

RESEARCH ARTICLE

Legislative communication and power: measuring leadership in the U.S. House of Representatives from social media data

Daniel Ebanks¹, Hao Yan², R. Michael Alvarez² , Sanmay Das³ and Betsy Sinclair²

¹California Institute of Technology, Pasadena, CA, USA, ²Washington University in St Louis, St Louis, MO, USA and ³George Mason University, Fairfax, VA, USA

Corresponding author: Betsy Sinclair; Email: betsysinclair@gmail.com

(Received 24 June 2023; revised 01 April 2024; accepted 15 April 2024; first published online 22 October 2024)

Abstract

Control over the legislative messaging agenda has important political, electoral and policy consequences. Existing models of congressional agenda-setting suggest that national polarization drives the agenda. At the same time, models of home style and formal models of leadership hypothesize that legislators shift their messaging as they balance coordination and information problems. We say the coordination problem dominates when conditions incentivize legislators to agree on the same message rather than fail to reach consensus. Conversely, the information problem is said to dominate in circumstances where legislators prefer to say nothing at all rather than reach consensus on the wrong political message. Formal theories predict that when coordination problems are pressing, legislative members follow the policy positions of party leaders. When their party's information problem is acute, party members instead rely on the wisdom of the caucus to set the party's agenda. To test these theories, we analyze the Twitter accounts of U.S. House members with a Joint Sentiment Topic model, generating a new understanding of House leadership power. Our analyses reveal complex leader-follower relationships. Party leaders possess the power to substantially affect the propensity of rank-and-file members to discuss topics, especially when the coordination problem dominates; these effects are pronounced even when coordination problems are pressing. That said, when the underlying politics are unclear, rank-and-file members exert influence on the discussion of a topic because the information problem is more acute. At the same time and for these uncertain topics, leadership influence decreases, consistent with theory. We show these results are robust to the underlying dynamics of contemporary political discussion and context, including leading explanations for party leadership power, such as national polarization.

Keywords: Legislative communication; leadership; social media; information; coordination

Introduction

In January 2019, the U.S. Congress was on the brink of crisis and a shutdown. Due to a legislative impasse and political infighting, the legislature could not agree on a compromise to fund the government. Legislative leaders in both parties had to reconcile an uncertain political environment, high policy stakes, and potentially long-lasting electoral consequences. Legislators needed then to balance both their desire to coordinate on a unified message with their desire to actually espouse the right message (with respect to politics, policy, and electoral concerns). The ability of party leaders to set the messaging agenda during this crisis rested on their capacity to balance these concerns. A failure to coordinate on the message or the costs of choosing

the wrong message could have resulted in dire political and electoral consequences for the party, as well as harm to the country from unsound policy. This recent example shows that understanding how the party settles on a message and when members choose to follow their party leaders is crucial for understanding how party leadership functions in a democracy.

Existing theories of Congress suggest party leadership power in modern American parties is best explained by national polarization and increased party cohesion that give rise to top-down party leaders. These existing studies focus on polarization as an explanation for agenda-setting power, and they often compare power relations across parties. In contrast, we study within-party power relations, and analyze how leadership arises within a party. Our analysis considers a formal theoretical framework that considers the predictions from a signaling and coordination model of Congress due to Dewan and Myatt (2007). Extending this theory to how party leaders influence member communications, the model suggests that parties balance tensions between coordination and information problems. Parties would like to coordinate around a unified message, but that is difficult because the underlying political, economic and social conditions of the world are uncertain. Leadership's role in this setting is to help facilitate coordination in the face of this uncertainty.

The recognition of this tension in simultaneously resolving these two problems guides our empirical research. Drawing on this theoretical insight, we develop and test a key hypothesis about party leadership in the contemporary U.S. House of Representatives. We test this hypothesis using social media data and unsupervised learning methods. Testing formal political theory with these data and methods is an important contribution of our research.

We focus on a hypothesis that illuminates this informational problem and connects the party members' need for policy direction with House leaders' willingness to initiate discussion. We show structural stability in the findings across a single presidential term, even when the party in power changes.

These expectations contrast with previous studies of congressional party leadership which are conditioned on ideology and legislative institutions. In fact, we believe our results confound expectations because we are focused on the domain of influence over communication on social media. For example, Aldrich and Rohde (2001) present a theory of conditional government, whereby strong party leaders emerge when parties are internally homogeneous, but are polarized with respect to other parties. As the parties polarize, members delegate more authority to their partisan leaders. This is consistent with theories of strategic party government, whereby parties choose to polarize in order to win elections (Koger and Lebo, 2020). That said, the predictions from Koger and Lebo emphasize that electoral and political incentives for individual politicians may be in tension with the goals of the party overall. The added benefit of employing Dewan and Myatt (2007)'s formal framework is that it explicates and formalizes the underlying mechanisms that drive the divergence between party leaders and members. That is, it is precisely when issues are muddled that members turn to the caucus's preferred stance, and the leaders follow. Under this framework, Dewan and Myatt (2007)'s theory predicts that when issues stances are clear, members are more inclined to follow their leader. At first, this might seem counter-intuitive, but the underlying logic is simple. When there is more uncertainty about which stance to take, parties revert to party consensus and aggregate all of their information (information aggregation). When there is more certainty over the correct stance, party members follow the leader (coordination). The key benefit of implementing tests of predictions of this theory is that we can estimate the tension between the information aggregation motivation relative to the coordination motivation.

Additionally, Aldrich and Rohde (1998) used DW-Nominate scores to quantify how parties have grown more polarized and ideologically homogeneous. Similarly, Gamm and Smith (2020) argue that modern parties are top-down institutions, with party leaders exerting control over legislation and committees, especially in the U.S. House of Representatives. Others have argued that modern congressional leadership is powerful: various authors have noted that leaders are empowered with the capacity to bypass committees (Bendix, 2016; Howard and Owens, 2020), to

directly negotiate policy (Curry, 2015; Wallner, 2013), set the agenda (Harbridge, 2015), and to limit floor debate (Tiefer, 2016). All of these papers focus on policies and agenda items where the parties already have clear stances or are predisposed to be polarized. In this study, we test theories of party leadership and coordination where opinion is neither yet necessarily formed nor polarized. By using a social media data set, we can study how parties respond to an issue that is uncertain and does not divide or unite the parties. That is, no one in the party knows which message around which to coordinate. In this case, we would expect the party caucus to lead if they perceive acute electoral effects, but may defer to the leader if their first priority is to coordinate on a message over “getting it right”.

We note two key distinguishing features of our analysis relative to earlier studies. First, we avoid the selection problems inherent in using roll call data to identify leadership influence. As party leaders are strategic and have agenda power, they control which bills reach the floor. Since they are unlikely to bring bills to the floor that divide their own party, the fact that leadership-supported bills obtain majorities could signal strength within the party (if leaders persuaded the rank-and-file to support a bill close to the leader’s preferred stance), or weakness (if the rank-and-file overrules the leader in the party conference vote). Social media communications are not subject to the same level of leadership control – members of Congress often cultivate their own online home styles. Second, the high-frequency nature of social media data allows us to capture changes in legislative behavior at a much more granular level than roll call data. In particular, social media offers rich data concerning the party leadership’s ability to direct legislative communication and public engagement around specific topics *among* their members, and in real time.

Our paper contributes to the literature in four ways. First, because we define House leadership influence as the ability of leaders to persuade rank-and-file members to adopt communication strategies similar to their own, we can exploit social media data to measure policy positions (Yan et al., 2019). Specifically, we quantify House leadership influence in terms of leaders’ ability to pull rank-and-file public stances on Twitter closer to the leadership’s messaging on those same policy positions. Second, we use high-frequency data that shows that the dynamics of leadership can change daily. This suggests that leaders’ influence over the party’s policy positions varies based on the issues dominating discussion at a particular time. Third, our data let us study the influence of House rank-and-file members on their party leaders. We find that House rank-and-file members exert influence on their leaders’ policy position messaging under certain conditions. Our results demonstrate that polarization alone is not sufficient to explain patterns of party leadership in the House. Finally, we show that Natural Language Processing methods and social media data provide insight into online home styles. Thus our work neatly dovetails with Fenno (2003), as it offers a quantitative approach to understanding how members of Congress communicate with their constituencies and one another.

We argue that understanding the role of communication in shaping institutional structures in the House is central to theoretical understandings of leadership, especially within political parties. In particular, parties balance coordinating around a unified policy position while trying to communicate the best policy position in an uncertain world. We show that political communications data from Twitter illuminates understudied aspects of institutions in the House. Twitter is now a key platform that political leaders use to communicate with their constituents and with other politicians, yielding data on their revealed preferences like roll call votes or newsletters to constituents.¹ We use data from the official Twitter accounts of U.S. House members, collected for the 115th and 116th Congresses, between January 1st, 2017 and January 3rd, 2021. After pre-processing these data, we use weakly supervised machine learning methods to

¹Twitter provides a public forum for members of Congress to interact with each other and the public (Hall and Sinclair, 2018). Past research suggests that congressional Twitter activity is part of a legislator’s strategic public communication plan that researchers can use to study legislative behavior (e.g., Barbera et al. (2019) and Kang et al. (2018)).

show that intra-party variation in our data is associated with observed member behavior, namely House of Representatives messaging mechanisms and the institutional structure within each party's conference. We next discuss the primary hypothesis which guides our analysis, detailing the tension between the coordination and information problems.

Theory of leadership communication and power

Our empirical analysis is framed around theoretical insights from the Dewan and Myatt (2007) signaling and coordination game of party leadership and communication – where leadership facilitates coordination on a position in response to uncertain issues. In the context of this framework, uncertainty could be the political or electoral popularity of taking a position, or uncertainty about the policy outcome of a stance. For example, the government shutdown of 2019 presented uncertainty of all three types: there were reasons to believe the electoral impact of a shutdown could be either strong or mild and reasons to believe a shutdown could either favor or disfavor the Democratic House Caucus. Further, the policy outcome of the shutdown was uncertain, as the stalemate occurred over border wall policy. The correct position for Democratic and Republican House members to communicate publicly and in real time on social media was not immediately clear. The theoretical framework notes that leaders help resolve this tension between the information and coordination problems faced by party leaders and rank-and-file by acting as a coordination device around a position in light of this uncertainty. In the context of the model, party leaders issue a public speech and then party members try to coordinate on a public position in an uncertain state of the world.

To clarify the theory, we return to the 2019 government shutdown debate. House Speaker Pelosi attempted to coordinate her party around a single stance and unite the moderate and progressive wings of her party. The government shut down when President Trump and House Democrats failed to agree on a government funding bill due to disagreements over financing the president's border wall with Mexico. The moderate wing had political incentives to break the impasse by appropriating funds for President Trump's border wall, while Democratic progressives desired a harder line of negotiation. In the meantime, House rank-and-file Democrats were privately discussing their sense of the party's mood around the most politically advantageous messaging strategy as they negotiated with a Republican president to resolve the crisis. These discussions occurred online, in person, and over conference calls. The private signals in this legislative coordination game represent these online and offline discussions.

