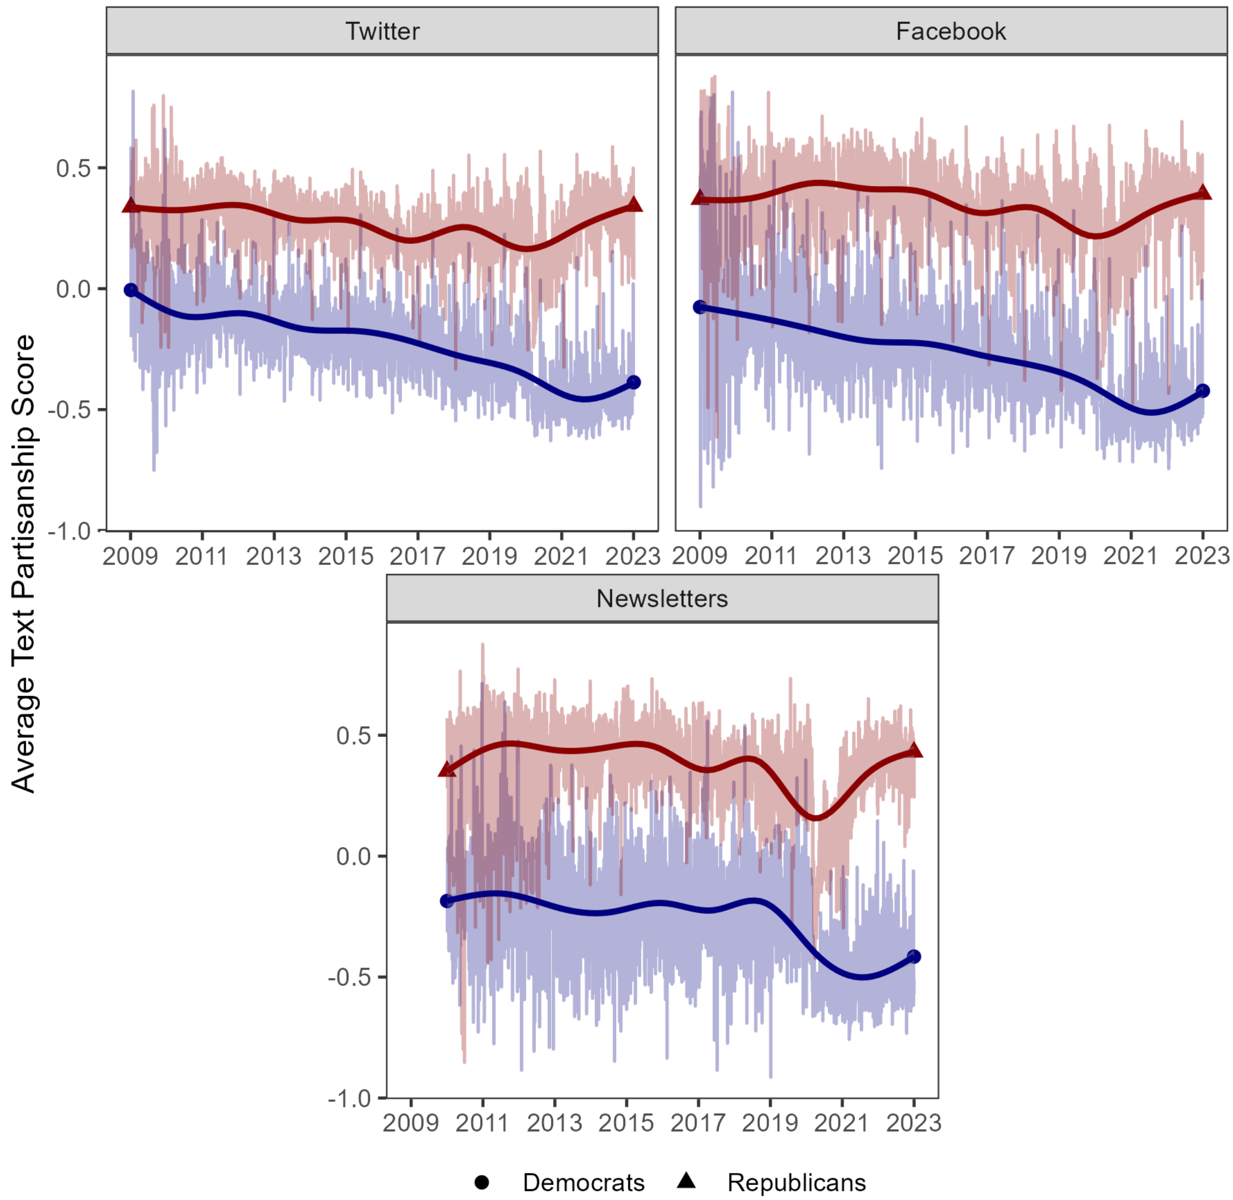

Note: The figure displays the daily average Text Partisanship Score for tweets, Facebook posts, and newsletter sentence bigrams (higher values indicate more Republican rhetoric). Dark lines indicate smoothed GAM regression lines of best fit. Trends shown separately for Democrats (blue; endpoints denoted with circles) and Republicans (red; endpoints denoted with triangles).

As can be seen in Figure 4, the partisan divide in Congressional rhetoric has grown steadily over this time period across all three forms of communication, but particularly on Facebook and Twitter. In 2009, the difference between the language used by Republicans and Democrats on social media was relatively small, approximately 0.4 points on the 2-point scale for both Twitter and Facebook. By 2022, this gap had more than doubled, to approximately 0.9 points. The trend is gradual, with the sole exception of a leftward shift by members of both parties towards the beginning of the COVID-19 pandemic. This short-term depolarization of rhetoric at the beginning of COVID has been identified in existing research using alternative methods (Heseltine 2024; see also Green et al. 2020).

One noteworthy feature of this polarization in congressional rhetoric is that the growth in partisanship is driven almost entirely by Democrats. Besides the dip in the immediate aftermath of COVID-19, there is no clear trend for Republicans. In contrast, there is a steady leftward shift in Democratic rhetoric on Twitter and Facebook. In email newsletters, partisanship does not change for Democrats until 2020.

To evaluate these across-time changes more rigorously, we estimate a series of OLS regression models at the level of an individual message where the dependent variable is the scaled partisan extremity of the tweet, Facebook post, or newsletter sentence bigrams, and the only independent variables are binary indicators for the Congressional session of the message. The intercept is omitted from these models, which are estimated separately for each communication form-party combination, meaning the coefficients represent the average partisan extremity of messages on a particular platform by a particular party in a given Congress. Standard errors are clustered by member. The results are shown in Figure 5.

The by-session regression estimates confirm the trends displayed in Figure 4. Partisan extremity in rhetoric is increasing for Democrats across the entire time period on Twitter and Facebook, and in the latter two sessions in the newsletters. There is no clear trend for Republicans on any platform.

FIGURE 5: PARTISAN EXTREMITY BY PARTY AND SESSION OF CONGRESS (111TH - 117TH)

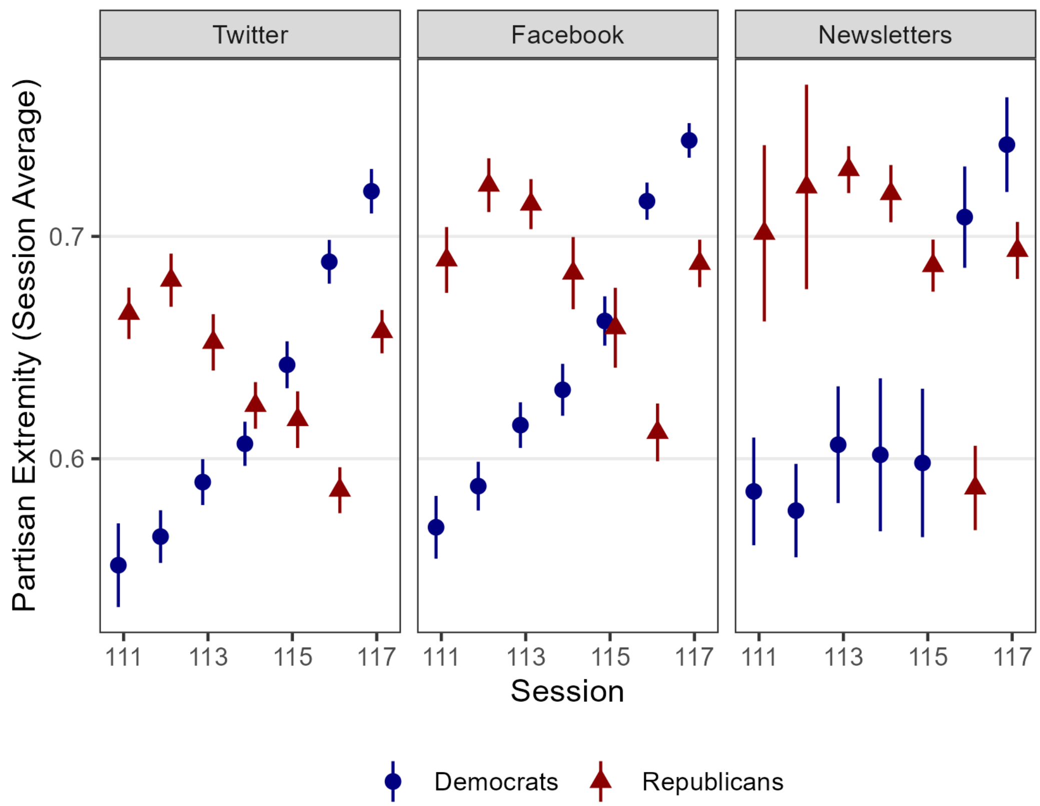

Note: The figure displays the average partisan extremity of tweets, Facebook posts, and newsletter sentence bigrams for each party in each session of Congress. Standard errors are clustered by member. Solid lines show 95% confidence intervals. Full regression results shown in Table G.1 in the Supplemental Materials.

From one perspective, this asymmetric polarization is surprising, given that (for a longer period of time than we have data here), roll call voting polarization in Congress appeared to be larger for Republicans than Democrats (McCarty 2019). On the other hand, the DW-NOMINATE scaling procedure has been criticized as failing to capture "ends-against-the-middle" voting dynamics that have become more common in recent years (Duck-Mayr and Montgomery 2023). Additionally, ideal point scaling of state legislators shows Democrats moving farther to the left than Republicans have moved towards the right during the time period studied here, supporting the idea that the 2010s was a decade of increasing liberalism among Democratic office-holders

(Shor and McCarty 2022).[19]

Supporting the finding that partisanship in general has become more common in congressional rhetoric in recent years, Figure 6 shows the trend in position taking and negative partisanship on Twitter, Facebook, and newsletters across time. On the two social media platforms, both position taking and negative partisanship have become more common. At the beginning of the time period, approximately 30% of tweets and 35% of Facebook posts by members featured position-taking, while 10% or fewer messages contained negative partisan attacks. By the end of the time period, between 40-45% of tweets and Facebook posts featured position taking, while approximately 15% of messages contained negative partisan attacks.

This increase in partisan attacks and position-taking appears to be quite similar between Democrats and Republicans. On the other hand, there is less evidence of a trend in either category in members' newsletters. Section H of the Supplemental Materials displays the across-time trends for the remaining four categories. As Figure H.1 shows, among these four categories, the clearest trend is a decrease in advertising across the two social media platforms. As in Figure 6, this appears common to both Democrats and Republicans.

To summarize, the partisanship scalings and message classifications tell a common story. In the years since social media has emerged and become ubiquitous both among the public as well as elected officeholders, online communication by members of Congress has become more partisan and more negative in tone, particularly on social media platforms themselves.

The Social Media Feedback Mirage

An advantage of the social media data is that posts and tweets contain information about public engagement in the form of likes, shares, retweets, and other metrics. These metrics repre-

[19] In contrast, Heseltine (2023) finds more polarization on social media by Republican members of Congress, as measured by the external news organizations they reference or link to in their tweets and posts. Thus while our measure shows Democratic members of Congress have been using more polarizing language, Heseltine (2023) shows Republican members have been sharing more partisan information sources.

FIGURE 6: TRENDS IN NEGATIVITY AND POSITION-TAKING (2010 - 2022)

