Legislative Communication and Power: Measuring Leadership in the U.S. House of Representatives from Social Media Data

Daniel Ebanks, Hao Yan, R. Michael Alvarez, Sanmay Das, Betsy Sinclair*

March 5, 2023

Abstract

Who leads and who follows in Congress? We analyze the Twitter accounts of U.S. House of Representative members to develop a new understanding of House leadership power. Formal theoretic work on congressional leadership hypothesizes that party members shift their policy stances as they balance coordination and information problems. The theoretical work predicts that when coordination problems are pressing, legislative members follow the policy positions of party leaders. When their party's information problem is more acute, party members instead give their leaders direction for the party's agenda. Our empirical study uses the Joint Sentiment Topic model and tests these implications of this theoretical model using our Twitter data. Our analyses reveal that leader-follower relationships are complex. Party leaders possess the power to substantially affect the propensity of rank-and-file members to discuss topics, especially when the coordination problem dominates; however, when rank-and-file members influence discussion of a topic, their effect on leadership's propensity to discuss it is far larger. These effects are particularly pronounced even when coordination problems are pressing. However, when the information problem is more acute, leadership influence decreases, consistent with

theory. We show these results are robust to the the underlying dynamics of contemporary political discussion and context.

1 Introduction

This paper uses data from the Twitter accounts of U.S. House members to study legislative communication and leadership influence. Previous studies have focused on traditional means by which party leadership exerts power over the members, such as through setting the legislative agenda, through committee assignments, or through the ability to influence members to adopt leadership talking points with formal press releases. Our paper takes the novel approach of examining how leaders exert influence over the policy stances that members of Congress communicate on social media. To guide the analysis, we have interpreted conventional theories of congressional leadership and communication in the context of social media. Our analysis is facilitated by the signalling and coordination model of Dewan and Myatt (2007). In analyzing the influence of party leadership over their members on social media, the relevant insight from the model is 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 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 key hypothesis which 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. To test this hypothesis, we construct a distinct measure of House member rank-and-file behavior from Twitter data – daily rank-and-file positions on issues that are being

discussed on social media.

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. 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 also 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).

We build on these previous studies by examining congressional leadership in a communication setting, particularly over social media. We note two key distinguishing features of our analysis. First, we are able to avoid the 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 which lack majority support, 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). Second, the high frequency nature of social media data allow 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 party leadership's ability to direct legislative communication and public engagement around specific

topics *among* their members.

Our paper contributes to these areas in three 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 stances (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 stances. 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 stances 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 stance messaging under certain conditions. Our results demonstrate that polarization alone is not sufficient to explain patterns of party leadership in the House.

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 stance while trying to communicate the best policy stance 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

---

[1] Twitter 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; Kang et al., 2018)).

3

and 116th Congresses, between January 1st, 2017 and January 3, 2021. After pre-processing these data, we use weakly supervised machine learning methods to 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, which we term "need for direction."

2   Mechanisms for Leadership Communication and Influence

Our empirical analysis is framed around a key theoretical insight from the Dewan and Myatt (2007) signalling and coordination game of party leadership and communication – where leadership facilitates coordination on a policy stance in response to uncertain issues. In the context of this framework, uncertainty could be the political and electoral popularity of taking a policy stance 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 policy stance 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 policy stance 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 policy stance in an uncertain state of the world.

To gain intuitive insight into the setting, we continue to discuss the 2019 government shutdown debate in more detail. 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 progressives in the Democratic Party 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 are not able to directly measure this quantity. In the model, the party will select one policy stance whose number of supporters is greater than some threshold. In our example, this might be House Speaker Pelosi's internal sense of the level of party support she needs in order to pursue a particular messaging strategy. In the case where neither policy stance 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. We might imagine she gauged internal support as sufficiently high for this strategy. This leads us to 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 our example from the 2019 government shutdown, some Democratic members

publicly indicated to the press 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 in fact 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 policy stance for each party member is communicated publicly on Twitter. To evaluate the ability of the party to coordinate around the leaders' stances, 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 gravity of the party coordinating around the "correct" policy stance – it is a multiplier on the payoff of coordinating on the "correct" policy stance. When need for direction is high, the information problem tends to dominate. This is because the merits of the policy stance are especially responsive to underlying fundamentals which are uncertain.