We explain the terms of our hypothesis in the context of our illustrative example; the precision of the private signals represents the variation over the moderate and progressive's internal discussions related to the messaging surrounding the border wall and government funding negotiations. As these signals are private, we do not measure this quantity directly. In the model, the party selects one position whose number of supporters exceeds a threshold. In our example, this is Speaker Pelosi's sense of the level of party support she needs to pursue a particular messaging strategy. In the case where neither position has sufficient support, the party fails to coordinate. In the government funding example, Speaker Pelosi initially struck a hardline messaging strategy, and her members followed her lead. She gauged internal support as sufficiently high for this strategy. This illustrates the concept of the *need of direction*. This concept represents the responsiveness of the messaging strategy to the fundamental political environment, and the gravity of choosing incorrectly. In our illustrative example, the *need for direction* is high, as failure to coordinate could result in prolonged national suffering and a calamitous electoral performance for the party assigned blame for the shutdown by the public.

To conclude the 2019 government shutdown example, some Democratic members publicly indicated they did not support the strategy pursued by their congressional leaders during the crisis, and feared political backlash for little electoral gain. We have no reason to believe that they

privately supported this strategy, as they actively advocated for countervailing messaging on social media. Nor is it likely that Democratic legislators adopted their leadership's messaging strategy if they thought it was doomed politically. Thus, the public signals reflected internal dissent and internal support for Speaker Pelosi's and her leadership team's proposed messaging strategy regarding the shutdown. This ultimately resulted in Speaker Pelosi making concessions to ideologically diverse factions within her party to ensure they coordinated around her stance on a critical issue. Ultimately, President Trump relented after 35 days and the House and Senate passed a funding bill by voice vote.

In our setting, the public position for each party member is communicated on Twitter. To evaluate the ability of the party to coordinate around the leaders' preferred messages, we construct a measure for the concept of *need for direction* that is discussed in detail in Sections 3 and 4.² Specifically, need for direction captures the importance of the party coordinating around the "correct" position. The importance is determined by individual electoral effects, party-level political effects, and national-level policy effects. By picking the correct position, we mean the position that best improves these effects as opposed to doing nothing or picking a different position altogether. When need for direction is high, the information problem tends to dominate. This is because the merits of the position are especially responsive to underlying fundamentals which are uncertain. In this case, we expect members to lead discussion. On the other hand, when need for direction is low, the coordination problem dominates, and we expect leaders to initiate discussion.

We now discuss the exact nature of need for direction in detail. As we reiterate in Table 1, for issues where need for direction is high, we expect House leaders to adopt the communication style of their rank-and-file. Here, the effects of picking the wrong message are outsized, and the party defers to the wisdom of the caucus³. Such topics include those where the policy choice is not immediately clear – for example, both parties have at various times chosen to shut down the government, fund the government, or engage in parliamentary maneuvers. All of these have potentially outsized effects electorally and in terms of policy. For example, a government shutdown could be seen as the party holding fast to their principles, and they will be rewarded by their base. Or, the voters may view it as the latest illustration of government dysfunction, and punish the party that shuts the government down. Here, we expect House leadership influence to be weaker, as the theory suggests that rank-and-file members will hedge against the potential for their leaders to choose an incorrect message, as the consequences for coordinating on the "wrong" message are high.

On issues where the party's need for direction is low, we expect House rank-and-file to adopt the positions of their leaders. Here, the stakes for choosing the wrong position are relatively low, and members prefer to coordinate around a unified policy – even if it is "incorrect" – rather than fail to coordinate at all. We define issues with low need for direction as those that strongly explain the variation in the propensity to discuss sentiment topics, such as the construction of a border wall – which Democrats generally oppose and Republicans generally favor. The "correct" stance on this type of issue for each party is clear. There is little *outsized* electoral payoff or cost in taking these stances. For example, Democrats are going to support abortion rights, oppose cutting welfare programs, and support taxes on high-income brackets. There is little additional cost or payoff for Democrats in tweaking their message. Voters have preconceived notions about the fundamental beliefs of the Democratic party, so the party would rather coordinate on a particular message than worry about crafting the perfect communication on an issue such as abortion or the size of the welfare state.

²Readers interested in details of the theory can refer to Dewan and Myatt (2007).

³The precise mechanism is a "wisdom of the crowd" where the members of the party aggregate all their information to arrive at the best possible message.

Table 1. Terminology

Concept	Revealed By	Empirical Analog
Need for Direction	Sentiment topics with outsized benefit or cost of coordinating	Classify top twenty topics for each party driving separation in sentiment-topic space as uncovered by PCA analysis as high in needing direction
Coordination Problem	Low Need for Direction	Leaders initiate discussion via IRFs on topics outside top 20 of PCA
Information Problem	High Need for Direction	Members initiate discussion via IRFs on top 20 topics of the PCA
Leadership Influence	Leaders' ability to convince rank-and-file members to follow their topics	Leaders have statistically significant IRFs on rank-and-file members

Table 1 presents the key theoretical concepts and their empirical measures. The first column describes the theoretical concepts as we have defined them in the preceding section, while the second column provides the theoretical meaning of each concept. The third column previews the empirical measures we derive from social media data, which we discuss in Section 3 of the paper. Then in Section 4, we discuss the methods we use to translate theoretical concepts into their empirical analogs, with results in Section 5, and the discussion and conclusion in Section 6.

Data and methodology

Data

In order to study the dynamics of communication, we examine legislators' Twitter posts. Using this high-frequency, individual-level data, we examine whether the House party rank-and-file discuss topics that are similar to their leaders' communications on social media or vice versa. We collect the Twitter handles of 511 representatives from January 3rd, 2017 to January 3rd, 2021, covering exactly the 115th and 116th sessions of Congress. We used the official Twitter handles list collected by C-SPAN⁴, following Barbera *et al.* (2019) who used the New York Times Congress Application Programming Interface to identify a list of handles for Members of Congress.

We do not include election, personal, or private accounts in our dataset. While many members have additional personal or campaign social media presences, in order to have a consistent method to collect Twitter data from members of Congress, we focus on their official Twitter accounts. It is precisely these accounts that best represent strategic interactions around substantive policy positions. Personal and electoral Twitter accounts often focus on non-policy issues, like personal family matters, sporting events or scheduling of specific campaign events (such as local town halls or rallies). We focus our study on social media posts that are most likely to discuss policy. Our dataset includes 738,066 tweets, including only original posts. Table SI 1 in the online appendix shows that on average House members tweeted 727.17 times, with notable inter-party variation. Democratic Party members tweeted on average 894.45 times, while Republican Party members tweeted on average 528.31 times.⁵

Twitter data motivation

A lack of granular data has inhibited empirical study of the role of communication in shaping institutions. To this end, political communications data could illuminate understudied aspects of institutions in Congress. Social media is a source of such data, as Twitter (now X, though we will use Twitter as that was the platform's name when we collected the data used in this paper) has

⁴<https://twitter.com/cspan/lists/members-of-congress/members>

⁵See Figure SI 1 for the overall distribution of tweets by House members for this period.

developed into a key platform that political leaders use to communicate with their constituents and with other politicians. In particular, Congressional leaders use Twitter to communicate with their constituents and their co-partisans in each of the party conferences. In this paper, we exploit congressional social media communications to test hypotheses from a game theoretic model where communication is directly tied to the institutional forms of party leadership.

Twitter data is useful to test these theories because it provides both high-frequency political communications data in the form of tweets and high-frequency data regarding legislative relationships in the form of retweets. The average congressional account tweets about 50 times per day, enabling us to study legislative behavior and organization at a granular level that is recorded in real time. Few previous empirical studies of legislative behavior have been able to use such high-frequency daily data in their studies of legislative organization and congressional institutions.

Twitter provides a platform for members of Congress to interact with other legislators and to send public signals (Hall and Sinclair, 2018). We argue that past research suggests that congressional Twitter activity is not merely babble, and instead that communications on Twitter is part of a strategic public communication plan by legislators that researchers can directly observe. Thus, researchers are able to discern patterns of debate and discussion among members of Congress (e.g., Barbera et al. (2019) and Kang et al. (2018)).

In turning to Twitter as a source of data on legislative behavior, we build on previous empirical work employing Twitter data, which has uncovered meaningful structure from Congressional accounts on the social media platform. For example, (Hemphill, Otterbacher and Shapiro, 2013) investigated the Twitter usage by members of Congress in late 2012. They also demonstrated that members of Congress used Twitter for self-advertising. In a related study, Vaccari and Nielsen (2013) looked into the factors that were associated with the popularity of politicians in social networks. They showed that open-seat race candidates were more popular. They also found that campaign funding and popularity in opinion polls have no positive correlation with a politician's popularity in social networks. Finally, Peng et al. (2016) showed that members of Congress were willing to communicate with other members who shared political ideologies, who were from the same home state, and who had similar political opinions. We go one step further by directly connecting our empirical work to an established theoretical debate; moreover, our conceptual framework is directly informed by a game theoretic model of communication, patterns of leadership, and legislative organization.

Methodology

In summary, our analysis proceeds in three steps. First, we analyze the original tweets using a Joint Sentiment Topic (JST) model, which we believe is new to legislative studies. We use this JST model to produce estimates of the daily propensity to discuss a sentiment topic for each legislator. Second, to uncover the topics in need of direction, we use principal components analysis (PCA) on the member-level average of the topic weights to identify which topics best explain the variation between members' preferred discussion topics. Finally, we use a daily average of the topical weights for House rank-and-file and for the House leaders to test whether House leaders exert influence and lead on the messaging regarding a policy position or whether House party rank-and-file exert influence and lead discussion.

Joint sentiment topic analysis: We estimate a topic mixture and sentiment mixture, the JST model, which we believe is new to the study of legislative communication and behavior. It is based on Latent Dirichlet Allocation (LDA), though it estimates an additional latent layer. However, unlike LDA (which estimates two latent layers, topic classification and words alone), the JST estimates three latent layers (sentiment orientation, topic classification, then word mixtures). Importantly, the JST model estimates the unconditional probability of each sentiment. Note that this model is weakly supervised, as we place a weak prior over the sentiments' orientations for a selection of common words.

In order to measure the structure of communication, we use the JST method to classify all tweets for all House members over both sessions of Congress at once. Previous work in political science has used topic analysis to classify open-ended survey responses (Roberts *et al.*, 2014), while Kim, Londregan and Ratkovic (2018) have used text to augment an ideological spatial model. Our strategy is an amalgamation of these two approaches. Our work captures the discussion space, without relying on assumptions regarding exogenous covariates to uncover the latent topics.

By accounting for both topic and sentiment, a key feature of the communication structure uncovered by JST is the clear variation in how Democrats and Republicans communicate on social media. By uncovering this inter- and intra-party variation, we are able to analyze behavior within and across parties. Moreover, this method uncovers partisan separation in party communication, evidence that the unsupervised method has external validity. We strongly expect there is a partisan element to discussion on social media from the patterns of communication, which should be especially strong for our sample of members of Congress.