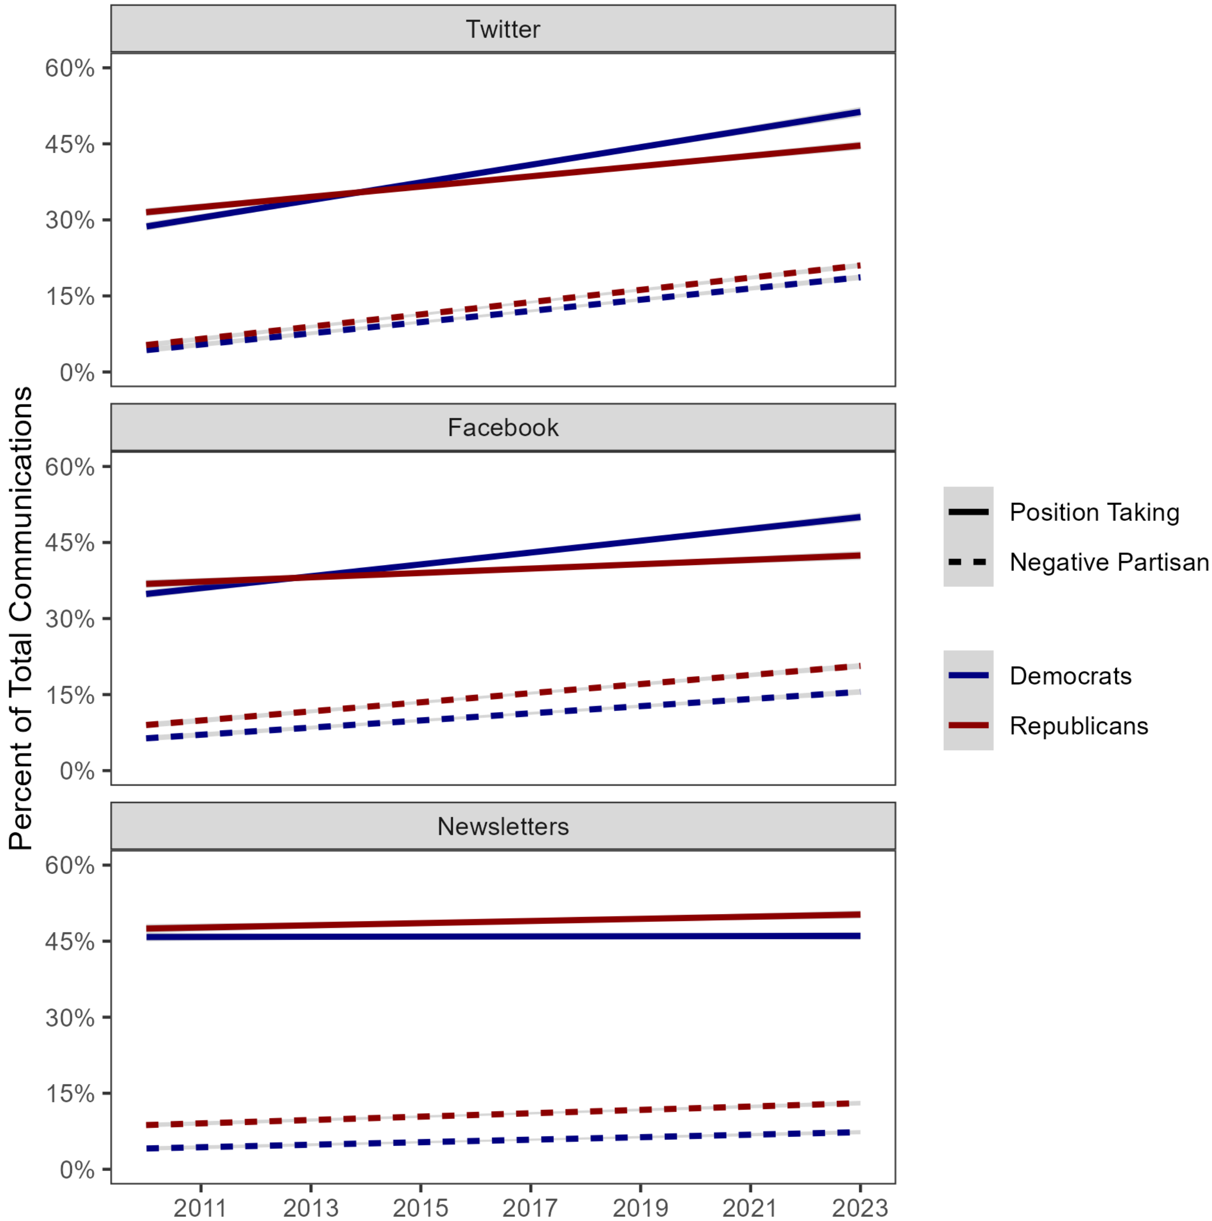

Note: The figure displays linear trends in the average percent of congressional communication (aggregating tweets, Facebook posts, and newsletters) classified into the position-taking and negative partisanship categories for both of the two major parties between 2010 and 2022.

sent the information members and their staff receive in real time, showing how followers respond to the messages they share. Which messages result in the most positive feedback for members?

Prior research has found that incivility or negativity on social media leads to more engagement in more limited data (Ballard et al. 2022; Macdonald, Russell, and Hua 2023; Yu, Wojcieszak, and Casas 2024). We replicate and extend this research, examining engagement as a function of partisan extremity and our six purposive categorizations on both Facebook and Twitter across a 14-year period.

To do so, we estimate a series of OLS regression models where the unit of observation is an individual tweet or post. For each message, we include six binary variables for whether the tweet is classified into each of our categories, as well as the continuous (0 to 1) Tweet Partisan Extremity Score. The dependent variable is the number of likes, retweets, or shares the message received (logged to address right-skew). Regressions include member fixed effects. As a consequence, coefficients can be interpreted as the difference in expected likes or retweets relative to the average tweet by a member that is not extreme and does not credit claim, advertise, position-take, etc. We also include legislative session fixed effects, to account for changes in engagement and messaging across time.

The results are displayed in Figure 7. There are clear differences in positive engagement across message types. The most striking pattern is the high levels of engagement that negative partisan messages receive. Negative partisan attacks on Twitter receive 86% more likes and 102% more retweets than a member's average tweet; similarly, negative partisan attacks on Facebook receive 32% more likes and 113% more shares than a member's average Facebook post. Position taking on social media also received much higher levels of engagement than the typical member message. A message's partisan extremity score is associated with significantly more retweets and shares, though not likes. Of the other categories, only bipartisan messages receive significantly higher engagement than the typical message, and only on Facebook. Credit claiming (of either type) and advertising tend to receive less positive engagement on social media.

Does this difference in message reception by social media users matter? On the one hand, reelection-focused members should care most about how communication styles shape the views

FIGURE 7: Message Content and Social Media Engagement

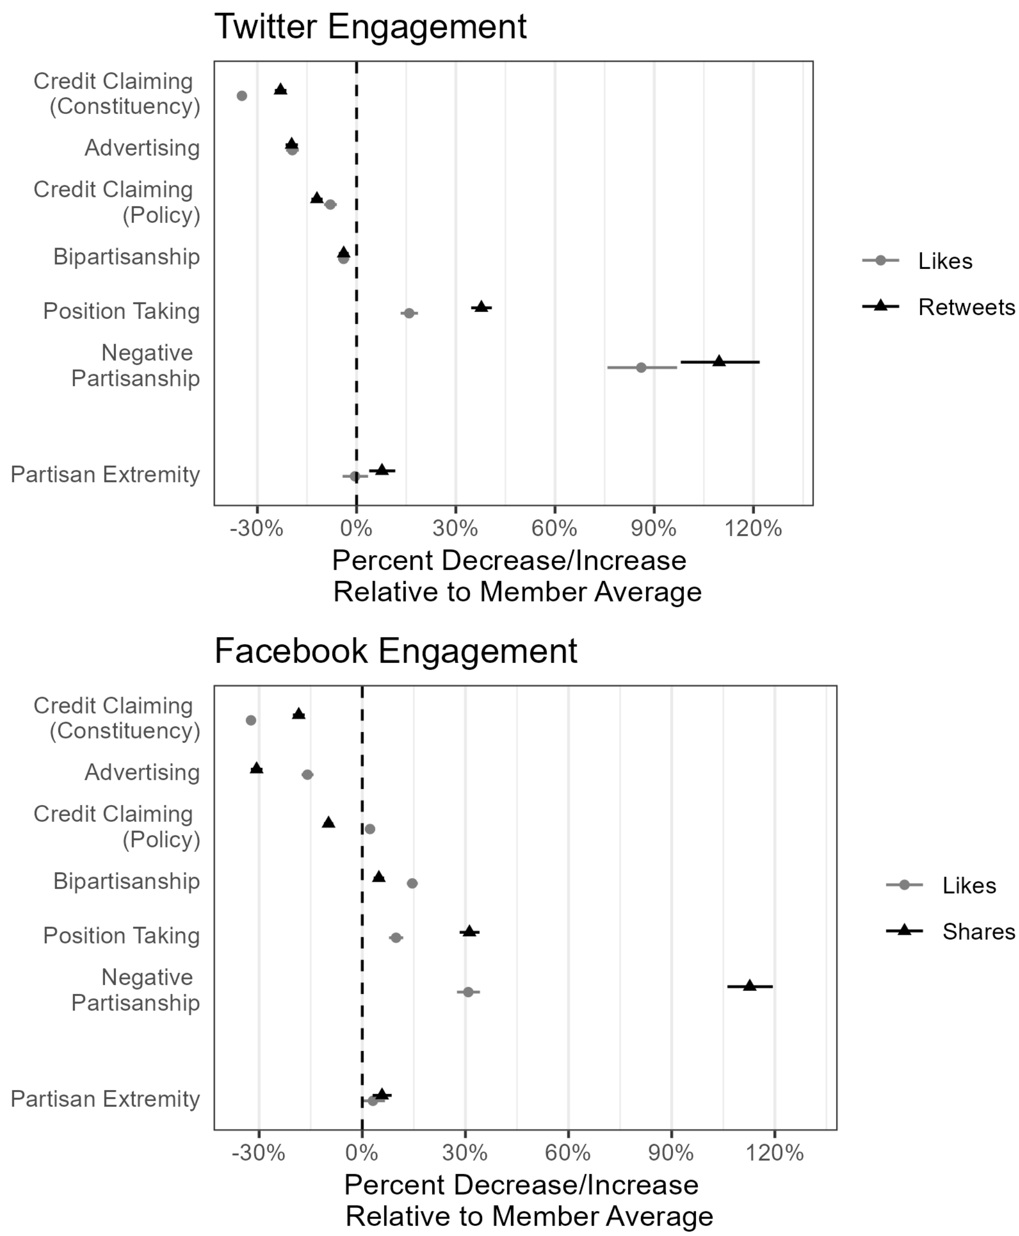

Note: The figure displays OLS coefficient estimates and 95% confidence intervals from the full model results shown in Table I.1. Coefficient estimates are transformed ($\exp(\hat{\beta})-1$) to percent differences for interpretability. Independent variables, shown on the y-axis, consist of message purpose and partisanship. All models include member fixed effects, meaning the coefficient represents the number of likes or retweets relative to a member's average. Standard errors are clustered by member. Full regression results shown in Table I.1 in the Supplemental Materials.

of voters in their districts and states. On the other hand, social media sites like Twitter and Facebook provide instantaneous feedback on message reception to legislators and their staff in a way that they do not receive from other forms of communication or other audiences. This could potentially lead to mistaking engagement on social media with broader approval.

Two pieces of evidence suggest that the general public does not respond as positively to negativity, position-taking, and partisanship as social media engagement metrics might suggest.

First, experimental work consistently finds that partisan messages are among the *least* preferred of the messaging types. For example, Simas et al. (2025) show that in almost all cases, even in-partisans report significantly greater satisfaction when exposed to legislators who post position-taking and credit-claiming messages than they do when exposed to legislators who engage in partisan posturing. Costa (2021) reports similar findings, as respondents across three studies all express lower satisfaction with and intention to vote for legislators who use negative partisanship in their messaging.[20]

Second, we merge our data on member communication on the three platforms with constituency approval ratings from Cooperative Election Study (CES) surveys. This allows us to evaluate whether messaging portfolios in aggregate are associated with higher or lower approval for members of Congress from their constituents. The full analysis is provided in Section J of the Supplemental Materials. In general, however, we find that position-taking, negativity, and partisan extremity are either not associated or are negatively associated with constituent approval.[21]

Based on these evidence sources, it does not seem that negativity and partisanship resonate with the general public, despite the positive reception by social media users. This potentially misleading discrepancy we term the *social media feedback mirage*. While we cannot fully dissect the reason for the divergent responses of social media users and the public, a likely culprit is

[20] While issue messages appear to receive significantly more positive evaluations, this is shown to be conditional on agreement with the position taken and the partisanship of the individual.

[21] The one exception is that same-party constituents (not general constituents or independents) approve more highly of members who use more negative partisan attacks.

selection effects. McCabe, Green, Goel and Lazer (2023) link administrative data to Twitter survey records to show that the average member of the 115th Congress was followed by only 2.4% of in-district Twitter users. These followers tended to have high levels of political interest and be the same party as their member of Congress.

Regardless of the source behind this divide, if legislators and their staff respond to the feedback they receive on social media messaging, social media has the potential to amplify content that ultimately alienates broader voters, potentially leading to broader dissatisfaction and cynicism. Without more study beyond what is possible in this paper, it is impossible to know whether politicians and their staff are influenced at all by the response to posts on social media or are able to completely adjust for these differences in mental calculations. Given the importance of this question for democratic discourse in contemporary politics, further investigation of this possibility – potentially via interviews or surveys of congressional communications staffers – is merited. It is noteworthy, however, that the messaging types that have increased in frequency the most during the era of growing social media usage (position taking, negative partisanship, and extreme partisan rhetoric, as shown in Figures 5 and 6) are the types of messages that receive the most positive engagement on social media.

How Primary Elections Affect Messaging Strategies

For a final application, we demonstrate how our dataset can be used to draw causal inferences about how and why representatives communicate. For this analysis, we take advantage of the staggered nature of primary election dates by state to conduct a difference-in-differences analysis comparing how members communicate before a primary election versus how members communicate after their primary has passed (but before the general election).

A major benefit to using social media posts and other messaging data to test theories of congressional behavior is that the data are high-frequency and the timing of each message can be determined (depending on the data source) to the second. Even analyzing social media messages by day offers advantages over other forms of data.[22]

We can take advantage of this to address a longstanding debate in the American politics literature: do primary elections contribute to polarization? If so, how? Recent scholarship presents conflicting evidence on this topic. Brady, Han, and Pope (2007) offer support for the polarizing influence of primaries, finding that primary voters favor ideologically extreme candidates, which pulls congressional candidates away from median district preferences. Similarly, Anderson, Butler, and Harbridge-Yong (2020) argue that fear of primary voter punishment leads legislators to reject compromises, exacerbating legislative gridlock. However, several studies challenge this primary-polarization link. Neither McGhee et al. (2014) nor Rogowski and Langella (2015) find that primary rules have an effect on legislator extremism. Boatright (2014) examines specific instances of congressional incumbents getting challenged in primaries, and finds no significant changes in voting patterns. Finally, Hirano and Snyder (2019) push back against the polarizing nature of primaries as well. They note that primaries were introduced decades before the recent increase in party polarization and provide evidence that primary voters can be strategic in choosing candidates with the best chance of winning the general election, as opposed to the most partisan candidates.

These studies all use roll call voting or ideal point measures based on roll call voting to quantify polarized behavior. Primaries may have polarizing effects on outcomes other than roll call votes, however, and these effects may be transient and short-lived. Do members of Congress change their messaging based on whether or not they have a primary election in the near future? Research that has used social media data to determine if members of Congress moderate post-

[22]For example, in a typical year in recent Congresses, members of the US House will cast 500 to 700 floor roll call votes, while senators will vote approximately 250 to 350 times (see www.congress.gov/roll-call-votes, for instance). And because Congress typically votes on multiple items during a single day when they're in session, rather than holding sessions every day of the year, votes only occur on approximately 80 to 90 calendar days. In contrast, social media posts and newsletters are sent out daily. In our data, the fewest number of messages (combining tweets, Facebook posts, and newsletters) was 66 separate messages. The median number of messages per day is 1,462 messages.

primary has examined only the 2020 election cycle (Cowburn and Sältzer 2024; Macdonald et al. 2025), in each case drawing mixed conclusions.

Instead, we use seven election cycles of primary elections to greatly increase our statistical power and evaluate any changes across time. Our research design takes advantage of the fact that primary elections are held on different dates depending on the state, meaning at any given point in an election cycle some congressional incumbents have a primary election coming up while others have passed the primary election and no longer need to target primary voters in their communications (either because they won, and must focus on general election voters, or they lost, and do not have to worry about any electoral considerations whatsoever).