We analyze our data at the individual sentiment-topic level. On issues where the party's need for direction is low, we expect House rank-and-file to adopt the policy stances of their leaders. Here, the stakes for choosing the wrong policy stance are relatively low, and members prefer to coordinate around a unified policy – even if it is "incorrect" – rather than fail to coordinate at

---

[2] Readers interested in details of the theory can see (Dewan and Myatt, 2007). In the Supplementary Information we present game details and details about how the theoretical concepts translate into empirical measures. Below, in Table 1 we connect the key theoretical concepts to their empirical analogues.

all. For issues where need for direction is high, we expect House leaders to adopt the policy stances of their rank-and-file. We define issues with low need for direction as those which drive the partisan divide between parties, 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 electoral payoff or cost in taking these stances. Conversely, need for direction is high when coordinating on the "correct" stance has out-sized electoral and policy effects, such as a government shutdown. Government shutdowns have resulted in severe policy and electoral consequences. Here, we expect House leadership influence to be weaker, as the theory suggests that rank-and-file members will hedge against the leaders and adopt their private stance publicly, as the consequences for coordinating on the "wrong" message are out-sized.

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 analogues. We present the results in Section 5, with the discussion and conclusion in Section 6.

Concept	Revealed By	Empirical Analogue
Need for Direction	Sentiment-topics with out-sized benefit or cost of coordinating	Classify top twenty topics for each party driving separation in sentiment-topic space as uncovered by PCA analysis as needing direction
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: Terminology

3 Data and Methodology

3.1 Data

In order to study the dynamics of communication, we exploit legislators' use of Twitter, which provides us with high-frequency text data. Using this granular data, we examine whether the House party rank and file anticipate 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. These dates coincide exactly with the 115th and 116th sessions of Congress. The list of accounts were collected based on the official Twitter handles list collected by C-SPAN[3], following Barbera et al. (2019) who used the NYT Congress API 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 stances. Personal and electoral Twitter accounts often focus on non-policy issues, like personal family matters or scheduling of specific campaign events (such as local town halls or rallies). They also tend to highlight the personal interests of the member of Congress, such as sports events or television programs. Our dataset includes 738,066 tweets, including only original posts. Table SI 1 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.[4]

---

[3] https://twitter.com/cspan/lists/members-of-congress/members

[4] See Figure SI 1 for the overall distribution of tweets by House members for this period

8

3.2 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 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) to estimate the latent structure in the sentiment-topic results. Finally we uses these estimates in a time-series analysis to test whether House leaders initiate the messaging regarding a policy stance or whether House party rank-and-file initiate discussion. These three steps are discussed in this section.

3.2.1 Joint Sentiment Topic Analysis

We employ a method of estimating both a topic mixture and sentiment mixture which we believe is new to political science and the study of legislative communication and behavior, the Joint Sentiment Topic (JST) model. It is based on LDA, though it estimates a conditional mixture for topics k given sentiment j. However, unlike LDA (which estimates two latent layers, topic classification and words alone), the JST estimates three latent layers (sentiment orientation, then topic classification, then word mixtures). Importantly, the JST model estimates the unconditional probability of each sentiment j. 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 full discussion space, but we do not rely on assumptions regarding exogenous covariates to uncover the latent space.

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, even when projected into a lower dimensional space. By uncovering this inter- and intra-party variation, we are able to analyze behavior within and across parties. Without variation within party, we would not be able to analyze the parties' respective communications over time. Without variation across party, we would not be able to compare the communications between party. Moreover, since this method uncovers clear partisan separation in party communication that suggests the unsupervised method has external validity, as we reveal the partisan nature of discussion on social media from the patterns of communication.

For each of the 738,066 tweets in the dataset, we produce 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 sentiment-topic probabilities and j sentiment probabilities for each tweet.[5]

Importantly, as we connect the JST model to political contexts, the model relies on exchangeability and is a bag-of-words approach to speech. That is, the order of the words in the document is not considered as sentiment-topics are uncovered. Although this is a simplistic model of speech, these assumptions allow for a tractable estimation of the topics at hand, with little cost to the coherence of the uncovered sentiment-topics. Explicitly, the underlying data-generation process for the documents is summarized as follows:

---

[5] Note that the sentiment-topic labels are independent, so that Sentiment1-Topic 3 has no relation to neither Sentiment 2-Topic 3 nor Sentiment 3-Topic 3.

1. For each tweet t, choose a distribution $\pi_t \sim \text{Dirichlet}(\gamma)$. Here, π_t is a multinomial distribution over sentiments for each document drawn from a Dirichlet prior.