For all tweets in the dataset, we estimate a probability distribution for every word and every tweet which can be decomposed as:

$$\Pr(\text{Word} = w, \text{Sentiment} = j, \text{Topic} = k) = \Pr(\text{Word} = w | \text{Sentiment} = j, \text{Topic} = k) \\ \Pr(\text{Topic} = k | \text{Sentiment} = j) \Pr(\text{Sentiment} = j)$$

This produces a vector of kj independent sentiment-topic probabilities and j sentiment probabilities for each tweet, which are analogous to the estimates one derives from mixture membership topic models, such as Latent Dirichlet Allocation.

As with many standard topic model approaches, as we connect the JST model to political contexts, the model relies on exchangeability and is a bag-of-words approach to speech, which allows for feasible, tractable estimation. We provide a full technical overview in SI Section 4.1.⁶

To calibrate the model, we optimize the coherence score of the model. SI Figure SI 2 suggests that the optimal number of topics is 60 topics, the local maximum in the coherence score metric we employ – normalized pointwise mutual information. This is a measure of the extent to which, on average, words we say are likely to be in a topic to be associated in the same topic are actually associated based on what we see in the data. This measure is among the most accurate for determining quantitative coherence for uncovered topics Röder, Both and Hinneburg (2015). For the number of sentiments, we fix the number at 3, following the paradigmatic prior in Lin and He (2009). This results in 84 conditional sentiment-topic probabilities, and three unconditional sentiment probabilities for each tweet.

SI Table SI 2 highlights the tweets with the highest probability of belonging to their sentiment-topic label. We report the pre-processed tweet and the associated author-generated labels. The tweets in Table SI 2 highlight that the JST model produces coherent topic structure, in addition to mathematical coherence.⁷

Measuring need for direction: In order to measure need for direction on a policy, we examine structural notions of leadership derived from a PCA analysis of the sentiment-topic space. This is distinct from the topic-by-topic analysis in the preceding section as here we look at measures of party behavior at the party level.

Communication decisions among House members are likely guided by exogenous events, party and peer effects, and personal preferences of legislators, which are not immediately obvious from looking at the raw mixtures at the document level. So to understand the individual-level data, we aggregate document-level data by averaging the topical weights for each member. By using PCA as a dimension reduction technique on this aggregate individual-level data, we identify topics that explain the variation in what members in Congress discuss relative to one another. Figure 1

⁶This is also reviewed in Lin and He (2009) and Lin *et al.* (2012).

⁷For additional details, see Supplementary Information 4.1.

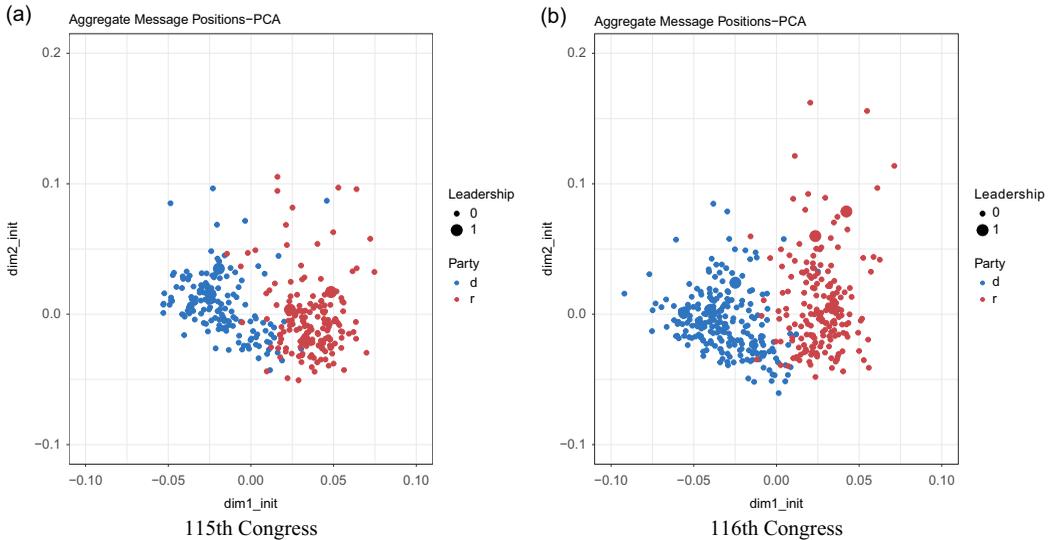


Figure 1. Aggregated legislator policy positioning in the two-dimensional topic space derived from the PCA analysis of the sentiment-topic propensities for the 115th (left) and 116th (right) Congresses. Red indicates a Republican member's policy position, blue indicates a Democratic member's policy position.

illustrates the sentiment-topic space for all members in our data, summarized by member for the entire period covered by the dataset. We call the coordinate pairs in this figure the policy position for each legislator.⁸

After computing the JST mixtures for each tweet, we find the average probability a House member tweeted about a particular sentiment topic for the 115th and 116th Congresses. We then employ PCA to reduce the sentiment-topic space to two dimensions. Next, using the PCA results, we examine which sentiment topics contribute strongly to the PCA solution – these are sentiment topics that tend to drive legislators to the extreme areas of the sentiment-topic space. These are considered sentiment topics that are low in need for direction. These are topics where the coordination problem dominates, and thus we expect leaders to initiate discussion.

Similarly, we look at the sentiment topics that do **not** contribute strongly to the PCA solution – these are sentiment topics that do not drive legislators toward the extremes of the space and thus are topics highly in need of direction. These are topics where the information problem dominates, and thus we expect members to lead. Later, in Section 4.1 we will show the topics that are in need of direction based on this PCA analysis.

We emphasize that these PCA results measure a position in sentiment-topic space over popular debates taking place on social media in real time. PCA analysis allows us to analyze more readily the hundreds of thousands of messages espoused by legislators on social media. PCA is useful when taking our JST model as input, as JST accounts for both sentiment orientation and topic content. This allows the latent partisan structure of the data to be detected, without imposing additional structure from potentially endogenous variables to induce this structure. The output of this mapping is a two-dimensional coordinate for each legislator in Twitter communication space for each Congress. From these individual-level measures of communication, we can identify topics that need policy direction or not. These topics form the basis of our empirical tests of the hypothesis regarding party leaders' ability to coordinate.

⁸We also estimate this measure restricted to only the respective Congress. Figure SI 4 shows the contrast of rank-and-file members' positions in the PCA-derived Twitter communication space when we estimate it separately. We show there that the main result is robust to changes in this estimation routine.

Dynamic analysis: Finally, we exploit the micro-level data to examine whether House leaders exert influence and lead discussion on Twitter within their party coalition (and thus exert influence and lead discussion over their rank-and-file), or whether they adopt their members' consensus. As we have stationary data (see SI Figures SI 7 and SI 8), we follow the time series strategy employed in Barbera *et al.* (2019). We measure daily propensity to discuss a sentiment topic in precisely the same way – except using the posterior probability estimates of sentiment-topic JST mixture weights. This is the daily average probability of a House member discussing a particular topic with a particular sentiment orientation. Here, influence is measured by the impulse response functions (IRF) from a vector autoregression (VAR), and we say members or party leaders exert influence and lead when these IRF estimates are statistically and substantively significant.

As our data are stationary, but censored between 0 and 1, as in Barbera *et al.* (2019), we follow Wallis (1987)'s logit specification for VAR. However, our specification contains only two endogenous variables: the average propensity to discuss a sentiment topic by leader and rank-and-file within each party. We make this choice for two reasons: first, because the theory makes predictions over which types of topics should facilitate the emergence of leadership within individual parties, we estimate VARs separately for each topic and party to evaluate the extent that party leaders emerge as theory predicts. Second, the parameter space is large. Thus, the system of equations may not be identified for a reasonable number of lags. The within-topic analysis allows us to identify more lags and improves computational tractability. At the same time, it also avoids introducing spurious correlations, given the highly interrelated nature of the data. Finally, in cases where the nature of the structural relationships is not known to the researcher, interpreting the results from a VAR regression is difficult. Our parsimonious specification allows for a more direct analysis.

For our specification, we fix a sentiment-topic label k where k can take on one of three possible values: positive, negative, and neutral. Let $x_{mem,t}^k$ and $x_{lead,t}^k$ denote the probability of the average member and average leader respectively discussing a sentiment-topic label k . Let $X_t^k = (x_{lead,t}^k, x_{mem,t}^k)$. Then let

$$Z = \log\left(\frac{X}{1-X}\right)$$

Our specification thus is:

$$Z_t^k = c^k + \sum_{p=1}^7 \beta_p Z_p^k + \varepsilon_p^k$$

Here c is a constant accounting for the fact the time series are stationary around a non-zero mean after taking logs. SI Figures SI 6 and SI 5 show for selected series that the time series in log-odds of daily propensity to discuss sentiment topics are stationary over our period of analysis. Furthermore, SI Figures SI 6 and SI 5 show that we reject, at the 1% level, a null of unit roots for the vast majority of our time series for the Democratic and Republican Parties across both the 115th and 116th Congresses. These are key assumptions of VAR analysis, and these results indicate that our data are consistent with them. Finally, we choose a lag of 2 days, which captures the length of the news cycle on Twitter.⁹

Finally, to capture the extent that House leaders or followers exert influence and lead discussion, we estimate generalized impulse response functions for each specification following

⁹We also tried a method where we selected the optimum lags based on an AIC criterion, but we found the optimal number was always around 2 days, so we chose to fix the number of lags, given that this fixed number induces a consistent number of lags across the specifications and did not substantively alter the results. In fact, choosing lags of 1, 5, and 7 days did not significantly alter the results.

Koop, Pesaran and Potter (1996).¹⁰ That is, we measure the effect of a two-standard deviation increase in a party leader's log-odds of discussing a given sentiment topic on the average members' log-odds of discussing that topic and vice versa. Using the median daily propensity to discuss a sentiment topic as a base rate, we convert the log-odds to relative risk. Using the relative risk, we estimate the change in daily propensity as a percentage point increase over the base rate in the contemporaneous period of the shock. We report 95% bootstrapped confidence intervals with 500 draws.

Testing the implications of the hypothesis

The theoretical framework from Dewan and Myatt (2007) suggests a clear hypothesis regarding how House party leadership influence relates to party communication. In this section, we connect the theoretical framework to our empirical setting. See Table 1 for a road map to our analyses.

Need for direction

To test the hypothesis that House leaders exert influence and lead discussion when the need for policy direction is low (and the coordination problem dominates) and high (when the information problem dominates), we first need to uncover when leaders exert influence and lead discussion and when rank-and-file members influence discussion. We chose a threshold approach here, as we wanted a principled, data-driven approach to delineating the topics.