To analyze how the presence of a forthcoming primary election matters, we estimate a series of two-way fixed effects regression models at the level of an individual message (tweet, post, or newsletter), where the outcome variable is the scaled Partisan Extremism of the message or one of the six purposive classifications. The treatment variable in this context is a binary indicator for *Post-Primary*, coded as 1 if the date of the primary election for the member who shared the tweet, post, or newsletter has already passed, and 0 otherwise.[23] We include fixed effects for each day to flexibly account for any factors at any given point in time that might make messages more or less partisan, and include member-cycle fixed effects to account for all differences between members that do not vary over the course of the election cycle. Standard errors are clustered by member.

The results of these difference-in-differences models are shown in Figure 8. We estimate these models for Democrats and Republicans separately, to allow for the possibility that members of the two parties target primary voters in different ways. We also estimate these models using the full timespan of our data (2010 to 2022) as well as estimating the models separately for each six single election cycles, reflecting the amount of communications data most prior work has used

[23]Data on primary election dates was collected from www.FEC.gov. In these regression models, we only include House members since, unlike Senators, all House members must run for re-election every two years. Only messages shared during an election year after the first primary date but before the last primary date are included, ensuring that at any given point, some messages are shared by members in the pre-primary stage of the election cycle and others are shared by members in the post-primary stage.

FIGURE 8: Differences in Message Content After Versus Before the Primary Election

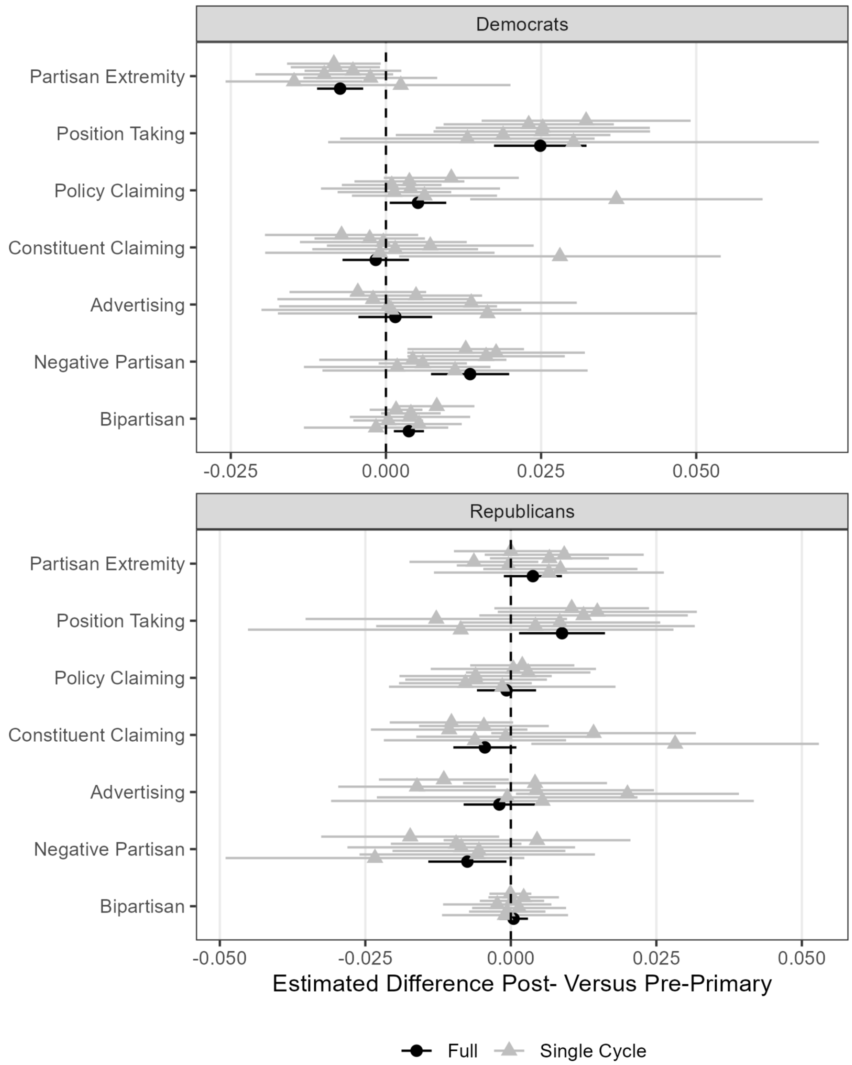

Note: The figure displays the coefficient estimates from difference-in-differences models comparing message content for members who, on a given day in an election cycle, have passed their primary election to members whose primary election is still upcoming. All models include member-session and calendar date fixed effects, with standard errors clustered by members. The coefficient estimates are thus interpretable as the estimated effect on messaging of being post-primary. Solid lines display 95% confidence intervals. Black point estimates and confidence intervals are from models using the full timespan of the data (2010 to 2022), while grey point estimates and confidence intervals are from single election cycles (order from earliest to latest, moving upwards). Full regression results shown in Table K.1 in the Supplemental Materials.

in their analyses, to demonstrate the practical advantages of possessing communications data as large as the dataset we introduce in terms of statistical power.

Figure 8 reveals that members do change their messaging depending on whether they have a primary election upcoming or not, but the changes are modest, and differ somewhat depending on party. Most relevant to the polarization question, when using the full timespan of the data we find that Democratic messaging becomes less extreme after members pass their primary election date. The change is small, however, approximately 0.01 points (on a scale that ranges from 0 to 1). Similarly, messages by Democratic members are more likely to promote bipartisanship after they pass their primary election date, consistent with what Anderson, Butler, and Harbridge-Yong (2020) find, although this effect is small as well. Messages are also more likely to include negative attacks against the other party post-primary, suggesting members may pivot to criticizing their opposite-party general election challenger after the primary election stage has concluded. For Republicans, there are no statistically significant differences in partisan extremity or bipartisanship, and only a small and marginally significant difference in negative partisanship. The largest messaging difference pre- versus post-primary for members of both parties is in position-taking, which becomes more common after the primary election, perhaps reflecting the dangers of taking strong positions on topics during the primary election that may alienate one wing of a member's party or another.[24]

All of these effect sizes are quite small, however, and – as the confidence intervals of the gray estimates show – most of the statistically significant effects we observe would not be statistically significant if researchers were to use data from a single election cycle alone. While we can rule out null hypotheses of precisely zero effect, the magnitudes of the effect sizes we observe generally rule out large short-term effects of primary elections.[25] Members of Congress simply do not change the content of their rhetoric much to appeal to primary voters, at least not in ways that vary within a single election cycle.

[24]It may also be the case that there is less room to differentiate on policy in a primary election, where candidates from the same party will presumably take more similar positions.

[25]This does not preclude that the threat of primary elections induce more stable changes in partisanship and ideology by members relative to a counterfactual world with no or different primary elections.

Conclusion

Communications from elected officials both facilitate effective representation (e.g. Pitkin 1967) and shape the subject and tenor of public conversation (Barberá et al. 2019; Lippmann 1922; Zaller 1992; Jacobs and Shapiro 2000). As such, studying the content of those communications is central to numerous strands of the political science literature. Work in this area has attempted to keep pace with the changes that have come with the wide-spread adoption of online communications, but efforts have been hampered by the challenge of collecting such large volumes of data over time and across platforms. Our Scaled and Classified Congressional Communication (SCCC) dataset, spanning approximately 5 million tweets, 2.5 million Facebook posts, and 184,000 email newsletters authored by members of Congress between 2009 and 2022, helps advance study in this area in several ways. The size and scope of this dataset not only resolves practical issues related to data collection, but also allows for more comprehensive investigations of how communications change over time and vary across modes. Moreover, the coding of both purpose and partisanship allow for more precise estimates of the effects of the various components of online messages. Altogether, the SCCC is one of the most comprehensive resources available for studying how characteristics of politicians' rhetoric affect a host of key political outcomes.

Indeed, the applications presented here demonstrate the vast potential of this dataset to make contributions beyond just the more focused study of online congressional communications. Our findings that negative and partisan messages are increasingly prevalent and attention-grabbing speak to work on polarization, while our findings that these same types of messages do not appear to be tailored toward primary voters nor rewarded by general election constituencies have implications for work on elections and responsiveness. And while each of these findings is important on its own, they also raise some interesting questions when considered in combination. Namely, why has online communication grown and evolved to be more partisan in nature when our analyses suggest that there are no apparent payoffs from voters? While some of this may be due to disconnect between the feedback received from social media users vs. direct constituents (i.e., the social media feedback mirage), it seems unlikely that savvy politicians would continue to pursue strategies that do not help and possibly hinder their goals in some form. This suggests that more work is needed to gain a complete understanding of the motivations behind and consequences of politicians' communications.

It should also not be lost that while a major strength of our approach is that it offers estimates of both style and partisanship. The partisan ratings alone are an extremely valuable tool for approximating legislators' positions, in some cases providing more facially valid estimates than alternative such as DW-NOMINATE. The two widely used measures of legislative positioning we compare our measures to above are derived from legislative voting and fundraising. Though the public has access to records of both, the majority of individuals lack knowledge about their representatives' activities in these areas. Our ratings, in contrast, are derived from visible and easily accessible communications that are crafted to portray a legislator as they want to be seen. So when used in combination with these other types of measures, our partisan scaling has the potential to offer a fuller picture of legislators' preferences and speak to questions of whether those preferences appear relatively consistent when estimated from these different sources.

Our dataset, as currently constructed, ends in 2022. In the very short time since then, social media usage by politicians has continued to evolve. One of the most notable examples is the rise of new social media platforms that are associated with particular points on the ideological spectrum: Truth Social and Bluesky. While the era from 2009 to 2022 was characterized primarily by monolith platforms in the social media space that catered to all political proclivities (including none), it is possible that the next era might feature even more fractionalization, with politically-oriented media consumers choosing platforms that cater in a focused way to their tastes. On the other hand, even among the population of social media users most do not follow their representatives on social media McCabe et al. (2023), perhaps limiting the impact of this type of political Balkanization. A clear path forward for scholars of political communication is unpacking the effects these new platforms have on rhetoric and discourse.

Though politics and political communication has changed drastically in the past half century, the fact remains that "if there is to be congruence between the policy preferences of the represented and the policy decisions of the representatives, however, two-way communication between them is a prerequisite" (Fenno 1978, p.241). We offer a comprehensive resource for the continuing and evolving study of representation in the U.S. The possibilities our data unlock are too many to list here; we are confident that scholars from across the discipline will benefit from having such a resource at their disposal, and use it in ways beyond what we imagine here.

References

Albert, Zachary. 2020. "Partisan Policymaking in the Extended Party Network: The Case of Cap-and-Trade Regulations." *Political Research Quarterly* 73 (2): 476–491.

Anderson, Sarah E., Daniel M. Butler, and Laurel Harbridge-Yong. 2020. *Rejecting Compromise: Legislators' Fear of Primary Voters.* Cambridge: Cambridge University Press.

Ashworth, Scott. 2012. "Electoral Accountability: Recent Theoretical and Empirical Work." *Annual Review of Political Science* 15 (1): 183–201.

Ballard, Andrew O., Ryan DeTamble, Spencer Dorsey, Michael Heseltine, and Marcus Johnson. 2023. "Dynamics of Polarizing Rhetoric in Congressional Tweets." *Legislative Studies Quarterly* 48 (1): 105–144.

Ballard, Andrew, Ryan DeTamble, Spencer Dorsey, Michael Heseltine, and Marcus Johnson. 2022. "Incivility in Congressional Tweets." *American Politics Research* 50 (6): 769–780.

Barberá, Pablo. 2015. "Birds of the Same Feather Tweet Together: Bayesian Ideal Point Estimation Using Twitter Data." *Political Analysis* 23 (1): 76–91.

Barberá, Pablo, Andreu Casas, Jonathan Nagler, Patrick J. Egan, Richard Bonneau, John T. Jost, and Joshua A. Tucker. 2019. "Who Leads? Who Follows? Measuring Issue Attention and Agenda Setting by Legislators and the Mass Public Using Social Media Data." *American Political Science Review* 113 (4): 883–901.

Bernhard, William, Daniel Sewell, and Tracy Sulkin. 2017. "A Clustering Approach to Legislative Styles." *Legislative Studies Quarterly* 42 (3): 477–506.

Bernhard, William, and Tracy Sulkin. 2018. *Legislative Style.* The University of Chicago Press.

Besley, Timothy. 2006. *Principled Agents? The Political Economy of Good Government.* Oxford: Oxford University Press.

Blum, Rachel, Lindsey Cormack, and Kelsey Shoub. 2023. "Conditional Congressional Communication: How Elite Speech Varies Across Medium." *Political Science Research and Methods* 11 (2): 394–401.

Boatright, Robert G. 2014. *Getting Primaried: The Changing Politics of Congressional Primary Challenges.*

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

Brady, David W., Hahrie Han, and Jeremy C. Pope. 2007. "Primary Elections and Candidate Ideology: Out of Step with the Primary Electorate?" *Legislative Studies Quarterly* 32 (1): 79–105.

Burke, Edmund. 1774. "Speech to the Electors of Bristol." In *The Works of the Right Honourable Edmund Burke.* Vol. 1 London: Henry G. Bohn pp. 446–448. Published 1854–56.

Butler, Daniel M., Thad Kousser, and Stan Oklobdzija. 2023. "Do Male and Female Legislators Have Different Twitter Communication Styles?" *State Politics & Policy Quarterly* 23 (2): 117–139.

Cohen, Joshua. 1989. "Deliberation and Democratic Legitimacy." In *The Good Polity: Normative Analysis of the State*, ed. Alan Hamlin, and Phillip Petit. New York: Blackwell.

Cormack, Lindsey. 2016*a*. "Extremity in Congress: Communications versus Votes." *Legislative Studies Quarterly* 41 (3): 575–603.

Cormack, Lindsey. 2016*b*. "Gender and Vote Revelation Strategy in the United States Congress." *Journal of Gender Studies* 25 (6): 626–640.

Cormack, Lindsey. 2017. "DCinbox - Capturing Every Congressional Constituent E-newsletter from 2009 Onwards." *The Legislative Scholar* 2 (1): 27–36.

Costa, Mia. 2021. "Ideology, Not Affect: What Americans Want from Political Representation." *American Journal of Political Science* 65 (2): 342–358.

Cowburn, Mike, and Marius Sältzer. 2024. "Partisan Communication in Two-Stage Elections: The Effect of Primaries on Intra-Campaign Positional Shifts in Congressional Elections." *Political Science Research and Methods* pp. 1–20.

Duck-Mayr, JBrandon, and Jacob Montgomery. 2023. "Ends Against the Middle: Measuring Latent Traits when Opposites Respond the Same Way for Antithetical Reasons." *Political Analysis* 31 (4): 606–625.

Evans, Heather K., and Jennifer Hayes Clark. 2016. ""You Tweet Like a Girl!": How Female Candidates Campaign on Twitter." *American Politics Research* 44 (2): 326–352.