2. For each sentiment label j under tweet t, choose a topic distribution $\theta \sim \text{Dirichlet}(\alpha)$. Here, θ is a multinomial distribution over topics for each tweet conditional a sentiment. This distribution is drawn from a Dirichlet prior.

3. For each word w_i in tweet t,

   (a) Choose a sentiment label j_i from π_t.

   (b) Choose a topic label k_i from θ_{t,l_i}.

   (c) Choose a word w_i from the distribution, ϕ_{j_i,k_i} over words defined by the topic k_i and sentiment j_i. Here, ϕ_{j_i,k_i} is a distribution over words given being in sentiment label j_i and topic label k_i under sentiment j_i.

This is a Bayesian hierarchical mixture model. In this study we consider the prior parameters α as the prior concentration of the sentiment-topic k_i for a document before having seen any documents. Similarly, β can be interpreted as the prior concentration of the sentiment-topic j for a word before having observed any words. Finally, λ can be interpreted as the prior concentration of sentiment labels sampled under a document before having observed any documents.

Observe that as β goes to 0, the model converges to a model of a single sentiment-topic. That is, one sentiment-topic label has probability 1, with all other labels being assigned 0. On the other hand, as β grows large, the limiting distribution is uniform over sentiment-topics. We expect that tweets, given their concise nature, are likely only to relate to very few topics at once, so we set these priors relatively small, following standard practice (such as in Lin and He (2009)). We provide a full technical overview in SI Section 4.1.[6]

---

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

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 max 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 most emblematic tweets for each sentiment-topic. These are the tweets with the highest probability of belonging to their sentiment-topic label. We report the stripped down tweet (which is the raw data) 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]

3.2.2 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.

---

[7]For additional details, see Supplementary Information 4.1. JST estimates two layers in addition to the word layer, so it is not necessarily the case that these two layers are indeed sentiment and topic. Nonetheless, our findings show that the second layer uncovers meaningful partisan separation in the data, as seen in Figure 1.
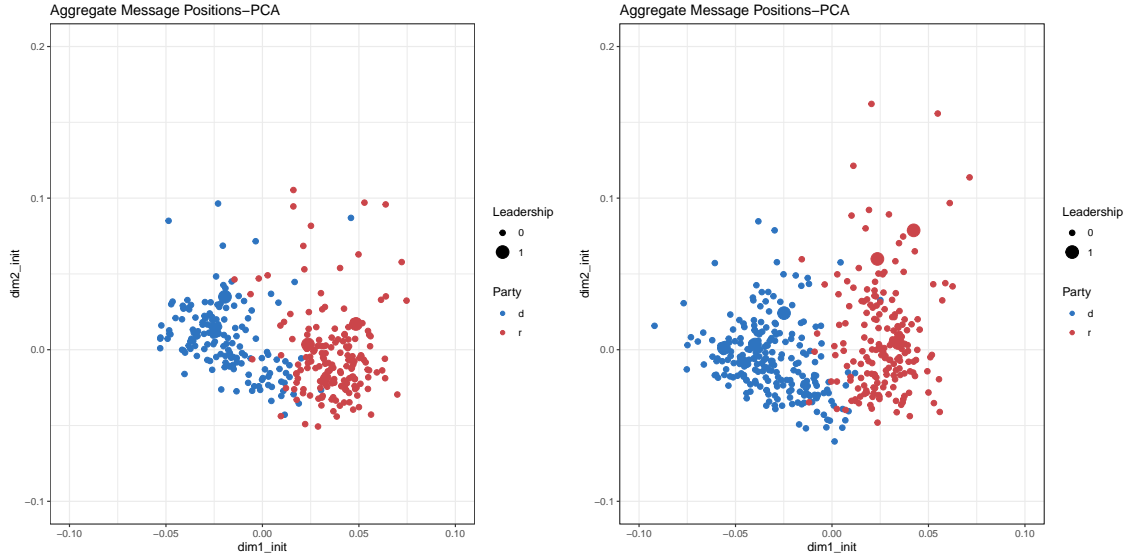
Figure 1: Policy Stance Positioning

Figure 1: Aggregated legislator policy stance 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.