Alternative theories

The key concept that we test from Dewan and Myatt's work is the tension between information aggregation and coordination when the stance the party should take is not clear. An intuitive story from extant theories such as conditional government or strategic party government (Aldrich and Rohde 2001; Koger and Lebo, 2020) is that when parties are cohesive, there is high party agreement. Then, when there is heterogeneity, we observe less party agreement. The key theoretical insights from these papers are cross-sectional, yet under these theories, the time dynamics are not well-theorized. In our case, we want to study instances where the parties start heterogeneous and become united, or start and stay united, or start and stay heterogeneous, or start united and then become heterogeneous. By employing Dewan and Myatt (2007)'s framework, we are better able to analyze conditions when this may happen. More importantly, we explicitly test when party members begin employing similar messages to their leaders (and vice versa), not explicitly party-level coherence on messaging. Thus, our results emphasize the time dynamics, not the cross-sectional variation between and within parties, which by itself has many potential explanations.

Moreover, the predictions from Dewan and Myatt's theory sometimes produce counter-intuitive results: when the parties are heterogeneous, they are more likely to unite behind the caucus consensus (since they defer to the "wisdom of the crowd"), whereas when they are united, they are more likely to defer to the leader (as they wish to coordinate). In both cases, the party presents a united front, but for very different underlying reasons. In the first, they begin disunited and rally around a party consensus. In the second, they begin united and rally behind the message of the party leader. That is, after a period of time, both types of messaging give the appearance of polarized party messaging, but the theoretical frame we employ teases out very different

¹⁰Generalized impulse-response functions IRFs are invariant to variable ordering, unlike orthogonalized IRFs, while still allowing the researcher to study relationships with non-zero entries in the variance-covariance matrix, unlike the forecast error IRF. The magnitude of this IRF is how we derive our second notion of leadership, as noted in Table 1. That is, for an n step-ahead response, we compute $\Theta_i^k(n) = \frac{\delta_i}{\sigma_\varepsilon^2} \sum_{\varepsilon} \beta$ where δ is two standard deviations of our data, approximately 10%.

Table 2. Principal components analysis topic contributions – Leader Driven 115th

Topic	Contribution
Tax policy benefits – Positive	13.79
Tax cuts – Positive	4.92
Enjoyable visit – Positive	4.44
Protect health insurance – Neutral	3.91
Tune in/Watch cable news – Positive	2.83
Family separations – Negative	2.55
NDAA passage – Negative	2.28
Middle class tax cut – Positive	2.28
Opioid task force – Negative	2.12
Enroll in ACA – Positive	1.96
Pro trump mobilization – Positive	1.80
Jobs/Economy – Positive	1.76
Agriculture – Positive	1.73
Signed legislation – Negative	1.72
Trump asylum Policy	1.66
Prevent gun violence – Negative	1.57
Abortion rights – Negative	1.53
Manufacturing jobs – Neutral	1.51
DACA policy – Positive	1.49
Trump/Russia investigation – Negative	1.39

underlying mechanisms for how we arrived there (one bottom-up, the other top-down). We now explore the specific conditions when we expect the parties to follow the leader or follow the members.

Coordination problem: In Tables 2 (115th Congress) and 3 (116th Congress) (Table 3), we show the sentiment topics that define issues where the coordination problem dominates.¹¹ Our criterion for determining whether each topic needs direction is based on this percent contribution to the variation of the top two components derived from the PCA. We take the top twenty topics that contribute to variation in the member-level propensity to discuss sentiment topics for each Congress, and classify those topics as being low in need for direction.¹² Sentiment topics with low contribution to the variation in the sentiment-topic propensities do not drive legislators toward the extremes of sentiment-topic space, while large contributions drive them to the extreme portion of the space. As Figure 1 shows, policy positions for House members on these sentiment topics often delineate membership in a particular party. Thus, for sentiment topics that drive separation in this space (for example, immigration), we expect little coordination from party leadership, regardless of party, precisely because these are policy positions that delineate belonging to a particular party. In theory, it is on these types of partisan topics that leaders have the most influence over the rank-and-file since the outsized costs or benefits of coordinating on the wrong messaging are low.

Information aggregation problem: We classify the topics not in the top twenty as sentiment topics as in high need of policy direction. These topics do not contribute to variation in the

¹¹In the supplementary information we provide the PCA topic contributions for member-driven topics: Table SI 3 for the 115th Congress and Table SI 4 for the 116th Congress.

¹²We use the top and bottom twenty topics from the PCA so that we have some principled way to determine which topics to focus on for subsequent analysis. We establish these thresholds because the theory we are relying upon does not provide guidance about how many issues in this multidimensional space we should include in our analysis. Examination of the PCA results shows that the top or bottom twenty include the topics that most subject-matter experts would consider to be low or high in the need for direction. Changing the threshold from top 20 to 15 or even 10 does not materially change the results. Generally the topics that are outside of these thresholds regard conversations with Happy Birthday messages to constituents and colleagues, or congratulating sports teams in the district after a victory. They typically are not topics of strategic importance for party communication.

Table 3. Principal components analysis topic contributions – 116th leader driven

Topic	Contribution
Tune In/Watch cable news – positive	12.53
Impeachment – Negative	11.74
USMCA/Trade deals – Positive	5.96
Republicans attack democrats as socialists – Negative	5.03
Humanitarian aid at border – Negative	3.09
Trump/Russia investigation – Negative	2.94
Tune in/Watch interview – Negative	2.86
COVID economic relief – Positive	2.52
Lowest unemployment rate – Positive	2.44
Census encouragement – Positive	1.90
Wear a mask – Negative	1.53
Religious freedom – Negative	1.46
Climate change – Positive	1.38
Partisan attacks on Trump/Biden-Negative	1.36
Border crimes – Negative	1.36
Criminal justice reform – Negative	1.31
Jobs/Economy – Positive	1.28
Racial inequality in health care – Positive	1.26
Public health and safety – Neutral	1.26
Snap benefits – Positive	1.18

propensity to discuss topics by rank-and-file members of the House. We argue these remaining sentiment topics, many of which explain less than 1% of the variation in the individual propensities to discuss sentiment topics, represent sentiment topics where the underlying political fundamentals of the topics are more uncertain, so the information aggregation problem dominates. In this case, failure to coordinate would be preferable to coalescing around the wrong message. For example, on arcane matters of budgetary politics, the optimal message is not immediately clear. The parties may not coordinate on any message, but that might be preferable to coordinating on a message that would be bad for the party. (For the Democrats, they might coordinate on raising taxes, or for Republicans, they might coordinate on cutting Social Security. Neither position would be particularly popular.)

House leadership influence

To test the hypothesis that party leaders exert influence and lead when the need for direction is high, for each party, we measure the autoregressive correlations between the average propensity to discuss a leader-driven topic with the average propensity of the rank and file. We would reject the hypothesis of leader or member influence if the observed number of significant effects in the expected direction was indistinguishable from zero. For example in Figure 2, there is a 1.6% probability of observing the realized number of significant effects in the expected direction. We report this statistic in the caption for each figure, and find all of the tests consistent with the theory. To quantify influence, we employ IRF analyses from a vector-autoregression, similar to Barbera et al. (2019). The IRFs enable us to quantify the ability of House leaders to exert influence and lead discussion. We regress the the party leadership's average daily propensity to discuss a sentiment topic on that of the party rank-and-file members, and vice versa. The IRF analysis then supposes a hypothetical shock to the leadership's propensity to discuss a sentiment topic and estimates the increase in the propensity of rank-and-file member's to discuss. If this shock is statistically significant, we say House leadership influences rank-and-file members' propensity to discuss a sentiment topic. We also test the reverse – the influence of rank-and-file members on leadership's propensity to discuss a topic.

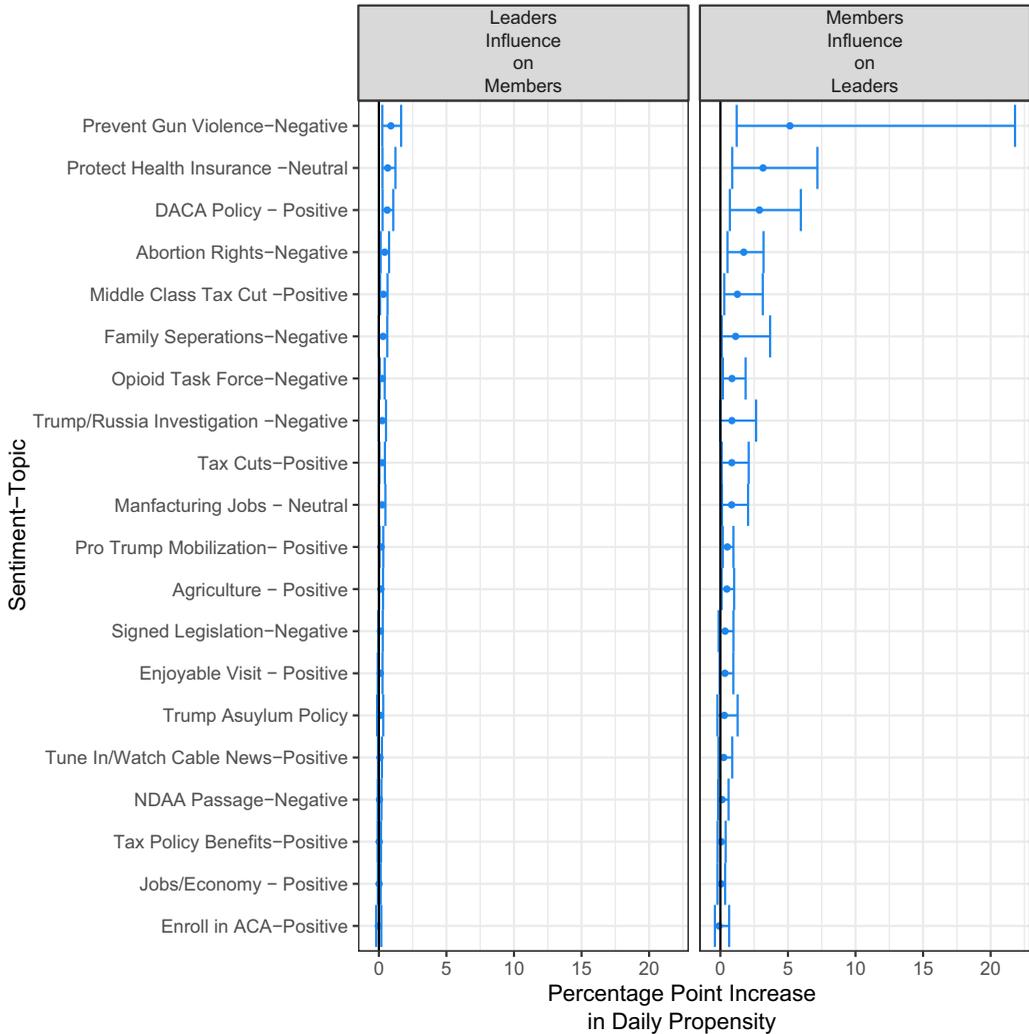


Figure 2. Democratic topics: need for direction predicted leader-driven 115th Congress. Impulse response functions for sentiment topics predicted to be leader-driven for the Democratic Party. Bootstrapped 95% confidence intervals are shown. The probability of observing the realized number of predicted effects in the expectation direction is 0.016.