Evans, Heather K., Victoria Cordova, and Savannah Sipole. 2014. "Twitter Style: An Analysis of How House Candidates Used Twitter in Their 2012 Campaigns." *PS: Political Science & Politics* 47 (2): 454–462.

Fenno, Richard F. 1978. *Home Style: House Members in Their Districts.* Scott, Foresman.

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.

Frankel, Laura Lazarus, and D. Sunshine Hillygus. 2017. "Niche Communications in Political Campaigns." In *The Oxford Handbook of Political Communication*, ed. Kate Kenski, and Kathleen Hall Jamieson. Oxford University Press.

Fu, Shu, and William G. Howell. 2020. "The Behavioral Consequences of Public Appeals: Evidence on Campaign Fundraising from the 2018 Congressional Elections." *Presidential Studies Quarterly* 50 (2): 325–347.
URL: *http://onlinelibrary.wiley.com/doi/abs/10.1111/psq.12645*

Gainous, Jason, and Kevin M. Wagner. 2013. *Tweeting to Power: The Social Media Revolution in American Politics*. Oxford University Press.

Gentzkow, Matthew, Jesse M. Shapiro, and Matt Taddy. 2019. "Measuring Group Differences in High-Dimensional Choices: Method and Application to Congressional Speech." *Econometrica* 87 (4): 1307–1340.
URL: *https://www.econometricsociety.org/doi/10.3982/ECTA16566*

Green, Jon, Jared Edgerton, Daniel Naftel, Kelsey Shoub, and Skyler J. Cranmer. 2020. "Elusive Consensus: Polarization in Elite Communication on the Covid-19 Pandemic." *Science Advances* 6 (28): eabc2717.
URL: *https://www.science.org/doi/10.1126/sciadv.abc2717*

Green, Jon, Kelsey Shoub, Rachel Blum, and Lindsey Cormack. 2024. "Cross-Platform Partisan Positioning in Congressional Speech." *Political Research Quarterly* 77 (3): 653 – 668.

Grimmer, Justin. 2013. *Representational Style in Congress: What Legislators Say and Why It Matters*. Cambridge University Press.

Grimmer, Justin, Sean J. Westwood, and Solomon Messing. 2015. *The Impression of Influence: Legislator Communication, Representation, and Democratic Accountability*. Princeton University Press.

Habermas, Jürgen. 1968. *The Theory of Communicative Action. Vol. 1. Reason and the Rationalization of Society*. Boston: Beacon Press.

Hayes, Danny, and Jennifer L. Lawless. 2015. "As Local News Goes, So Goes Citizen Engagement: Media, Knowledge, and Participation in US House Elections." *The Journal of Politics* 77 (2): 447–462.

Hayes, Danny, and Jennifer L. Lawless. 2017. "The Decline of Local News and Its Effects: New Evidence from Longitudinal Data." *The Journal of Politics* 80 (1): 332–336.

Hemphill, Libby, Annelise Russell, and Angela M. Schöpke-Gonzalez. 2021. "What Drives U.S. Congressional Members' Policy Attention on Twitter?" *Policy & Internet* 13 (2): 233–256.

Heseltine, Michael. 2023. "Asymmetric Polarization in Online Media Engagement in the United States Congress." *The International Journal of Press/Politics* 30 (1): 299–325.
URL: *https://doi.org/10.1177/19401612231211800*

Heseltine, Michael. 2024. "Polarizing Online Elite Rhetoric at the Federal, State, and Local Level During the COVID-19 Pandemic." *American Politics Research* 52 (3): 187–202.

Heseltine, Michael, and Bernhard Clemm von Hohenberg. 2024. "Large language models as a substitute for human experts in annotating political text." *Research & Politics* 11 (1): 20531680241236239.
URL: *https://doi.org/10.1177/20531680241236239*

Hirano, Shigeo, and James M. Snyder. 2019. *Primary Elections in the United States.* Political economy of institutions and decisions Cambridge: Cambridge University Press.

Hunt, Charles R., and Kristina C. Miler. 2025. "How Modern Lawmakers Advertise Their Legislative Effectiveness to Constituents." *The Journal of Politics* 87 (1): 231–245.
URL: *https://www.journals.uchicago.edu/doi/10.1086/730732*

Jacobs, Lawrence R, and Robert Y Shapiro. 2000. *Politicians don't pander: Political manipulation and the loss of democratic responsiveness.* University of Chicago Press.

Jungherr, Andreas. 2014. "The Logic of Political Coverage on Twitter: Temporal Dynamics and Content." *Journal of Communication* 64 (2): 239–259.

Jungherr, Andreas, and Ralph Schroeder. 2022. *Digital Transformations of the Public Arena.* Elements in Politics and Communication Cambridge University Press.
URL: *https://www.cambridge.org/core/elements/digital-transformations-of-the-public-arena/6E4169B5E1C87B0687190F688AB3866E*

Kaslovsky, Jaclyn, and Michael R. Kistner. 2025. "Responsive Rhetoric: Evidence from Congressional Redistricting." *Legislative Studies Quarterly* 49 (4).

Laver, Michael, Kenneth Benoit, and John Garry. 2003. "Extracting Policy Positions from Political Texts Using Words as Data." *American Political Science Review* 97 (2): 311–331.

Lippmann, Walter. 1922. *Public opinion.* Harcourt, Brace, and Company.

Macdonald, Maggie, Annelise Russell, and Whitney Hua. 2023. "Negative Sentiment and Congressional Cue-Taking on Social Media." *PS: Political Science & Politics* 56 (2): 201–206.
URL: *https://www.cambridge.org/core/journals/ps-political-science-and-politics/article/negative-sentiment-and-congressional-cuetaking-on-social-media/1E0DF82E947BBE25B35D28AEB203B2B2*

Macdonald, Maggie, Megan A. Brown, Joshua A. Tucker, and Jonathan Nagler. 2025. "To Moderate, or Not to Moderate: Strategic Domain Sharing by Congressional Campaigns." *Electoral Studies* 95.

Macdonald, Maggie, Whitney Hua, and Annelise Russell. 2024. "Constrained Communication and Negativity Bias: Gendered Emotional Appeals on Facebook." *Journal of Women, Politics & Policy* 45 (2): 261–274.

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

Mayhew, David R. 1974. *Congress: The Electoral Connection.* Yale University Press.

McCabe, Stefan, Jon Green, Pranav Goel, and David Lazer. 2023. "Inequalities in Online Representation: Who Follows Their Own Member of Congress on Twitter?" *Journal of Quantitative Description: Digital Media* 3.

McCarty, Nolan. 2019. *Polarization: What Everyone Needs to Know.* Oxford': Oxford University Press.

McGhee, Eric, Seth Masket, Boris Shor, Steven Rogers, and Nolan McCarty. 2014. "A Primary Cause of Partisanship? Nomination Systems and Legislator Ideology." *American Journal of Political Science* 58 (2): 337–351.
URL: *http://onlinelibrary.wiley.com/doi/10.1111/ajps.12070/abstract*

Mill, John Stuart. 1861. *Considerations on Representative Government.* Cambridge: Cambridge University Press.

Munger, Kevin. 2023. "Temporal Validity as Meta-Science." *Research & Politics* 10 (3): 1–10.

Nguyen, Dat Quoc, Thanh Vu, and Anh Tuan Nguyen. 2020. BERTweet: A pre-trained language model for English Tweets. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*, ed. Qun Liu, and David Schlangen. Online: Association for Computational Linguistics pp. 9–14.
URL: *https://aclanthology.org/2020.emnlp-demos.2/*

Payson, Julia, Andreu Casas, Jonathan Nagler, Richard Bonneau, and Joshua A. Tucker. 2022. "Using Social Media Data to Reveal Patterns of Policy Engagement in State Legislatures." *State Politics & Policy Quarterly* 22 (4): 371–395.

Perry, Patrick O., and Kenneth Benoit. 2017. "Scaling Text with the Class Affinity Model." *arXiv* .
URL: *https://arxiv.org/abs/1710.08963*

Peterson, Erik. 2021. "Paper Cuts: How Reporting Resources Affect Political News Coverage." *American Journal of Political Science* 65 (2): 443–459.

Pitkin, Hanna F. 1967. *The Concept of Representation.* University of California Press.

Poole, Keith T., and Howard Rosenthal. 2000. *Congress: a Political-Economic History of Roll Call Voting.* Oxford University Press.

Rogowski, Jon C., and Stephanie Langella. 2015. "Primary Systems and Candidate Ideology: Evidence From Federal and State Legislative Elections." *American Politics Research* 43 (5): 846–871.
URL: *https://doi.org/10.1177/1532673X14555177*

Russell, Annelise. 2018*a*. "The Politics of Prioritization: Senators' Attention in 140 Characters." *The Forum* 16 (2): 331–356.

Russell, Annelise. 2018*b*. "U.S. Senators on Twitter: Asymmetric Party Rhetoric in 140 Characters." *American Politics Research* 46 (4): 695–723.

Russell, Annelise. 2021*a*. "Gendered Priorities? Policy Communication in the U.S. Senate." *Congress & the Presidency* 48 (3): 319–342.

Russell, Annelise. 2021*b*. *Tweeting is Leading: How Senators Communicate and Represent in the Age of Twitter.* Oxford University Press.

Scherpereel, John A., Jerry Wohlgemuth, and Audrey Lievens. 2018. "Does Institutional Setting Affect Legislators' Use of Twitter?" *Policy & Internet* 10 (1): 43–60.

Schroeder, Ralph. 2018. *Social Theory after the Internet: Media, Technology, and Globalization.* UCL Press.
URL: *http://www.jstor.org/stable/10.2307/j.ctt20krxdr*

Shor, Boris, and Nolan McCarty. 2022. "Two Decades of Polarization in American State Legislatures." *Journal of Political Institutions and Political Economy* 3 (3–4): 343–370.

Simas, Elizabeth N., David Hilden, Michael R. Kistner, and Jamie Wright. 2025. "Credit Claiming and Accountability for Legislative Effectiveness." *Center for Effective Lawmaking Working Paper Series* .
URL: *https://thelawmakers.org/wp-content/uploads/2025/01/Credit_Claiming.pdf*

Slapin, Jonathan B., and Sven-Oliver Proksch. 2008. "A Scaling Model for Estimating Time-Series Party Positions from Texts." *American Journal of Political Science* 52 (3): 705–722.
URL: *https://www.jstor.org/stable/25193842*

Smith, Sarah A., and Annelise Russell. 2022. "Different Chambers, Divergent Rhetoric: Institutional Differences and Policy Representation on Social Media." *American Politics Research* 50 (6): 792–797.

Straus, Jacob R., Raymond T. Williams, Colleen Shogan, and Matthew Glassman. 2016. "Congressional Social Media Communications: Evaluating Senate Twitter Usage." *Online Information Review* 40 (5): 643–659.

Tausanovitch, Chris, and Christopher Warshaw. 2017. "Estimating Candidates' Political Orientation in a Polarized Congress." *Political Analysis* 25 (2): 167–187.

Tillery, Alvin B. 2021. "Tweeting Racial Representation: How the Congressional Black Caucus Used Twitter in the 113th Congress." *Politics, Groups, and Identities* 9 (2): 219–238.

Tucker, Joshua, Andrew Guess, Pablo Barbera, Cristian Vaccari, Alexandra Siegel, Sergey Sanovich, Denis Stukal, and Brendan Nyhan. 2018. "Social Media, Political Polarization, and Political Disinformation: A Review of the Scientific Literature." *SSRN Electronic Journal* .
URL: *https://www.ssrn.com/abstract=3144139*

Volden, Craig, and Alan E. Wiseman. 2014. *Legislative Effectiveness in the United States Congress: the Lawmakers.* New York: Cambridge University Press.

Volden, Craig, and Alan E. Wiseman. 2018. "Legislative Effectiveness in the United States Senate." *The Journal of Politics* 80 (2): 731–735.

Warner, Seth B. 2023. "Analyzing Attention to Scandal on Twitter: Elites Sell What Supporters Buy." *Political Research Quarterly* 76 (2): 841–850.
URL: *https://doi.org/10.1177/10659129221114145*

Yiannakis, Diana Evans. 1982. "House Members' Communication Styles: Newsletters and Press Releases." *The Journal of Politics* 44 (4): 1049–1071.

Yu, Xudong, Magdalena Wojcieszak, and Andreu Casas. 2024. "Partisanship on Social Media: In-Party Love Among American Politicians, Greater Engagement with Out-Party Hate Among Ordinary Users." *Political Behavior* 46 (2): 799–824.

Zaller, John. 1992. *The nature and origins of mass opinion.* Cambridge university press.

Supplementary Materials

Measuring Partisanship and Representation in Online Congressional Communications

Contents

A Past Research on Online Congressional Communication

TABLE A.1: Summary of Online Congressional Communication Studies

Author(s)	Medium	Chamber(s) & Time Period
Ballard et al. (2022)	Twitter	Both, 2009–2020
Ballard et al. (2023)	Twitter	Both, 2009–2020
Barbera et al. (2019)	Twitter	Both, 2013–2014
Cormack (2016a)	Newsletters	Both, 2009–2010
Cormack (2016b)	Newsletters	Both, 2009–2010
Cowburn & Saltzer (2024)	Twitter	House, 2020
Davis & Russell (2024)	Twitter	Senate, 2013–2023
Evans & Clark (2016)	Twitter	House, 2012
Evans et al. (2014)	Twitter	House, 2012
Fowler et al. (2021)	Facebook Ad Library (API)	Both, 2018
Fu & Howell (2020)	Twitter and Facebook	Both, 2018
Greene (2024)	Facebook	Both, 2016–2022
Green et al. (2024)	Twitter, Facebook, newsletters, press releases, and floor speeches	House, 2019–2021
Hemphill et al. (2020)	Twitter	Both, 2017–2019
Heseltine (2023)	Twitter and Facebook	Both, 2011–2022
Heseltine (2024)	Twitter and Facebook	Both, 2020–2022
Hunt & Miler (2025)	Newsletters	House, 2009–2020
Kaslovsky & Kistner (2024)	Twitter	House, 2019–2022
LaPlant et al. (2023)	Twitter	House, 2020
Macdonald et al. (2023)	Facebook	House, 2019–2020
Macdonald et al. (2024)	Facebook	House, 2019–2020
Macdonald et al. (2025)	Twitter	House, 2019–2020
McKee et al. (2021)	Twitter	Both, 2019
Russell (2018a)	Twitter	Senate, 2013, 2015
Russell (2018b)	Twitter	Senate, 2013, 2015
Russell (2021a)	Twitter	Senate, 2015
Smith & Russell (2022)	Twitter	Both, 2017–2019
Straus et al. (2016)	Twitter	Senate, 2014
Tillery (2018)	Twitter	Both, 2013–2014
Warner (2023)	Twitter	Both, 2013–2014
Yu et al. (2024)	Twitter	Both, 2016–2020

B Training Materials Provided to Research Assistants

Classifying Congressional Tweets

Training Materials

May 2023

Summary of the Research Project

Studying Communication by Congress

A long-studied question: How do members of Congress communicate policy positions and craft their personal brand?