The processes for determining communications decisions are likely driven by exogenous events, party and peer effects, and personal preferences of legislators, which are not immediately obvious from looking at the raw mixtures. By using PCA as a dimension reduction technique, we compactly capture this latent structure in the data. We are then able to compute an intuitive measure of the variance of the topics being discussed within each party by computing the variance of the policy stance scores. Figure 1 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 stance for each legislator.[8]

To compute these scores, we employ the PCA in the following fashion to compute a "policy

---

[8] We also estimate this measure restricted to just the respective Congress. Figure SI 4 shows the contrast of rank-and-file members' position in the PCA-derived policy stance space when we estimate it separately. We show the main result is robust to changes in this estimation routine.

stance" score for each legislator. After computing the JST mixtures for each tweet, we find the average probability a House member tweeted about a particular sentiment-topic k and sentiment j by taking the rolling average over an 8-week period. We choose this time window because it ensures that every legislator has on average 50 tweets in the given time frame, which we assume is a sufficient amount of data to identify the true sentiment-topic distribution of a given legislator's communication strategy. We drop any legislator with less than 10 tweets over this period, usually about 3 legislators out of 435 per period.[9]

We emphasize that our policy stance scores measure a position in sentiment-topic space over popular debates taking place on social media in real time. PCA analysis allows us to analyze public policy stances 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 "policy stance" space for each time period. From these individual-level measures of communication, we can compute party-level measures of messaging focus, which form the basis of our empirical tests of the hypothesis regarding party leaders' efficacy in coordinating the party around key policy stances.

3.2.3 Dynamic Analysis

Finally, we exploit the micro-level data to examine whether House leaders initiate discussion on Twitter within their party coalition (and thus exert influence over their rank-and-file), or whether they adopt their members' consensus, helping to create a focal point around which to coordinate. 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, with some key modifications. First, we measure daily propensity

---

[9] We show that in SI Figure SI 3 that the final results are invariant to the cutoff choice.

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.

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 vector autoregression (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 reason: first, because the theory makes predictions over which types of topics should facilitate the emergence of leadership within individual parties, we estimate VAR's separately for each topic and party to evaluate the extent that party leaders emerge as theory predicts. Second, given the large number of sentiment-topics (84 total) and the fact we are looking at leaders and non-leaders, the parameter space is large. Thus, the system of equations may not be identified for a reasonable number of lags. Although assuming the topics are not directly related is a strong assumption, it allows us to identify more lags and improves computational tractability. It also avoids introducing potentially many spurious correlations, given the highly interrelated nature of the data. Finally, in cases where the nature of the structural relationships are not known to the researcher, interpreting the results from a VAR regression is difficult. Our parsimonious specification allows for a more direct examination of whether leaders lead or follow.

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

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

Our specification thus is:

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

Here c is a constant accounting for the fact the time series are stationary around a non-zero mean after taking logs. Appendix Figures SI 7 and SI 8 show for selected series that the times 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 we reject at the 1 percent 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 Vector Auto Regression analysis, so these dynamic properties of the data in question suggest our data are consistent with the the key prerequisite assumptions of the methods we employ. Finally, we choose a lag of 2 days, which captures the length of the news cycle on Twitter.[10]

Finally, to capture the extent that House leaders lead, or followers initiate, discussion, we estimate generalized impulse response functions for each specification following Koop, Pesaran and Potter (1996).[11] That is, we measure the effect of a two standard deviation increase in a

[10] We 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 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.

[11] Generalized 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,

16

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-percent bootstrapped confidence intervals with 500 draws.

The final step is our determination of the topics on which House leaders lead versus those on which followers lead. We employ Granger tests to determine those topics for which leaders' daily propensity to discuss a sentiment-topic precipitates their members' daily propensity as well as those for which the reverse holds. As we state in Table 1, the Granger tests measure our first notion of leadership, which is purely temporal– that leaders precipitate their rank-and-file members' messaging strategies. Then, as we defined in Table 1, we compare the initial IRF responses from leaders to members and members to leaders on the topics where leaders are predicted to influence discussion. This measures a second notion of leadership – the ability of leaders to alter discussion.

4   Operationalizing 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. Importantly, to test our hypothesis related to need for direction, we look at a topic-by-topic analysis of the data using temporal notions of House leadership. See Table 1 for a road map to our analyses.

---

we compute $\Theta_i^k(n) = \frac{\delta_j}{\sigma_j^2}\Sigma_\epsilon\beta$ where δ is two standard deviations of our data, approximately 10 percent.