Results

Need for direction by leadership – coordination problem

We find evidence consistent with the theory outlined in the previous sections. The IRF analysis suggests leaders can increase the rank-and-file’s propensity to discuss the most partisan topics by between 0.1% and 1% for each standard deviation increase in the leadership’s daily propensity to discuss a topic. These are substantively large – shocks of 3 or 4 standard deviations (40% to 60%) on the daily propensity to discuss a topic are common, so finding discernible effects at the more conservative level of 1 standard deviation suggests the result is stronger under conditions that are normal for social media. This reflects the nature of conversation on Twitter, which reacts sensitively to the news cycle. This result is consistent across parties and time, even when the party in power changes. This consistency is evidence that the result is robust across these same dimensions, during the period of 2017 to 2021.

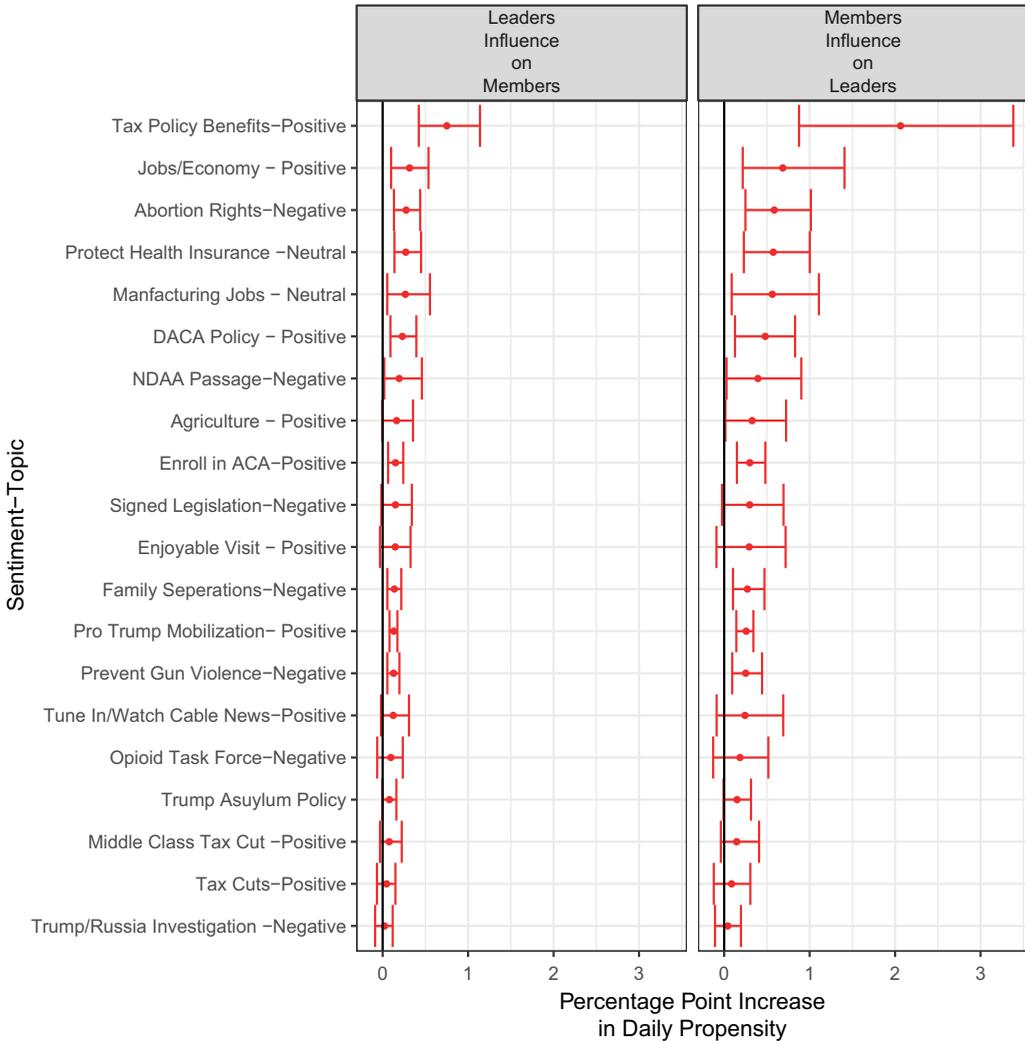


Figure 3. Republican topics: need for direction predicted leader-driven 115th Congress. Impulse response functions for sentiment topics predicted to be leader-driven for the Republican Party. Bootstrapped 95% confidence intervals are shown. The probability of observing at least the realized number of predicted effects in the expectation direction is 2^{-11} .

In Figure 2, we show the impulse response functions in the first period for the Democrats in the 115th Congress for topic-sentiments that are low in needing direction. Democratic leaders in this period exert statistically significant levels of influence for messaging around preventing gun violence, protecting health insurance, abortion rights, and DACA policy. These topics make sense as having low need for direction – in these cases, the Democrats desired retaining the status quo (preserving Obamacare, DACA) or were discussing topics that are central to Democratic Party ideology, such as abortion and gun violence. In both cases, the party needs little direction in terms of their stances on these issues, so the party would rather coordinate on some message than no message at all.

For Republicans in the 115th Congress, Figure 3 shows that economic sentiment topics are statistically significant. Given the overall strength of the economy from 2017 to 2018, the Republicans benefited politically from raising the salience of the economy. We interpret this result as evidence that mis-calibrating the message on the positive economy was less costly than not coordinating on any message at all.

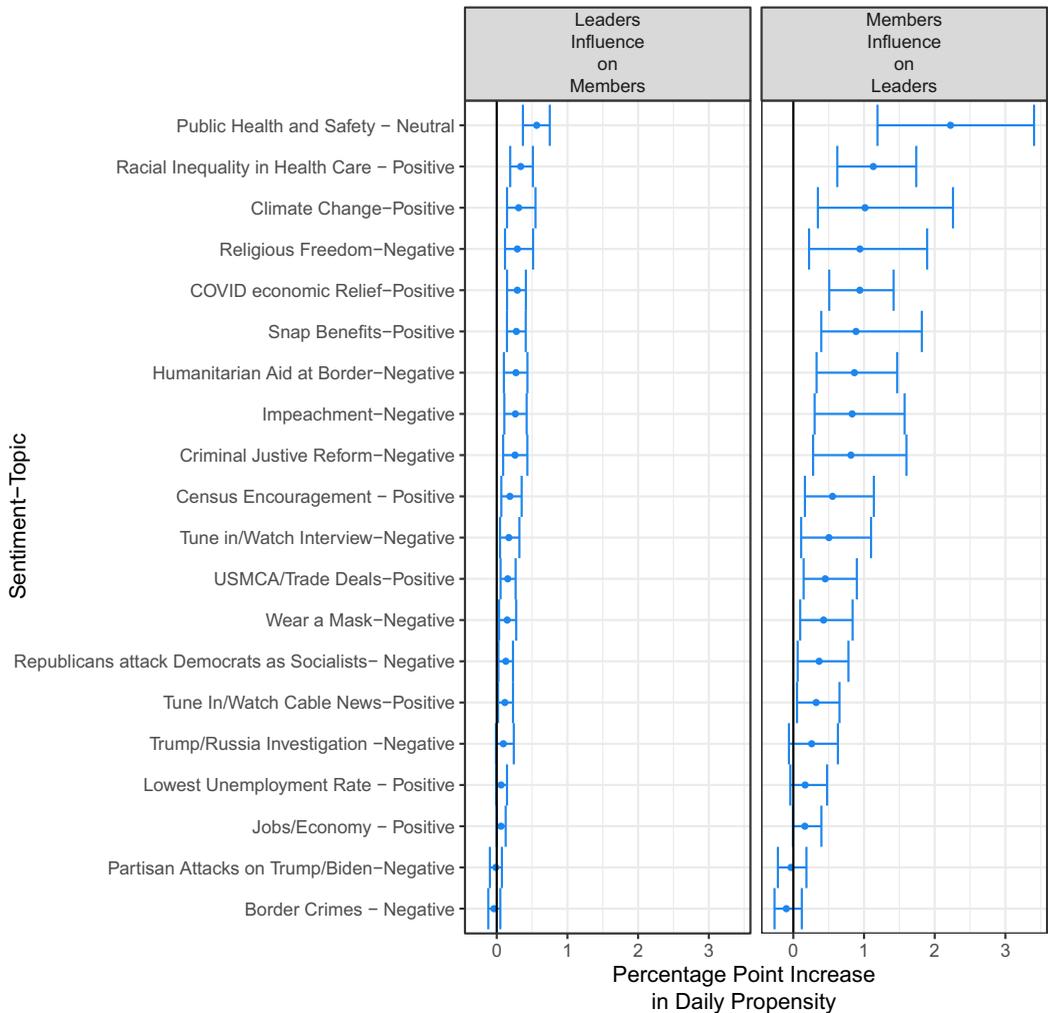


Figure 4. Democratic topics: need for direction predicted leader-driven 116th congress. Impulse response functions for sentiment topics predicted to be leader-driven for the Democratic Party. Bootstrapped 95% confidence intervals are shown. The probability of observing at least the realized number of predicted effects in the expected direction is 2^{-11} .

In Figure 4, we show the impulse response functions for the Democrats in the 116th Congress for topic-sentiments that are low in needing direction. Democratic leaders in this period exert statistically significant levels of influence for messaging around public health topics, COVID economic relief, climate change, and impeachment. Similar to the 115th Congress, these topics are consistent with being in low need of direction. In these cases, the Democrats discussed two types of such issues. In the first type, they raised the salience of issues where Republicans faced political downside risk (for example impeachment). Second, they discussed topics that are central to the Democratic Party’s ideology, such as racial equality and public health.