- E.g., Mayhew (1974), Fenno (1978), Grimmer et. al. (2015)

With a new twist…

- How do MCs communicate in the age of *social media*?
- Do they focus on credit claiming? Personal advertising? Partisan attacks?
- Under what conditions do constituents listen?

Members of Congress on Twitter

Our platform of interest: Twitter

- Usage by MCs is near universal
- Typical member has ~40k followers
- Journalists, politicos, etc. amplify tweets further

Congressional communication on Twitter has been studied before (e.g., Russell 2021), but never before at scale (broad categorizations, both House and Senate, multiple sessions)

The Strategy

Last summer, David Hilden, Jamie Wright, Beth Simas and I classified 2,000 tweets from the 117th Senate.

- The *Center for Effective Lawmaking* has given us a grant to expand classification to include both House and Senate for the 115th, 116th, and 117th Congresses

Can’t manually classify each and every tweet, however.

How to Categorize the Tweets

Classifying Congressional Tweets

Our goal: Classify Congressional tweets into one of 7 broad, previously established categorizations

- Based on Mayhew (1974) and Russell (2021)

The categories: Position Taking, Advertising, Policy Credit Claiming, Constituent Credit Claiming, Negative Partisan Posturing, Positive Partisan Posturing, Bipartisan Posturing

- Categories are neither *mutually exclusive* nor *collectively exhaustive*

Category 1. Position Taking

Definition of Position Taking: Any tweet mentioning a policy area, a specific piece of legislation, or making a statement on a public issue

- Must involve the member taking a position of some sort, either for or against.
- Has to be more specific than just endorsing a commonly-shared value. Expressing support for “governmental effectiveness” or “honesty” is not taking a position on an area of conflict.

Example of a Position Taking Tweet

```
[1] "The Biden Administration believes if they dont talk about the border
crisis it will go away.  Its not going away its only getting worse."
```

Another Example

```
[1] "Today in a Senate Energy and Natural Resources Cmte hearing I questioned
the Director of the @BOEMDOI Lefton to explain the administrations decision to
update the so-called Arctic Rule completely disregards the latest science
regarding offshore resource development.
https//twitter.com/lisamurkowski/status/1393245146200817666/video/1"
```

Category 2. Advertising

Definition of Advertising: Any tweet discussing a public-facing event or communication by the member such as a media appearance, a committee hearing, a press release, or a town hall.

- Tweet must highlight the member specifically (e.g., a tweet simply mentioning a committee hearing without focusing on the member's participation/actions would not be advertising)

Example of an Advertising Tweet

[1] "Last night Senator Hawleys staff presented a Congressional Record to the father of Marine Lance Corporal Jared Schmitz. Schmitz died while heroically serving alongside 12 other U.S. service members in Afghanistan this past August. https//twitter.com/SenHawleyPress/status/1501932299021012993/photo/1"

Another Example

```
[1] "Sen. Paul talks debt jobs in visit to Henry Co.
https//www.hclocal.com/content/sen-paul-talks-debt-jobs-visit-henry-co"
```

Category 3. Policy Claiming

Definition of Policy Claiming: Any tweet by the member taking responsibility for any public policy effort (sponsoring a bill, influencing a federal agency, etc.).

- Policy claiming can be either retrospective (“I helped pass X”) *or* prospective (“I will fight for X”), but must highlight the member’s actions.
- Action taken outside of government (e.g., businesses, universities) should NOT be classified as policy claiming.
- Can include action against a policy

Example of a Policy Claiming Tweet

[1] "Last month I urged HSGAC Chair SenGaryPeters to investigate taxpayer-funded royalty payments collected by govt scientists at the NIH. I wont stop fighting for accountability. We cannot let American taxpayers be on the hook for this potential conflict of interest. https//twitter.com/FoxBusiness/status/1534769808612401153"

Another Example

[1] "I look forward to serving as Ranking Member on the Senate Homeland Security &amp Govt Affairs Committee &amp working with @SenGaryPeters to conduct bipartisan oversight of our federal agencies and craft legislation to safeguard Ohio and our country from growing threats."

Category 4. Constituent Claiming

Definition of Constituent Claiming: Any tweet by the member that mentions benefits to the constituents or state

- *Does not* have to reference a specific action (linking policy and benefits to constituents is sufficient), but *does* have to make a constituency-specific benefit clear
- A tweet that mentions meeting with constituents to discuss issues of local concern *is* Constituency Claiming
- Constituency Claiming can not be negative (i.e., not “Bill X will hurt our state”)

Example of a Constituent Claiming Tweet

```
[1] "UPDATE Since the Paycheck Protection Program reopened last month 4666
Maine small businesses have been approved for 371 million in forgivable loans
supporting jobs across our state.
https//www.collins.senate.gov/newsroom/nearly-4700-maine-small-employers-have-
been-approved-371-million-forgivable-loans-pppE28099s"
```

Another Example

```
[1] "Happy Great American Outdoors Day Michigan ??????? ?????   Proud to help
pass the Great American Outdoors Act one year ago so we can protect Michigans
treasured places for generations to come.
https//twitter.com/SleepingBearNPS/status/1422953449013002241"
```

Category 5. Negative Partisan Posturing

Definition of Negative Partisan Posturing: Any tweet by the member that attacks or negatively references the opposing party

- Includes mentions of the opposing party as well as any political actor or policy tied to the party
- The connection to the opposing party must be clear, however

Example of a Negative Partisan Tweet

```
[1] "??Ill be on foxnews with HARRISFAULKNER immediately after JoeBidens lie-
filled speech to talk about Bidens raging inflation crisis &amp why he must
RESIGN to fix it. Tune in
https//twitter.com/SenRickScott/status/1524047996727574529/photo/1"
```

Another Example

[1] "Democrats backwards approach to their reckless tax-and-spending spree is an insult to the American people especially while inflation reaches a 31-year high. American families have to be responsible with their spending and their government should do the same."

Category 6. Positive Partisan Posturing

Definition of Positive Partisan Posturing: Any tweet by the member with a supportive reference to the member's party or a representative of the party

- Requires a positive-valence, i.e., must cast copartisan in a positive light (not just mention)

Example of a Positive Partisan Tweet

[1] "Im glad to see POTUS taking this step. The student debt crisis is preventing millions of working Americans from being able to thrive. Its critical that we CancelStudentDebt. https//www.nbcnews.com/politics/white-house/biden-review-executive-authority-cancel-student-debt-n1262791"

Another Example

```
[1] "Tonight @POTUS laid out his plan to get our supply chains back on track
including through reforms in my bipartisan shipping bill with @SenJohnThune.
Im glad to have his support and will keep working to ensure exporters can get
their goods to market quickly and for a fair price."
```

Category 7: Bipartisan Posturing

Definition of Bipartisan Posturing: Any tweet by the member mentioning compromise or bipartisanship

- Requires an explicit or implicit endorsement of the value of bipartisanship. Must use bipartisanship or compromise in a positive light.

Example of a Bipartisan Tweet

```
[1] "Im glad to have the strong support of so many Democrats Republicans &amp
public figures for my bipartisan effort to build a memorial honoring those who
fought in our nations longest war.  The time is now to get the GWOTMemorial
built on our National Mall.https//www.ernst.senate.gov/public/index.cfm/press-
releases?IDE7D93540-DBEF-45AD-8A5D-845C93FF1A6A"
```

Another Example

[1] "According to research done by @CAPAction the American Rescue Plan has BROAD bipartisan support including 91 of Dems 69 of Independents and 53 of Republicans. Its time for GOP representatives to get on board with their constituents and for Congress to pass @POTUSs plan."

C Full BERTweet Classification Reports

C.1 Twitter

TABLE C.1: Advertising Classifier Classification Metrics, Twitter

	precision	recall	f1-score	support
0	0.93	0.94	0.93	1097
1	0.73	0.71	0.72	276
accuracy			0.89	1373
macro avg	0.83	0.82	0.83	1373
weighted avg	0.89	0.89	0.89	1373

TABLE C.2: Bipartisan Classifier Classification Metrics, Twitter

	precision	recall	f1-score	support
0	1.00	0.99	1.00	1331
1	0.84	0.86	0.85	42
accuracy			0.99	1373
macro avg	0.92	0.93	0.92	1373
weighted avg	0.99	0.99	0.99	1373

TABLE C.3: Constituent Claiming Classifier Classification Metrics, Twitter

	precision	recall	f1-score	support
0	0.97	0.97	0.97	1277
1	0.61	0.53	0.57	96
accuracy			0.94	1373
macro avg	0.79	0.75	0.77	1373
weighted avg	0.94	0.94	0.94	1373

TABLE C.4: Negative Partisan Classifier Classification Metrics, Twitter

	precision	recall	f1-score	support
0	0.97	0.97	0.97	1122
1	0.85	0.85	0.85	251
accuracy			0.95	1373
macro avg	0.91	0.91	0.91	1373
weighted avg	0.95	0.95	0.95	1373

TABLE C.5: Policy Claiming Classifier Classification Metrics, Twitter

	precision	recall	f1-score	support
0	0.97	0.96	0.97	1238
1	0.68	0.76	0.72	135
accuracy			0.94	1373
macro avg	0.83	0.86	0.84	1373
weighted avg	0.94	0.94	0.94	1373

TABLE C.6: Position taking Classifier Classification Metrics, Twitter

	precision	recall	f1-score	support
0	0.89	0.84	0.86	653
1	0.86	0.91	0.88	720
accuracy			0.87	1373
macro avg	0.88	0.87	0.87	1373
weighted avg	0.88	0.87	0.87	1373

C.2 Facebook

TABLE C.7: Advertising Classifier Classification Metrics, Facebook

	precision	recall	f1-score	support
0	0.91	0.93	0.92	2016
1	0.64	0.59	0.62	445
accuracy			0.87	2461
macro avg	0.78	0.76	0.77	2461
weighted avg	0.86	0.87	0.86	2461

TABLE C.8: Bipartisan Classifier Classification Metrics, Facebook

	precision	recall	f1-score	support
0	0.99	1.00	0.99	2387
1	0.85	0.81	0.83	74
accuracy			0.99	2461
macro avg	0.92	0.90	0.91	2461
weighted avg	0.99	0.99	0.99	2461

TABLE C.9: Constituent Claiming Classification Metrics, Facebook

	precision	recall	f1-score	support
0	0.93	0.97	0.95	2387
1	0.66	0.40	0.50	74
accuracy			0.91	2461
macro avg	0.79	0.69	0.72	2461
weighted avg	0.90	0.91	0.90	2461

TABLE C.10: Negative Partisan Classification Metrics, Facebook

	precision	recall	f1-score	support
0	0.99	0.95	0.97	2231
1	0.65	0.88	0.75	230
accuracy			0.95	2461
macro avg	0.82	0.92	0.86	2461
weighted avg	0.96	0.95	0.95	2461

TABLE C.11: Policy Claiming Classification Metrics, Facebook

	precision	recall	f1-score	support
0	0.97	0.96	0.97	2203
1	0.69	0.74	0.72	258
accuracy			0.94	2461
macro avg	0.83	0.85	0.84	2461
weighted avg	0.94	0.94	0.94	2461

TABLE C.12: Position taking Classification Metrics, Facebook

	precision	recall	f1-score	support
0	0.90	0.82	0.86	1299
1	0.82	0.90	0.85	1162
accuracy			0.86	2461
macro avg	0.86	0.86	0.86	2461
weighted avg	0.86	0.86	0.86	2461

C.3 Newsletters

TABLE C.13: Advertising Classifier Classification Metrics, Newsletter

	precision	recall	f1-score	support
0	0.95	0.91	0.93	2241
1	0.42	0.58	0.49	259
accuracy			0.87	2500
macro avg	0.69	0.75	0.71	2500
weighted avg	0.90	0.87	0.88	2500

TABLE C.14: Bipartisan Classifier Classification Metrics, Newsletter

	precision	recall	f1-score	support
0	0.99	1.00	0.99	2410
1	0.87	0.77	0.82	90
accuracy			0.99	2500
macro avg	0.93	0.88	0.91	2500
weighted avg	0.99	0.99	0.99	2500

TABLE C.15: Constituent Claiming Classifier Classification Metrics, Newsletter

	precision	recall	f1-score	support
0	0.82	0.96	0.88	1915
1	0.70	0.29	0.40	585
accuracy			0.80	2500
macro avg	0.76	0.62	0.64	2500
weighted avg	0.79	0.80	0.77	2500

TABLE C.16: Negative Partisan Classifier Classification Metrics, Newsletter

	precision	recall	f1-score	support
0	0.99	0.96	0.98	2369
1	0.55	0.87	0.67	131
accuracy			0.96	2500
macro avg	0.77	0.92	0.82	2500
weighted avg	0.97	0.96	0.96	2500

TABLE C.17: Policy Claiming Classifier Classification Metrics, Newsletter

	precision	recall	f1-score	support
0	0.91	0.97	0.94	1966
1	0.85	0.65	0.74	534
accuracy			0.90	2500
macro avg	0.88	0.81	0.84	2500
weighted avg	0.90	0.90	0.90	2500

TABLE C.18: Position taking Classifier Classification Metrics, Newsletter

	precision	recall	f1-score	support
0	0.92	0.65	0.76	1455
1	0.66	0.92	0.77	1045
accuracy			0.77	2500
macro avg	0.79	0.79	0.77	2500
weighted avg	0.81	0.77	0.77	2500