Republicans in the 116th Congress exhibit similar behavior to the Democrats in the 116th Congress. For the Republicans, Figure 5 shows that shocks to leaders’ daily propensity to discuss a particular issue generally results in a less than 1% increase in the rank-and-file members’ daily propensity to discuss that issue. In particular, Republican leaders induced a ~ 1 percentage point increase in their rank-and-file members’ propensity to discuss impeachment and

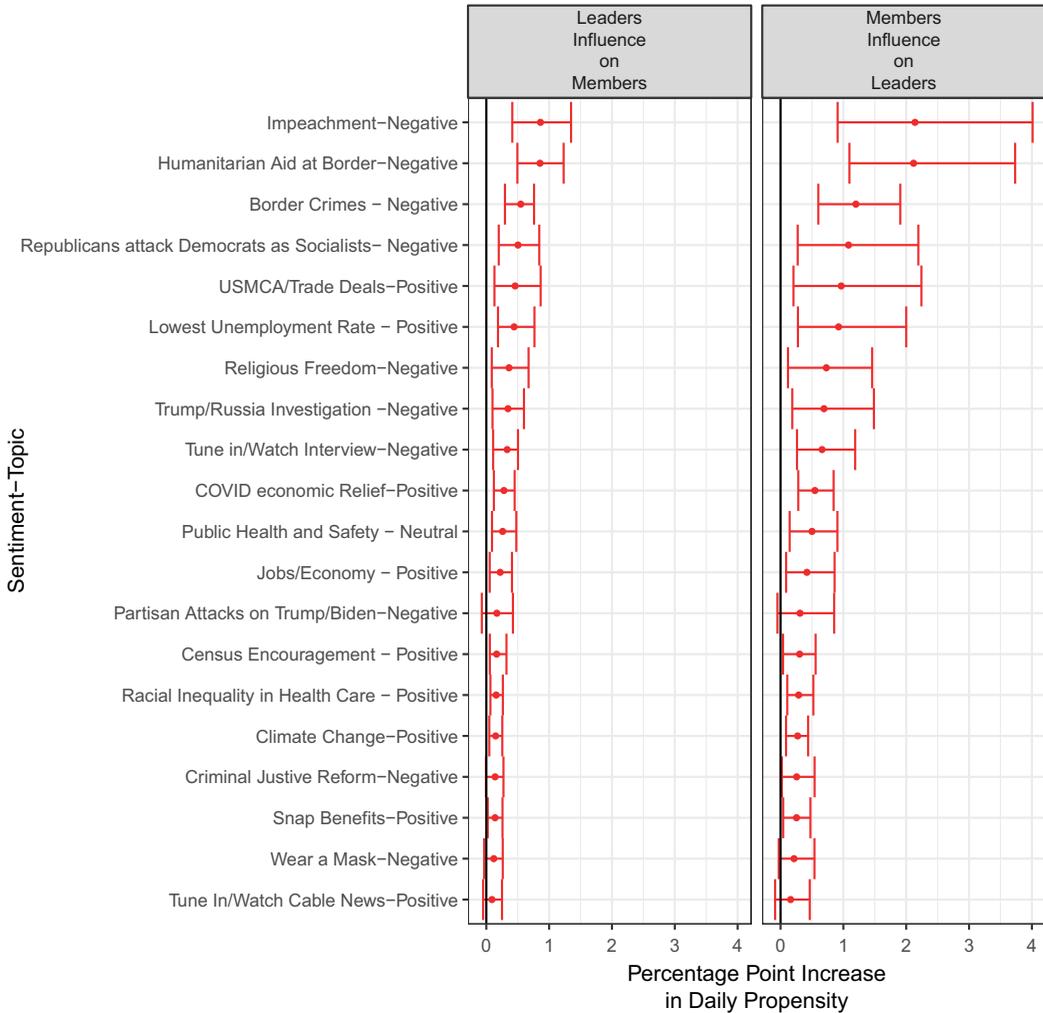


Figure 5. Republican topics: need for direction predicted leader-driven 116th Congress. Impulse response functions for sentiment topics predicted to be leader-driven for the Republican Party. Bootstrapped 95% confidence intervals are shown. The probability of observing at least the realized number of predicted effects in the expectation direction is 1.7^{-14} .

freedom/sacrifice, and border security. Leaders induced a 0.5 to 1 percentage point increase for impeachment, crimes at the border, attacking the Democrats as socialists, USMCA, and lauding the low unemployment rate. Figure 5 also shows that members induced a ~ 2 percentage point increase in their leadership’s propensity to discuss impeachment and humanitarian aid at the border. Members exerted a ~ 1 percentage point increase in their leaders’ propensity to discuss crimes at the border and attacking the Democrats as socialists. Additionally, they exerted a nearly 1 percentage point increase for trade deals and USMCA, and lauding the low unemployment rate. Again, members’ influence is an order of magnitude larger than the leadership’s influence. Notably, the magnitudes derived for Republicans leadership and rank-and-file members are similar to those for Democratic leaders and members. This suggests that party leaders and members are similarly responsive to each other with respect to their messaging regarding their propensity to discuss sentiment topics, regardless of party.

These results show consistent patterns in legislators' social media behaviors. Party leaders exert influence over the messaging agenda in precisely the topics that are consistent with the theory. In fact, the results for the coordination problem are consistent across time periods, parties and the changes in the party which controls the House of Representatives.

Need for direction by membership – information problem

Next, we examine in detail the behavior of congressional parties for topics where we believe the information aggregation problem dominates. Intuitively, the information aggregation problem dominates the political environment when there are large costs to the party for choosing the wrong policy. This problem tends to arise when there is more uncertainty in the political environment, be it related to the nature of the political problem, the eventual policy outcome, or the electoral ramifications of taking a policy stance. For example, in a government shutdown scenario, whether to continue the shutdown carries large risks. It may galvanize the base of the party taking the strong stance and increase turnout in favor of the party. Or potentially just as likely, this stance may harm the economy and thus dissuade swing voters from supporting the party. In either case, the potential risks are large. In the case when the information problem dominates, the party relies on “the wisdom of the crowd” of the party at large. By aggregating information, the party hopes to coordinate on the “correct” message, even if this risks not coordinating on any message at all. In these cases, the costs of coordinating on the wrong message outweigh the costs of failing to coordinate.

Our results for topics predicted as member-driven are consistent with this theory. Specifically, Figure 6 shows that Democratic House members exerted the most influence over the propensity to discuss Supreme Court nominations (approximately a 4 percentage point increase for each standard deviation shock) and wishing thoughts and prayers after a crisis (a ~ 2.8 percentage point increase). However, across these same topics, leaders' influence is either statistically insignificant at traditional levels or is near 0. Notably, the effect sizes for members on leaders are an order of magnitude greater than the leadership's influence on rank-and-file members.

The Republicans messaging between leaders and rank-and-file is more tightly correlated, but we see that the influence exerted by members is less than influence exerted by Democratic rank-and-file members on their leadership. Rank-and-file members drive a 1.5 increase in both the propensity for leaders to discuss the low unemployment rate and also thoughts and prayers around a tragedy. Notably, as illustrated by Figure 7, rank-and-file members exert a $\sim 1\%$ increase on the propensity to discuss important meetings. We hypothesize this is an obfuscation messaging strategy. Given the majority party runs the risk for being blamed for negative economic and social conditions in the country, this result is preliminary evidence that majority parties find it advantageous to engage in measurable amounts of political deflection.

The results for the 116th Congress follow a similar pattern for both parties. Figure 8 shows that the Democratic rank-and-file membership exerts a 2% to 3% effect on the topics that are in need of direction, whereas leaders exert little influence on these same topics. In the 116th Congress, Democrats became the majority party. Despite this change in institutional control, party communication behavior on social media is consistent with the 115th Congress. Notably, decrying partisan votes – an obfuscation and deflection message – is now one of the key topics where rank-and-file Democratic members exert influence on their party leaders. This is similar to the obfuscation tactics among the Republicans rank-and-file when they were in the majority in the 115th Congress. This supports the prediction from the theoretical framework that parties would rather fail to coordinate than coordinate on the wrong message.

In the 116th Congress, the Republican rank-and-file behaves a lot like they did the 115th – and a lot like their contemporaneous Democratic colleagues during the 116th Congress. Figure 9 shows that impulses of a standard deviation to the leaders' daily propensity to discuss a particular issue generally results in an approximately 0.5% to 1% increase in the rank-and-file members' daily

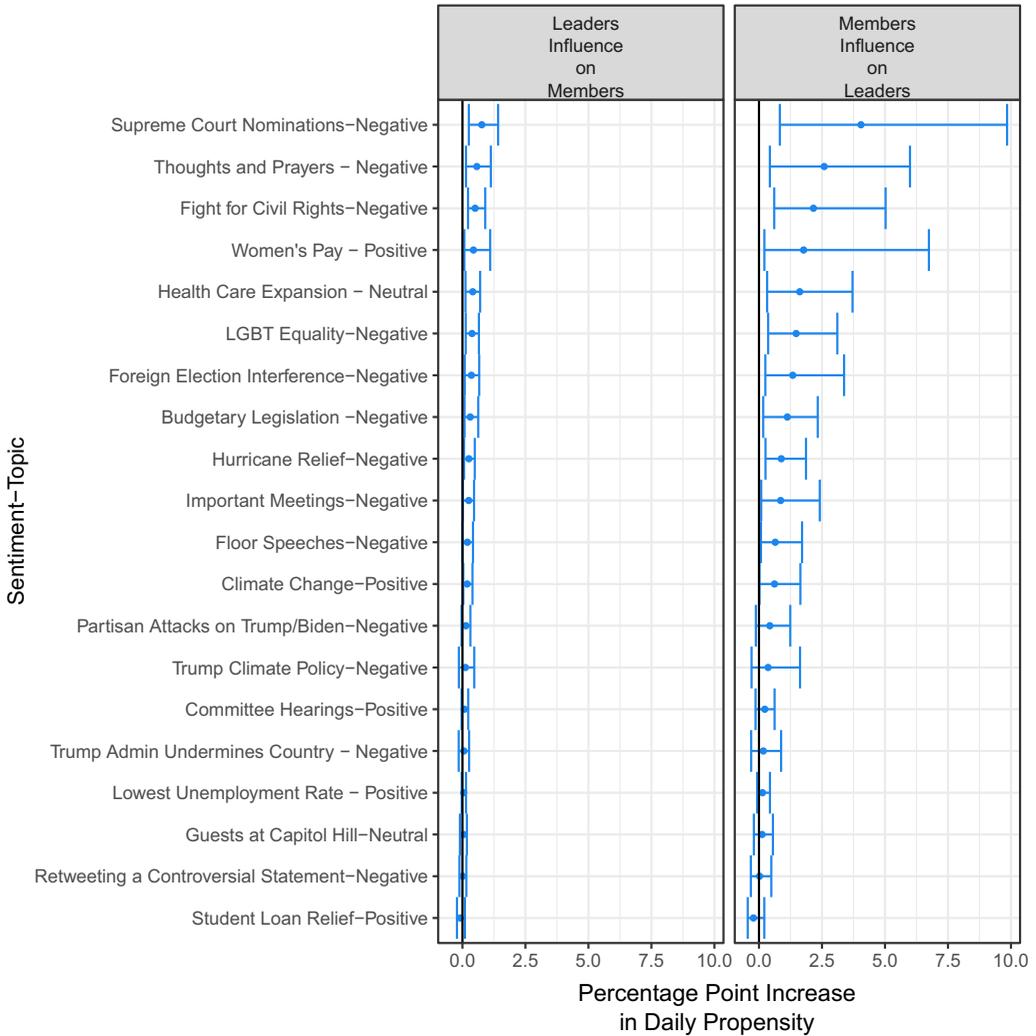


Figure 6. Democratic topics: need for direction predicted member-driven 115th Congress. Impulse response functions for sentiment topics predicted to be member-driven for the Democratic Party. Bootstrapped 95% confidence intervals are shown. The probability of observing at least the realized number of predicted effects in the expected direction is 6⁻¹³.

propensity to discuss that issue. As in the 115th Congress, leaders and rank-and-file members both exert influence over these topics, but rank-and-file members’ influence is an order of magnitude larger than the leadership’s influence. Notably, the magnitudes derived for Republican’s leadership and rank-and-file members are smaller than those for Democratic leaders and members. This suggests that party leaders and members are similarly responsive to each other in relative terms between members and leaders, though the magnitude of that influence varies between parties. Additionally, the Republicans, who controlled the presidency, continued to obfuscate, decrying partisan votes and discussing positive constituent visits to their congressional offices.

Discussion

We highlight the consistency of these findings across the parties and the substantive robustness: on issues where House rank-and-file influence discussion, their effect on leaders is larger in

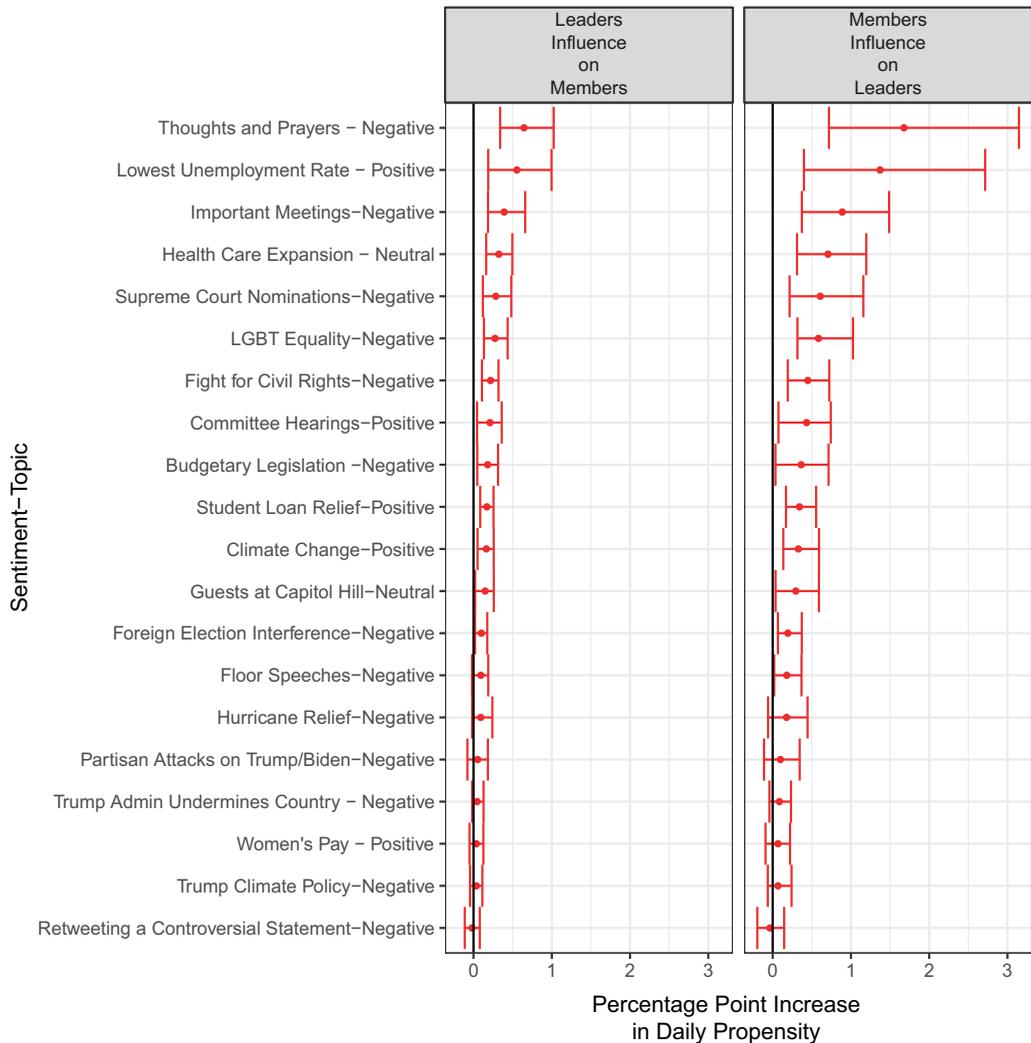


Figure 7. Republican topics: need for direction predicted member-driven 115th Congress. Impulse response functions for sentiment topics predicted to be member-driven for the Republican Party. Bootstrapped 95% confidence intervals are shown. The probability of observing at least the realized number of predicted effects in the expected direction is 5^{-10} .

magnitude than on issues where leaders lead. This is true across topic types, as illustrated in Figures 2, 3, 4, and 5. So, while leaders and rank-and-file influence each other, the measurable effects from rank-and-file are stronger than those on leaders for issues where they respectively had influence. In each of the above instances, we find evidence that the results are purely random noise is highly improbable, and thus we argue these tests are consistent with Dewann and Myatt’s theory; however, we note that some tests are significant for both members and leaders. One gap in our operationalization is that we make no claims regarding how members will respond to leaders on member-driven topics (and vice versa). Although these are outside the framework from which we generate our hypotheses, we include those results because they allow us to compare the relative influence of members and leaders on the same topic. On our IRF metrics, we find members consistently exert more influence over leaders than the reverse. This is true across parties, across congresses, and across topics. This is perhaps a surprising finding and one we wish to highlight for

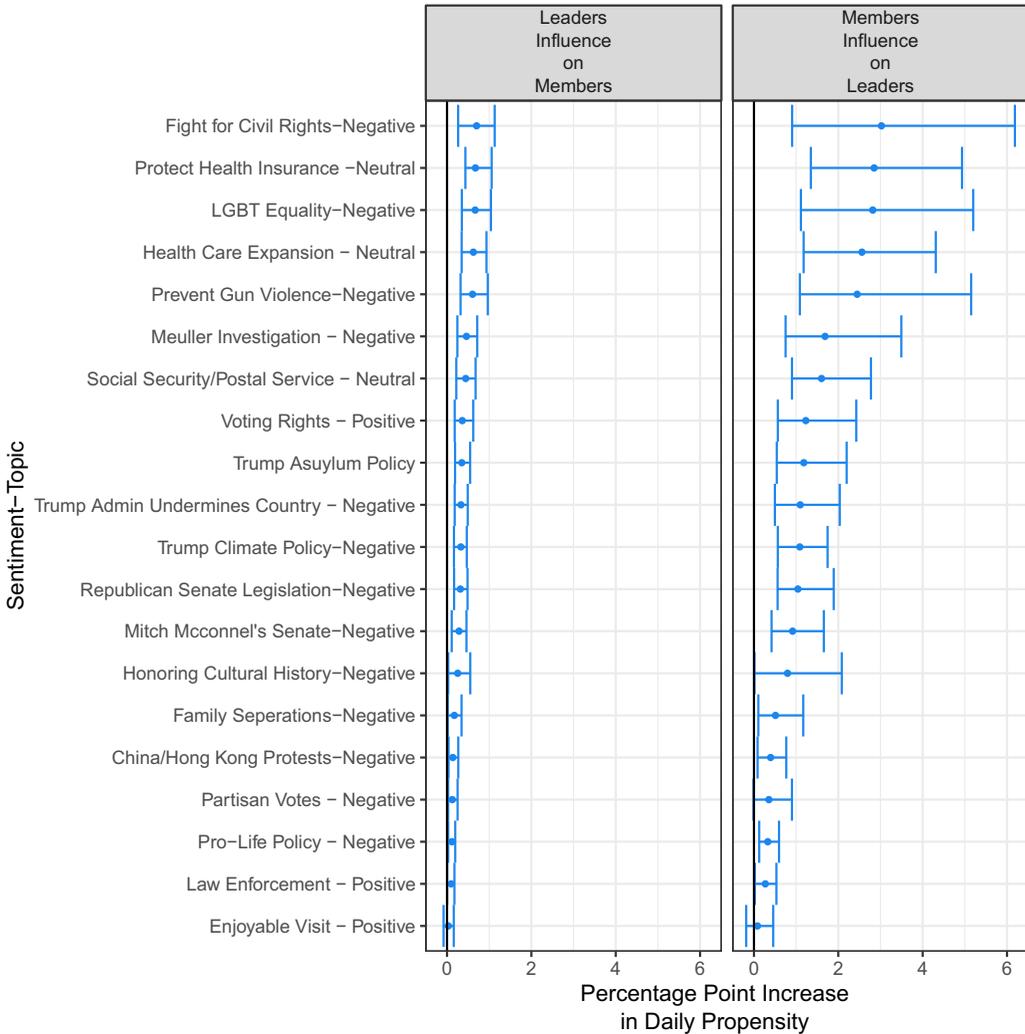


Figure 8. Democratic topics: need for direction predicted member-driven 116th Congress. Impulse response functions for sentiment topics predicted to be member-driven for the Democratic Party. Bootstrapped 95% confidence intervals are shown. The probability of observing at least the realized number of predicted effects in the expected direction is 1.77⁻¹⁴.

theoretical and empirical researchers studying how influence flows in Congress. In these cases, we hope researchers continue to push the methodological frontiers in the time series dynamics of highly correlated series.

Finally, in Table SI 5 we note that leaders exert on average more influence than the most followed accounts in each party. On average, leaders exert double the influence as the most followed accounts from within the same party. This highlights the strength of institutional leadership within the party caucus relative to the influence of members of the party who are popular with the public on social media.¹³

¹³We also show in Table SI 5 that leaders exert nearly double the influence on their own members than leaders from the other party exert on the members of the opposing party, suggesting the result is not due to trends on social media. Instead, this result suggests that the role of leaders within their own party explains the result.