D Comparing Communication Styles to Legislative Styles

FIGURE D.1: Differences in Member Rhetoric by Legislative Style

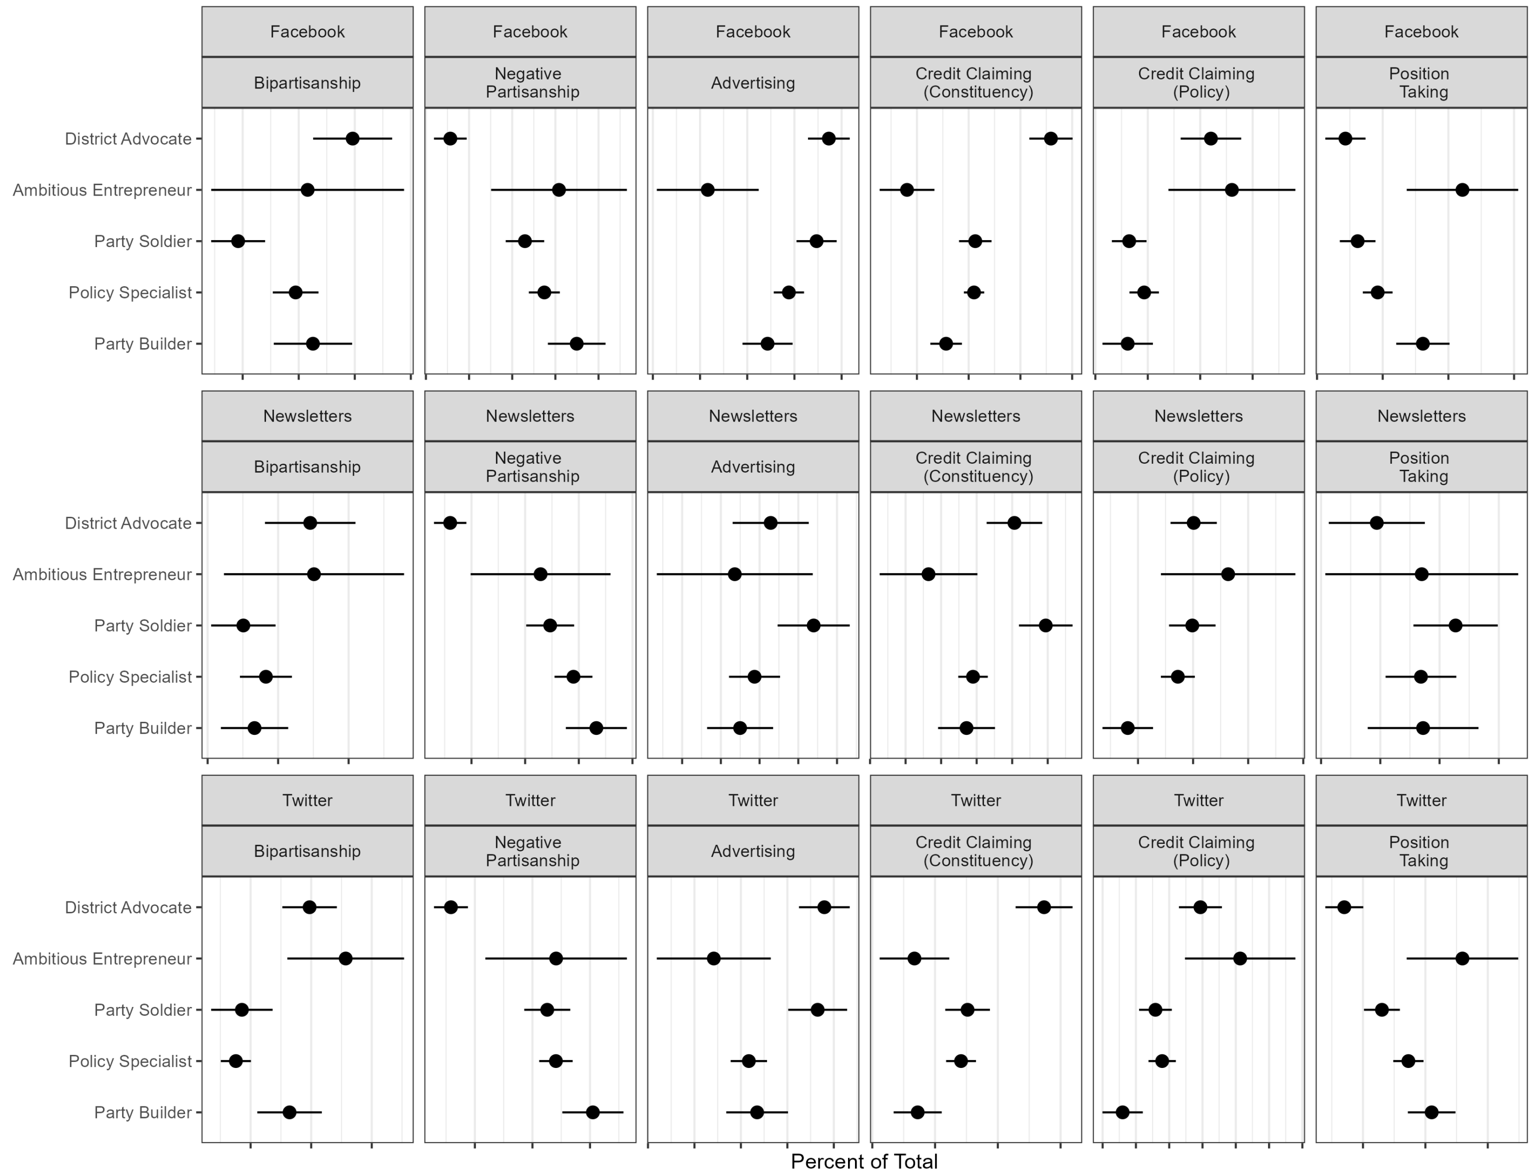

Note: The figure displays the average percent of a member's communications (tweets, Facebook posts, and newsletter bigrams) classified as each of the six categories for the five legislative styles as measured by Bernhard and Sulkin 2018). Error bars display the 95% confidence interval for each point estimate.

E Comparing Alternative Scaling Approaches

Our preferred scaling procedure, for ease of interpretation, performance, and computational efficiency, is the class affinity scaling model (Perry and Benoit 2017). However, to compare the performance of our chosen scaling method with others, and determine how similar resulting partisan scalings would be to alternatives, in this supplemental materials section, we present correlation results with alternative scaling methods.

The first alternative method we consider is the *Wordscore* text scaling method. Wordscores (Laver, Benoit, and Garry 2003) is a supervised text scaling method that positions political texts by comparing word frequencies in unsmeasured texts to those in reference texts with known positions. The technique calculates scores for individual words based on their relative frequencies in the reference texts, then uses these word scores to estimate positions of unscored documents by taking the weighted average of the scores of words they contain.

The second alternative method we consider is using a supervised machine learning model that produces predicted probabilities of class membership (i.e., Republican versus Democratic affiliation) and using those probabilities as a scaled measure of partisanship. This approach has been used by previous researchers who have scaled congressional ideology based on tweets and other messages (Cowburn and Sältzer 2024; Green et al. 2024, 2020). In these examples, text has been scaled at the level of an individual member of Congress. More technically, rather than construct a document feature matrix with a row for each tweet, post, etc., these researchers construct a document feature matrix with a row for each member of Congress, and each cell aggregates the total usage of particular words or features for all messages authored by the member. This improves computational efficiency considerably, and produces higher correlations with other member-level ideology measures (e.g., DW-NOMINATE), but does not allow one to scale individual messages as we do.

For this reason, we cannot use the multinomial inverse regression approach (Green et al. 2020, 2024), as it is computationally infeasible to apply to a dataset with millions of individual texts. We can, however, use the Naive Bayes predicted probability method of Cowburn and Salzer 2024), which we compare here.[26]

We thus replicate our scaling procedure using the Wordscore and Naive Bayes predicted probability approaches with our tweet data. First, Figure E.1 shows the correlation between the partisanship estimates these two methods produce and those produced by our Class Affinity method. The Naive Bayes classifier is extremely highly correlated with our method, with a Pearson's r of 0.99. The Wordscore scalings are also highly correlated with the outputs of both alternative methods, but the correlation is not quite as strong (r = 0.74 with our class affinity approach). In short, it appears that the estimated partisanship of texts does not vary much based on which scaling approach is utilized.

To evaluate whether either of these alternative scaling procedures is more correlated with other member-level ideology measures, we next take the average for individual members and compare to the member's DW-NOMINATE score (1st dimension) and their CF score, as described in the main text. The resulting within-party correlations are displayed in Figure E.2. Unsurprisingly given how strongly correlated the Naive Bayes and Class Affinity results are, both approaches perform approximately as well at predicted members' DW-NOMINATE and CF-scores. The Wordscore approach, in contrast, does a better job predicting member-level ideology in some cases (DW-NOMINATE for Democrats), a worse job in other cases (CF scores for Democrats), and performs roughly the same in still other cases (both CF scores and DW-NOMINATE for Republicans). We thus feel confident that, for the entirety of the data, the scaling method we use is at least as good as feasible alternatives. For some specific purposes, however, researchers may wish

[26]Another method we do not consider here is Wordfish, a scaling method developed by Slapin and Proksch (2008). Wordfish is a statistical model for analyzing political text that estimates policy positions of political actors (like parties or politicians) based on word frequencies in documents. Unlike Wordscores, Wordfish is an unsupervised scaling method that doesn't require reference texts or human coding that simultaneously estimates both document positions and word weights. Unfortunately, the computational time of Wordfish scales exponentially with corpus size, meaning that it too is computationally infeasible for our purposes.

FIGURE E.1: Correlation of Estimated Partisanship for Individual Messages Across Scaling Procedures

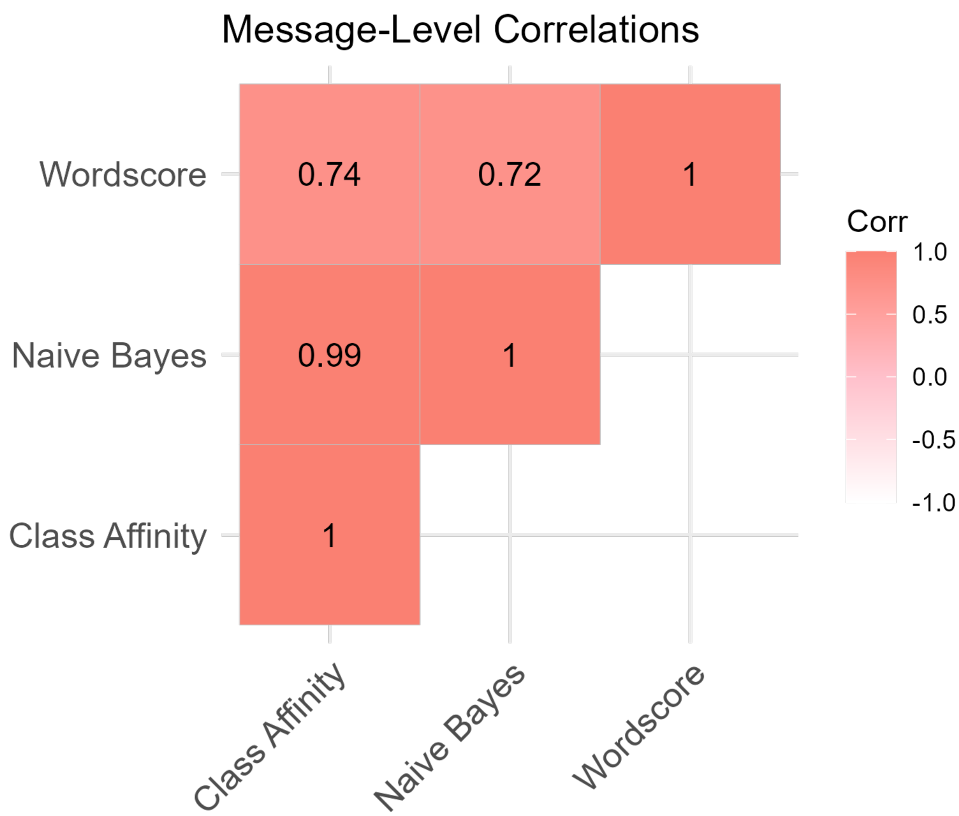

to consider using Wordscore as an alternative scaling approach.

FIGURE E.2: Correlation of Estimated Partisanship and Ideology, Aggregated to the Member-Level

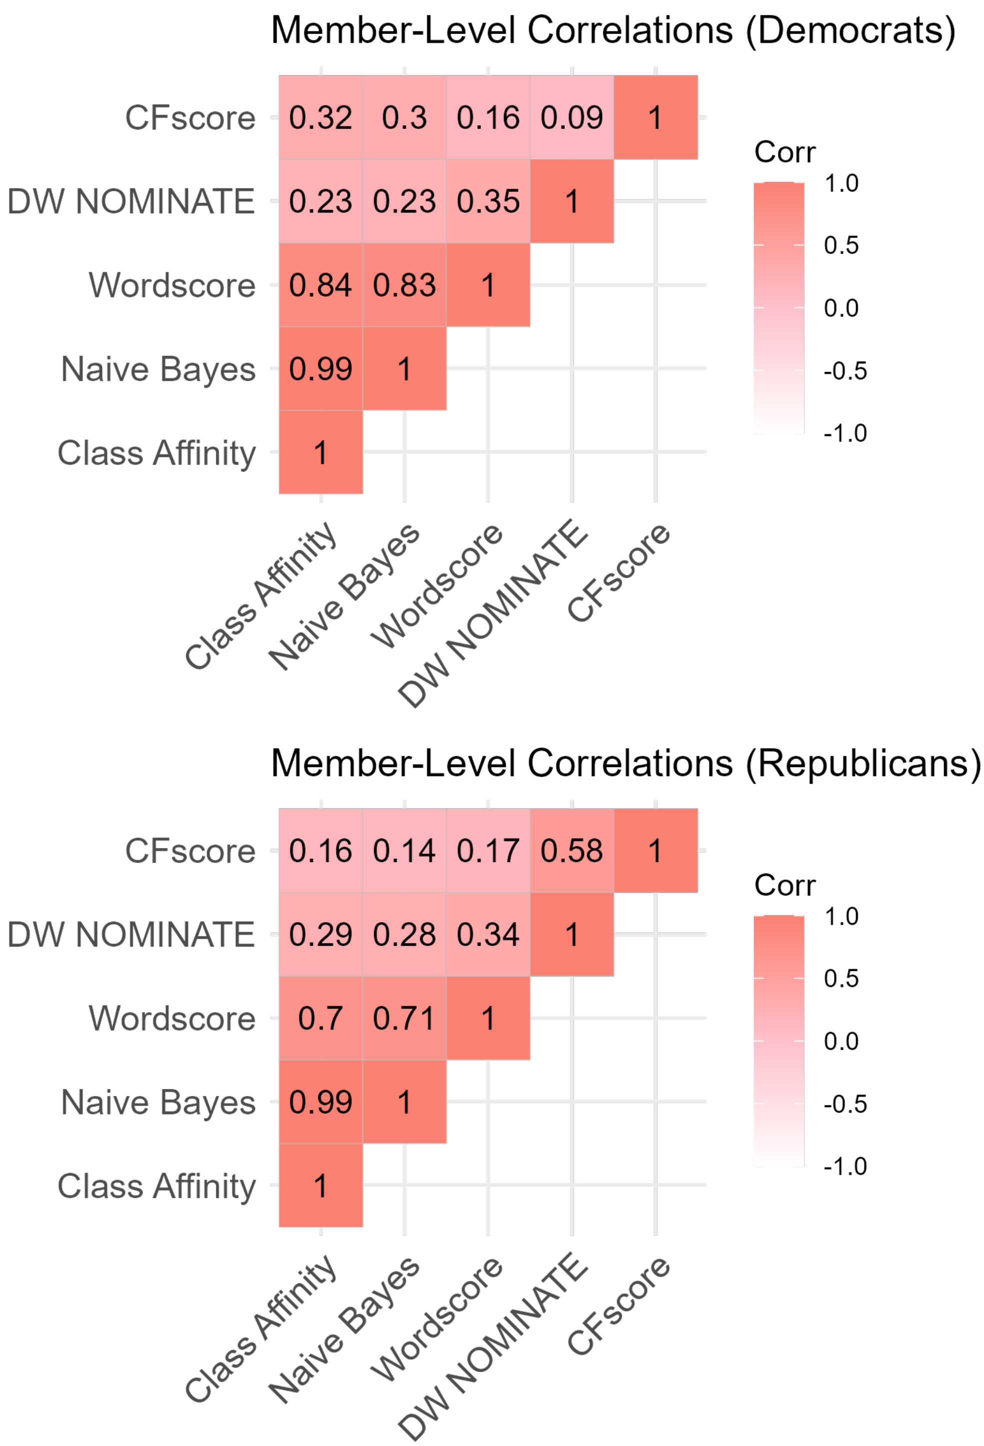

F Keyword Plot for Scaling

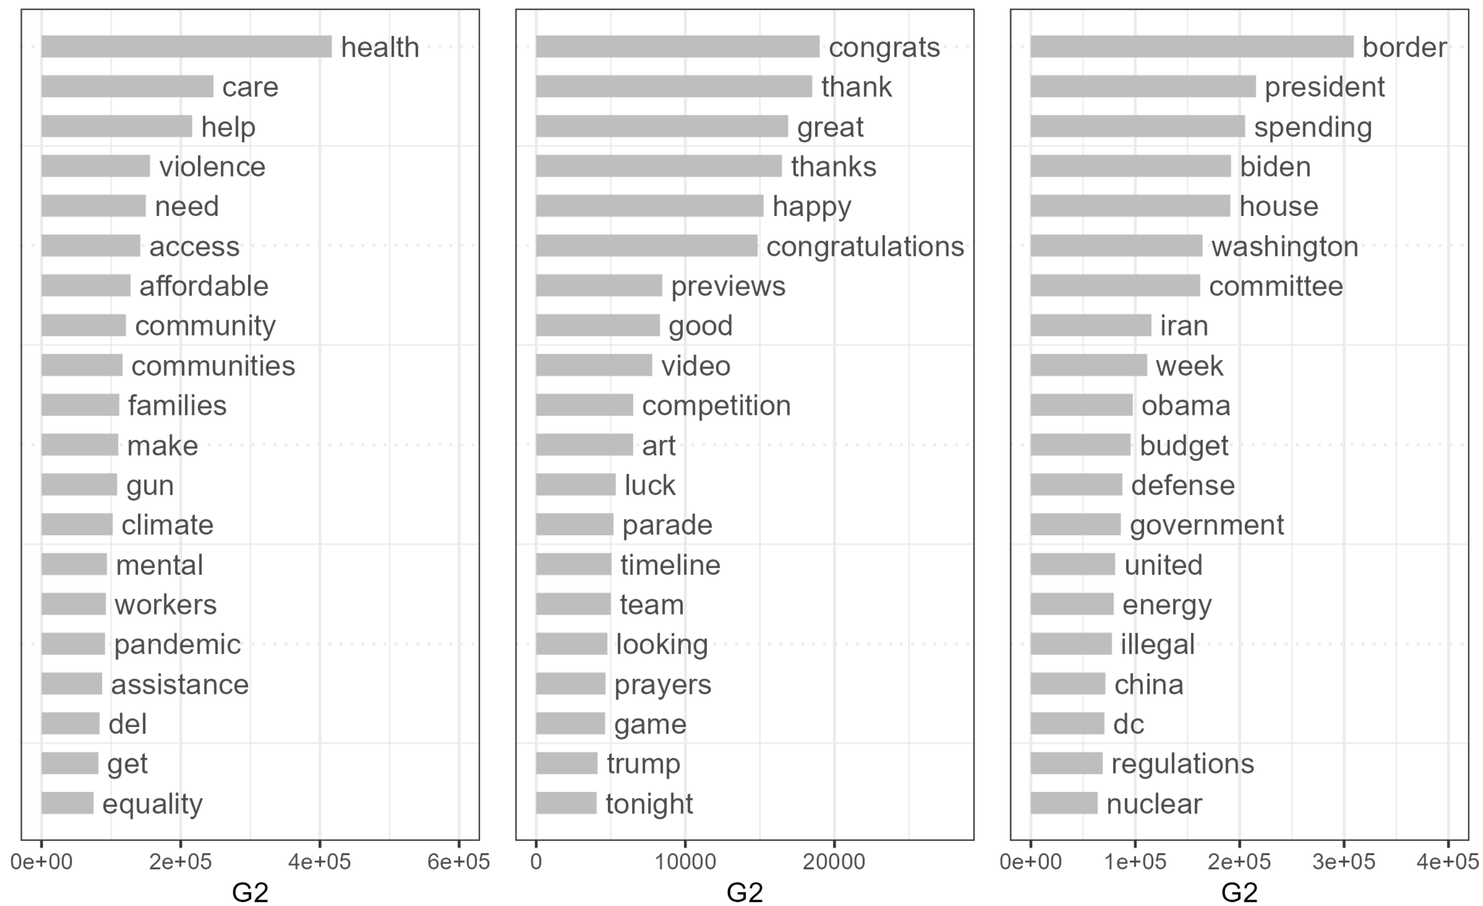

FIGURE F.1: MOST PREDICTIVE WORDS. TEXT PARTISANSHIP SCALE

Note: Frequency of key terms for scaling at various points. The left panel shows Democratic rhetoric (a Text Partisanship Score of -1 to -1/3), the middle panel shows mixed rhetoric (a Text Partisanship Score of -1/3 to +1/3), and the right panel shows Republican rhetoric (a Text Partisanship Score of +1/3 to 1).

G Table of By-Session Partisanship Regression Estimates

TABLE G.1: FULL TABLE OF RESULTS FOR FIGURE 5

	DV: Partisan Extremity Score					
	Twitter		Facebook		Newsletter Bigrams	
Session111	0.552**	0.665**	0.569**	0.689**	0.585**	0.701**
	(0.010)	(0.006)	(0.007)	(0.008)	(0.012)	(0.020)
Session112	0.565**	0.680**	0.588**	0.723**	0.577**	0.722**
	(0.006)	(0.006)	(0.006)	(0.006)	(0.011)	(0.023)
Session113	0.590**	0.652**	0.615**	0.714**	0.606**	0.730**
	(0.005)	(0.006)	(0.005)	(0.006)	(0.013)	(0.005)
Session114	0.607**	0.624**	0.631**	0.683**	0.602**	0.719**
	(0.005)	(0.005)	(0.006)	(0.008)	(0.018)	(0.007)
Session115	0.642**	0.618**	0.662**	0.659**	0.598**	0.687**
	(0.005)	(0.006)	(0.006)	(0.009)	(0.017)	(0.006)
Session116	0.689**	0.586**	0.716**	0.612**	0.709**	0.587**
	(0.005)	(0.005)	(0.004)	(0.007)	(0.012)	(0.010)
Session117	0.720**	0.657**	0.743**	0.688**	0.741**	0.694**
	(0.005)	(0.005)	(0.004)	(0.005)	(0.011)	(0.007)
Num.Obs.	2,781,449	2,146,893	1,330,239	1,202,403	1,496,074	2,915,905
R2 Adj.	0.040	0.015	0.036	0.017	0.049	0.028
Party	D	R	D	R	D	R

Table displays coefficient from OLS models.
Standard errors clustered by member shown in parentheses. *p<0.05; **p<0.01

H Across Time Trends for Categories Not Shown in Main Text

FIGURE H.1: CHANGE IN CATEGORY FREQUENCY FOR CATEGORIES NOT SHOWN IN MAIN TEXT

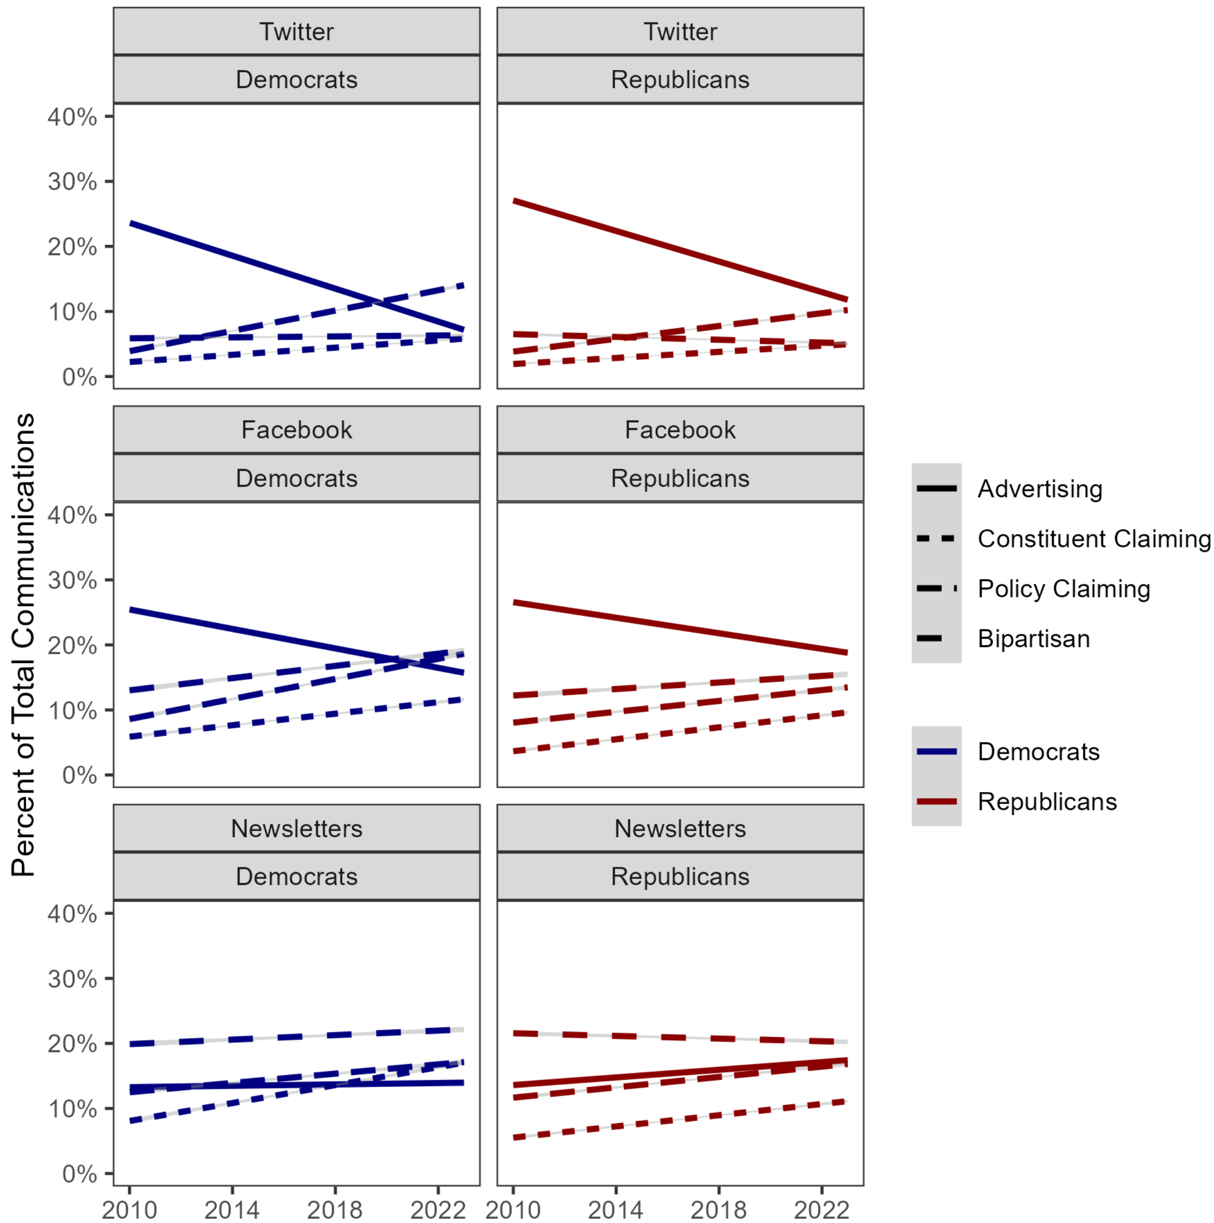

I Table of Engagement Analysis Regression Estimates

TABLE I.1: FULL TABLE OF ESTIMATES FOR FIGURE 7

	Dependent Variable			
	log(Likes + 1)	log(Retweets + 1)	log(Likes + 1)	log(Shares + 1)
Advertising	-0.216**	-0.219**	-0.174**	-0.367**
	(0.015)	(0.014)	(0.013)	(0.015)
Credit Claiming (Constituency)	-0.426**	-0.261**	-0.391**	-0.204**
	(0.015)	(0.014)	(0.011)	(0.013)
Credit Claiming (Policy)	-0.083**	-0.128**	0.022**	-0.103**
	(0.013)	(0.012)	(0.009)	(0.009)
Position Taking	0.148**	0.320**	0.094**	0.271**
	(0.014)	(0.014)	(0.012)	(0.014)
Negative Partisanship	0.621**	0.740**	0.269**	0.755**
	(0.035)	(0.036)	(0.016)	(0.019)
Bipartisanship	-0.040**	-0.040**	0.136**	0.047**
	(0.011)	(0.010)	(0.009)	(0.009)
Text Extremity Score	-0.005	0.074**	0.031	0.056**
	(0.024)	(0.023)	(0.021)	(0.016)
Platform	Twitter	Twitter	Facebook	Facebook
Member FEs	Y	Y	Y	Y
Session FEs	Y	Y	Y	Y
Num.Obs.	4,928,342	4,928,342	2,532,552	2,532,552
R2 Adj.	0.650	0.530	0.499	0.461

Table displays coefficient from OLS models. Standard errors clustered by member shown in parentheses. *p<0.05; **p<0.01

J CES Approval Analyses

As discussed in the main text, to evaluate whether messaging styles in aggregate are associated with higher or lower approval by constituents, we take advantage of constituency approval ratings contained in the Cooperative Election Study (CES) surveys.[27] The CES samples respondents from across the United States but identifies each respondent's Congressional district. The surveys ask each respondent how much they approve of their representative and two senators on a four-point scale ranging from Strongly Approve to Strongly Disapprove.