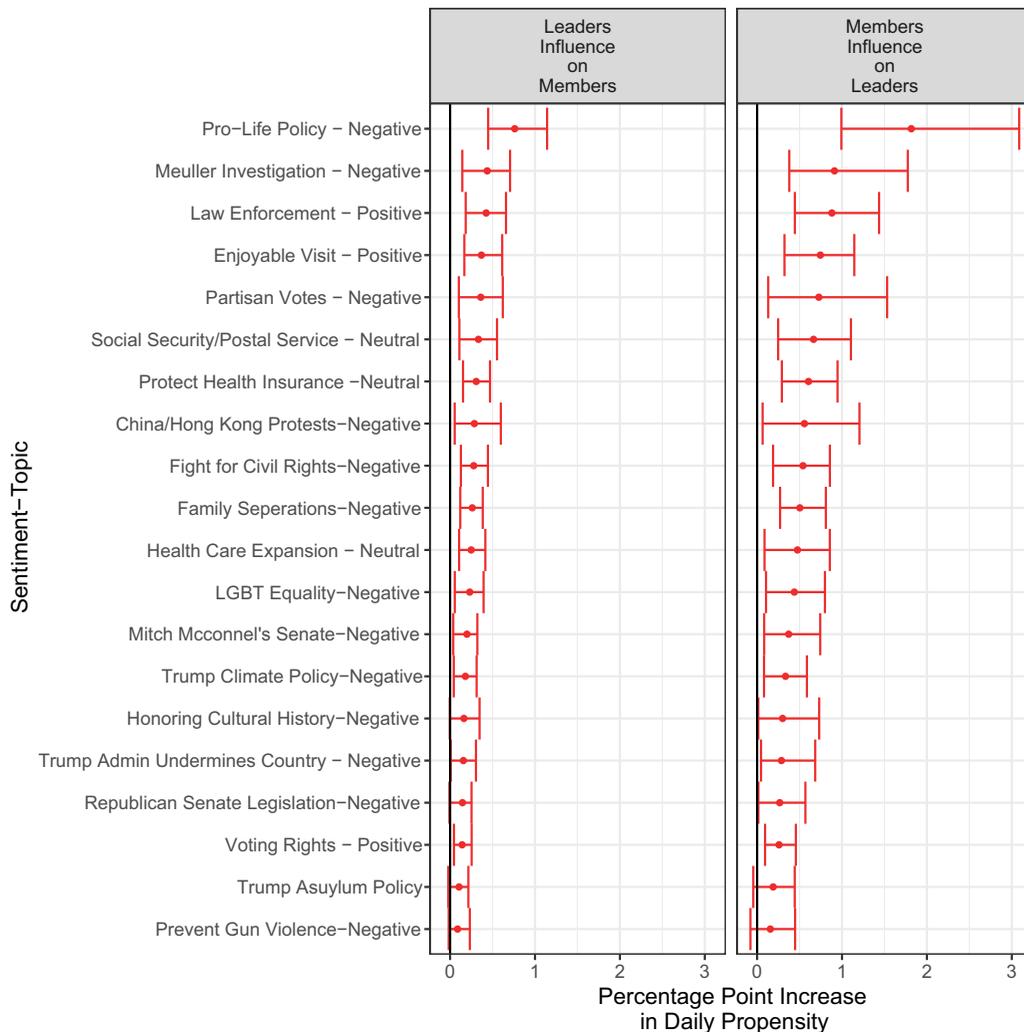


Figure 9. Republican topics: need for direction predicted member-driven 116th Congress. Impulse response functions for sentiment topics predicted to be member-driven for the Republican Party. Bootstrapped 95% confidence intervals are shown. The probability of observing at least the realized number of predicted effects in the expected direction is 6.09^{-18} .

Conclusion

Who controls the legislative messaging agenda has important consequences in a democracy. Currently, the literature on legislative agenda setting suggests that the agenda is driven by national polarization. However, other theories, such as formal models of legislative leadership, assert that legislative messaging strategies depend importantly on the information and political environment. In particular, these formal theories argue that legislators shift their messaging as they balance coordination and information problems. Thus these formal theories predict that when coordination problems are pressing, legislative members follow the policy positions of party leaders.

Our research contributes to the study of legislative leadership, messaging and agenda setting by putting a formal theory of party leadership to the test. We have presented evidence using social media data that the Dewan and Myatt (2007) theoretical framework of party leadership helps explain patterns of communication and leadership in the U.S. House of Representatives by

highlighting the tensions between the need of congressional political parties to coordinate around a unified policy position and the uncertain nature of politics. We present empirical support for our hypothesis that House party leaders exert influence and lead discussion on topics that do not need policy direction, while members exert influence discussion on topics where topics do need policy direction, mediated by information aggregation. To this end, we find that, given a large enough shock to the House leadership's propensity to discuss a sentiment topic where the coordination problem dominates, leaders exert a statistically significant influence in the short-run over their rank-and-file members' propensity to discuss that sentiment topic. Notably, this effect also operates when the information aggregation problem dominates, with influence flowing from rank-and-file to leaders. Moreover, when House rank-and-file members experience a shock to their propensity to discuss a sentiment topic, leaders are more strongly impacted than in the reverse case. For a standard deviation (~10 percentage point) shock to leadership's propensity to discuss, we might observe 0.5% to 2% increases in rank-and-file's propensity to discuss. For the reverse, we see a standard deviation (~10 percentage point) shock to House rank-and-file's propensities to discuss a sentiment topic results in a 1 to 3 percentage point increase in leadership's propensity to discuss a sentiment topic.

This suggests a complex interplay between leaders and members, which is in line with the theory and consistent across parties, changes in partisan control of the legislative institutions, and fundamental changes in the underlying political environment. We find evidence from the IRFs suggesting that leaders exert influence over their members on topics that come to dominate social media discussion. Furthermore, in those cases where members influence leaders, their effect on the messaging of leadership is nearly double that of leadership on rank-and-file members. That is, House leadership and rank-and-file messaging on Twitter influence each other. However, when rank-and-file members drive discussion, their effect is far larger than that of leadership. Thus, using this theoretical model to specify the coordination-information trade-off, we use our data to shed light on the situations where legislative party members resolve tensions between a coordination problem and an information problem.

We believe this theoretical framework provides a blueprint for studying how communication on social media reveals legislative party behavior, and our work demonstrates ways to measure and test a relevant hypothesis derived from the theory. Future work should more precisely classify topics in need of direction versus those that are not. They may also test notions of leadership.

Our research helps demonstrate that social media data is useful for studying legislative behavior and organization. We test formal political theory with social media data using machine learning methods, in line with the recent trend to more closely connect formal political theory with strong quantitative testing (Bueno de Mesquita and Fowler, 2021; Granato, Lo and Wong, 2021). Using formal political theory to guide our data collection and analytical methods is an important contribution of our research, which we hope provides direction for ways that social media data and advanced quantitative methods can be used to test political theories.

Supplementary material. To view supplementary material for this article, please visit <https://doi.org/10.1017/S1755773924000146>.

Author contributions. All authors contributed to the production of this paper.

References

- Aldrich, John, and David Rohde. *Measuring Conditional Party Government*. Working paper. Duke University, 1998. <https://www.researchgate.net/publication/251842516_Measuring_Conditional_Party_Government>
- Aldrich, John and David Rohde. The logic of conditional party government. In *Congress Reconsidered*, Ed. L. C. Dodd and B. I. Oppenheimer. Washington, DC: CQ Press, pp. 269–92, 2001.
- Barbera, Pablo, Andreau Casas, Jonathan Nagler, Patrick J. Egan, Richard Bonneau, John T. Jost, and Joshua A. Tucker. "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 (2019): 883–901.

- Bendix, William.** "Bypassing Congressional Committees: Parties, Panel Rosters, and Deliberative Processes." *Legislative Studies Quarterly* 41.3 (2016): 687–714.
- Bueno de Mesquita, Ethan and Anthony Fowler.** *Thinking Clearly with Data: A Guide to Quantitative Reasoning and Analysis*. Princeton University Press, 2021.
- Curry, James M.** *Legislating in the Dark: Information and Power in the House of Representatives*. Chicago: The University of Chicago Press, 2015.
- Dewan, Torun and David P. Myatt.** "Leading the Party: Coordination, Direction, and Communication." *American Political Science Review* 101.4 (2007): 827–45.
- Fenno, R.F.** *Home Style: House Members in Their Districts*. Classics Series Longman, 2003. <<https://books.google.com/books?id=p4UrAQAAMAAJ>>.
- Gamm, Gerald and Steven Smith.** The Dynamics of Party Government in Congress. In *Congress Reconsidered, 11th Edition*, Ed. L. C. Dodd and B. I. Oppenheimer. Washington, DC: CQ Press, 2020, pp. 197–224.
- Granato, Jim, Melody Lo, and M. C. Sunny Wong.** Empirical Implications of Theoretical Models in Political Science. Cambridge University Press, 2021.
- Hall, Thad E., and Betsy Sinclair.** *A Connected America: Politics in the Era of Social Media*. Oxford University Press, 2018.
- Harbridge, Laurel.** *Is Bipartisanship Dead? Policy Agreement and Agenda-Setting in the House of Representatives*. New York: Cambridge University Press, 2015.
- Hemphill, Libby, Jahna Otterbacher, and Matthew Shapiro.** "What's Congress Doing on Twitter?" In *Proceedings of the 2013 Conference on Computer Supported Cooperative Work*. ACM. pp. 877–86, 2013.
- Howard, Nicholas O., and Mark E. Owens.** "Circumventing Legislative Committees: The US Senate." *Legislative Studies Quarterly* 45.3 (2020): 495–526.
- Kang, T., E.F. Fowler, M.M. Franz, and T.N. Ridout.** "Issue Consistency? Comparing Television Advertising, Tweets, and E-mail in the 2014 Senate Campaigns." *Political Communication* 35 (2018): 32–49.
- Kim, In Song, John Londregan, and Marc Ratkovic.** "Estimating Spatial Preferences from Votes and Text." *Political Analysis* 26.2 (2018): 210–29.
- Koger, Gregory, and Matthew J. Lebo.** *Strategic Party Government: Why Winning Trumps Ideology*. University of Chicago Press, 2020.
- Koop, Gary, M. Hashem Pesaran, and Simon M. Potter.** "Impulse Response Analysis in Nonlinear Multivariate Models." *Journal of Econometrics* 74.1 (1996): 119–47.
- Lin, C., Y. He, R. Everson and S. Ruger.** "Weakly supervised joint sentiment-topic detection from text." *IEEE Transactions on Knowledge and Data Engineering* 24.6 (2012): 1134–45.
- Lin, Chenghua, and Yulan He.** "Joint Sentiment/Topic Model for Sentiment Analysis." In *Proceedings of ACM International Conference on Information and Knowledge Management (CIKM-2009)*, 2009.
- Peng, Tai-Quan, Mengchen Liu, Yingcai Wu, and Shixia Liu.** "Follower-Followee Network, Communication Networks, and Vote Agreement of the US Members of Congress." *Communication Research* 43.7 (2016): 996–1024.
- Roberts, Margaret E., Brandon M. Stewart, Dustin Tingley, Christopher Lucas, Jetson Leder-Luis, Shana Kushner Gadarian, Bethany Albertson, and David G. Rand.** "Structural Topic Models for Open-Ended Survey Responses." *American Journal of Political Science* 58.4 (2014): 1064–82.
- Röder, Michael, Andreas Both, and Alexander Hinneburg.** "Exploring the Space of Topic Coherence Measures." In Proceedings of the Eighth ACM International Conference on Web Search and Data Mining. WSDM '15, New York, NY, USA: Association for Computing Machinery, pp. 399–408, 2015. <https://doi.org/10.1145/2684822.2685324>
- Tiefer, Charles.** *The Polarized Congress: The Post-Traditional Procedure of Its Current Struggles*. Lanham, MD: University Press of America, 2016.
- Vaccari, Cristian, and Rasmus Kleis Nielsen.** "What Drives Politicians' Online Popularity? An Analysis of the 2010 US Midterm Elections." *Journal of Information Technology & Politics* 10.2 (2013): 208–22.
- Wallis, Kenneth F.** "Time Series Analysis of Bounded Economic Variables." *Journal of Time Series Analysis* 8.1 (1987):115–23.
- Wallner, James I.** *The Death of Deliberation: Partisanship and Polarization in the United States Senate*. New York: Lexington Books, 2013.
- Yan, Hao, Sanmay Das, Allen Lavoie, Sirui Li, and Betsy Sinclair.** "The Congressional Classification Challenge: Domain Specificity and Partisan Intensity". In *Proceedings of the 2019 ACM Conference on Economics and Computation. EC '19*. New York, NY, USA: ACM, pp. 71–89, 2019.