For each member in each legislative session from the 111th session of Congress onwards, we evaluate whether differences in communication style are associated with higher or lower approval ratings. Specifically, we use a series of multivariate OLS regressions to examine whether the number of tweets or Facebook posts (logged to address right-skew), the average Text Partisan Extremity Score, or the percent of messages in each of the six categories for each of the two communication forms predict how approving constituents are of their representatives. A variety of control variables and fixed effects are included to account for other factors that may affect both approval rating and member communication.[28] We control for the partisanship of a member's district or state (measured using presidential vote share), the member's seniority within the chamber, whether the member is a party leader, whether they are a committee chair, whether the member is a woman, whether the member is black, and whether the member is Hispanic. We include party-session fixed effects, which address any across-time fluctuations in approval for Democrats and Republicans, and also account for differences in approval between members in the majority versus the minority. Finally, we include state fixed effects to account for differences in average approval rating by state. Standard errors are clustered by individual members to account for non-independence of errors when the same members are in the dataset for multiple

[27]Date are publicly available at www.cces.gov.harvard.edu/

[28]All control variables come from updates to Volden and Wiseman (2014; 2018), and can be downloaded at https://thelawmakers.org/data-download.

sessions.

FIGURE J.1: Communication Style and Constituent Approval

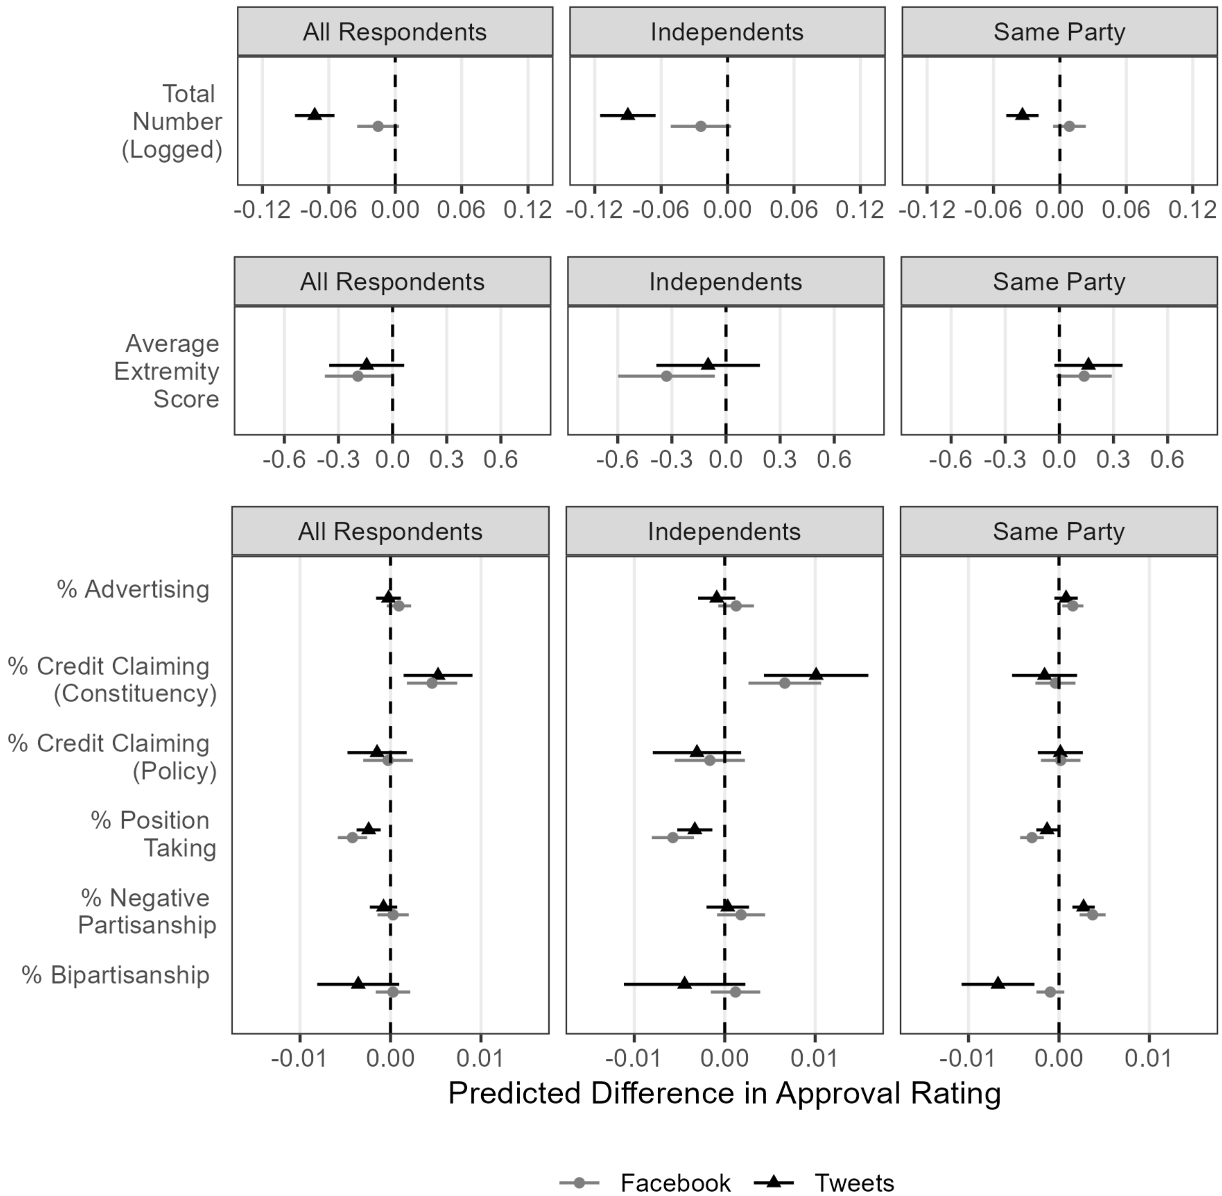

Note: The figure displays OLS coefficient estimates and 95% confidence intervals from the full model results shown in Table J.1. The dependent variable is the average CES approval rating (for all respondents, independent respondents, and respondents of the same party as the member, respectively). All models include control variables and fixed effects, with standard errors clustered by member.

The results of interest are displayed in Figure J.1. Separate models are estimated using a member's average approval rating among all respondents (mean = 2.9, SD = 0.3), independents (mean = 2.7, SD = 0.4) and same party respondents (mean = 3.3, SD = 0.2), as members may particularly care about the opinions of these latter two groups, given their importance in the general election and primary election respectively.

Figure J.1 reveals several key points. First, members who tweet more have lower approval ratings among their constituents, on average, a finding that is consistent regardless of whether one considers all constituents, just independents, or just same-party constituents. Members who use Facebook more also appear to have lower approval ratings, although these results are only marginally significant.

Second, members who use more extreme rhetoric on Twitter and Facebook seem to receive lower approval ratings from constituents writ large and independents in particular, although the estimate is only statistically significant for the latter group on Facebook. Among same-party constituents, however, there is some evidence of a positive relationship, although the estimates are not statistically significant here either.

Finally, if we consider how often members send messages of the various representational and partisanship categories we measure, we find that using more credit claiming messages is associated with higher approval ratings. Position taking is consistently associated with lower approval ratings. Negative partisanship is positively associated with higher approval ratings among constituents of the same party, but no other constituent groups.

TABLE J.1: Full Table of Estimates for Figure J.1

	DV: Average Approval Rating					
	All Respondents		Independents		Same Party	
	(1)	(2)	(3)	(4)	(5)	(6)
Total Number	-0.066**	-0.007	-0.077**	-0.011	-0.028**	0.016
	(0.009)	(0.010)	(0.014)	(0.015)	(0.008)	(0.008)
Average Extremity Score	-0.172	-0.343**	-0.060	-0.448**	0.106	0.023
	(0.109)	(0.097)	(0.150)	(0.139)	(0.094)	(0.089)
Pct. Advertising	-0.001	0.000	-0.001	-0.001	0.000	0.001*
	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)
Pct. Credit Claiming (Constituency)	0.002	0.005**	0.001	0.006**	-0.001	0.001
	(0.001)	(0.001)	(0.002)	(0.002)	(0.001)	(0.001)
Pct. Credit Claiming (Policy)	0.006**	0.008**	0.011**	0.010**	0.004*	0.006**
	(0.002)	(0.002)	(0.003)	(0.002)	(0.002)	(0.001)
Pct. Position Taking	-0.005**	-0.005**	-0.009**	-0.008**	-0.003**	-0.004**
	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)
Pct. Negative Partisanship	0.001	0.001	0.002	0.003*	0.004**	0.005**
	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)
Pct. Bipartisanship	-0.010*	-0.015**	-0.013*	-0.018**	-0.016**	-0.013**
	(0.004)	(0.002)	(0.005)	(0.003)	(0.004)	(0.002)
District Dem. Pres. Voteshare	0.100	0.019	0.219	0.107	-0.163*	-0.156*
	(0.093)	(0.090)	(0.132)	(0.131)	(0.075)	(0.077)
District Partisan Favorability	0.991**	1.035**	0.580**	0.665**	0.204**	0.170*
	(0.088)	(0.086)	(0.126)	(0.127)	(0.071)	(0.076)
Woman	-0.012	-0.022	-0.043	-0.043	0.029*	0.022
	(0.017)	(0.017)	(0.025)	(0.026)	(0.014)	(0.014)
AfricanAmerican	-0.023	-0.026	0.007	0.002	-0.003	0.006
	(0.027)	(0.025)	(0.042)	(0.041)	(0.021)	(0.021)
Hispanic	-0.039	-0.063*	-0.052	-0.089*	-0.076**	-0.078**
	(0.032)	(0.032)	(0.043)	(0.044)	(0.026)	(0.026)
Party Leader	-0.048	-0.077*	-0.097*	-0.123**	0.002	-0.023
	(0.031)	(0.033)	(0.043)	(0.046)	(0.026)	(0.028)
Committee Chair	-0.113**	-0.134**	-0.155**	-0.174**	-0.060*	-0.073**
	(0.024)	(0.025)	(0.034)	(0.037)	(0.023)	(0.025)
Data	Twitter	Facebook	Twitter	Facebook	Twitter	Facebook
Party-Session FEs	Y	Y	Y	Y	Y	Y
State FEs	Y	Y	Y	Y	Y	Y
Seniority FEs	Y	Y	Y	Y	Y	Y
Num.Obs.	2432	2245	2431	2245	2432	2245
R2 Adj.	0.477	0.490	0.305	0.306	0.253	0.267

Table displays coefficient from OLS models. Standard errors clustered by member shown in parentheses. *p<0.05; **p<0.01

K Table of Primary Analysis Regression Estimates

TABLE K.1: FULL SET OF REGRESSION ESTIMATES FOR FIGURE 8

Party	Message Content	2010 - 2022	2022	2020	2018	2016	2014	2012	2010
Democrats	Partisan Extremity	-0.009	-0.009	-0.009	-0.006	-0.011	-0.003	-0.015	0.002
		(0.004)	(0.004)	(0.004)	(0.004)	(0.005)	(0.005)	(0.006)	(0.009)
Democrats	Position Taking	0.033	0.033	0.024	0.026	0.023	0.021	0.015	0.030
		(0.008)	(0.008)	(0.008)	(0.008)	(0.009)	(0.009)	(0.011)	(0.020)
Democrats	Policy Claiming	0.010	0.010	0.003	0.002	0.005	0.000	0.006	0.037
		(0.005)	(0.005)	(0.005)	(0.004)	(0.007)	(0.005)	(0.006)	(0.012)
Democrats	Constituent Claiming	-0.005	-0.005	-0.004	-0.002	0.010	0.000	-0.001	0.028
		(0.006)	(0.006)	(0.004)	(0.007)	(0.008)	(0.007)	(0.009)	(0.013)
Democrats	Advertising	-0.004	-0.004	0.006	-0.002	0.013	0.001	0.000	0.016
		(0.005)	(0.005)	(0.005)	(0.008)	(0.008)	(0.009)	(0.011)	(0.017)
Democrats	Negative Partisan	0.012	0.012	0.021	0.015	0.002	0.007	0.002	0.011
		(0.005)	(0.005)	(0.008)	(0.006)	(0.007)	(0.004)	(0.008)	(0.011)
Democrats	Bipartisan	0.007	0.007	0.002	0.003	0.004	0.001	0.005	-0.002
		(0.003)	(0.003)	(0.002)	(0.002)	(0.005)	(0.003)	(0.003)	(0.006)
Republicans	Partisan Extremity	0.000	0.000	0.009	0.007	-0.002	0.000	0.008	0.006
		(0.005)	(0.005)	(0.006)	(0.005)	(0.006)	(0.005)	(0.006)	(0.010)
Republicans	Position Taking	0.009	0.009	0.018	0.007	-0.013	0.008	0.003	-0.008
		(0.006)	(0.006)	(0.008)	(0.009)	(0.011)	(0.009)	(0.013)	(0.018)
Republicans	Policy Claiming	0.002	0.002	0.001	0.002	-0.007	-0.006	-0.009	-0.002
		(0.004)	(0.004)	(0.007)	(0.005)	(0.007)	(0.006)	(0.006)	(0.010)
Republicans	Constituent Claiming	-0.009	-0.009	-0.005	-0.006	0.009	0.000	-0.003	0.029
		(0.005)	(0.005)	(0.005)	(0.007)	(0.009)	(0.008)	(0.008)	(0.013)
Republicans	Advertising	-0.010	-0.010	0.002	-0.012	0.005	0.019	0.002	0.005
		(0.005)	(0.005)	(0.006)	(0.007)	(0.010)	(0.010)	(0.011)	(0.018)
Republicans	Negative Partisan	-0.018	-0.018	0.004	-0.010	-0.008	-0.006	-0.005	-0.023
		(0.007)	(0.007)	(0.008)	(0.006)	(0.010)	(0.007)	(0.010)	(0.013)
Republicans	Bipartisan	-0.001	-0.001	0.003	0.000	-0.002	0.001	0.000	-0.001
		(0.002)	(0.002)	(0.003)	(0.003)	(0.005)	(0.004)	(0.003)	(0.006)

Table displays coefficient from OLS models. Standard errors clustered by member shown in parentheses. $^{*}p<0.05$; $^{**}p<0.01$