4.1 Need for Direction

Our key hypothesis necessitates a topic-by-topic level analysis. To test the hypothesis that House leaders initiate discussion when the need for policy direction is low (and the coordination problem dominates), we first need to uncover when leaders initiate discussion and when rank-and-file members influence discussion.

Coordination Problem

In Tables 2 and 3, we show the sentiment-topics which we use to define those issues in low need of direction – that is, where the coordination problem dominates. To classify sentiment-topics which need little direction, we take the top twenty sentiment-topics for each party which drive partisan separation in sentiment-topic space as measured by the principal components of the propensity to discuss sentiment-topics.

Table 2: PCA 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 Seperations-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 Asuylum Policy	1.66
Prevent Gun Violence-Negative	1.57
Abortion Rights-Negative	1.53
Manfacturing Jobs - Neutral	1.51
DACA Policy - Positive	1.49
Trump/Russia Investigation -Negative	1.39

Table 3: PCA 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
GOP 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 Justive 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

Our criterion for determining whether each topic needs direction is based on the percent contribution to the variation of the top two components derived from the principal components analysis. We take the top twenty topics that contribute to each party's half of the sentiment-topic space, and classify those topics as being low in need for direction. Sentiment-topics with low contribution to the partisan space do not drive legislators toward the extremes of policy sentiment-topic space, and large contributions drive them to the extreme portion of the space. As we can see in Figure 1, policy stances 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 stances which define 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 next 20 sentiment-topics as in high need of direction. These topics tend to have political significance, but do not contribute to partisan separation. We argue these sentiment-topics represent topics where the underlying political fundamentals of the topics are more uncertain, so the information aggregation problem dominates for the party. The parties are incentivized to find the "correct" message, and in this case, failure to coordinate would be preferable to coalescing around the wrong message.

Table 4: PCA Topic Contributions - Member Driven 115th

Topic	Contribution
Lowest Unemployment Rate - Positive	1.26
Guests at Capitol Hill-Neutral	1.24
LGBT Equality-Negative	1.23
Fight for Civil Rights-Negative	1.15
Retweeting a Controversial Statement-Negative	1.11
Climate Change-Positive	1.11
Partisan Attacks on Trump/Biden-Negative	1.01
Important Meetings-Negative	0.99
Trump Admin Undermines Country - Negative	0.97
Budgetary Legislation -Negative	0.86
Committee Hearings-Positive	0.86
Hurricane Relief-Negative	0.86
Trump Climate Policy-Negative	0.85
Health Care Expansion - Neutral	0.84
Foreign Election Interference-Negative	0.82
Women's Pay - Positive	0.82
Supreme Court Nominations-Negative	0.73
Thoughts and Prayers - Negative	0.71
Floor Speeches-Negative	0.68
Student Loan Relief-Positive	0.67

Given these topics carry additional outsized risks to coordinating on the wrong message, we classify them as needing direction. For example, LGBT equality was a topic that was relatively less divisive during the Trump administration. For Republicans, the optimal LGBT policy was

Table 5: PCA Topic Contributions - 116th Member Driven

Topic	Contribution
Family Seperations-Negative	1.18
Pro-Life Policy - Negative	1.14
China/Hong Kong Protests-Negative	1.11
Republican Senate Legislation-Negative	1.10
Prevent Gun Violence-Negative	1.08
Trump Admin Undermines Country - Negative	0.96
Fight for Civil Rights-Negative	0.93
Meuller Investigation - Negative	0.88
Trump Asuylum Policy	0.86
Enjoyable Visit - Positive	0.74
LGBT Equality-Negative	0.71
Social Security/Postal Service - Neutral	0.70
Health Care Expansion - Neutral	0.70
Trump Climate Policy-Negative	0.68
Mitch Mcconnel's Senate-Negative	0.67
Partisan Votes - Negative	0.60
Voting Rights - Positive	0.56
Law Enforcement - Positive	0.55
Honoring Cultural History-Negative	0.55
Protect Health Insurance -Neutral	0.54

entirely unclear because of the rapidly shifting political landscape related to LGBT rights. By the 115th and 116th Congresses, the 'correct" conservative consensus on LGBT rights was unclear, with marriage rights for LGBT Americans experiencing strong public support in public opinion surveys and the GOP political base raising the salience of other issues, such as immigration and border policy.

4.2 House Leadership Influence

We employ a dynamic notion of House leadership influence. To test the main hypothesis we exploit temporal dynamics at the individual sentiment-topic level to derive our empirical notion of leadership. We analyze the impulse responses for these same sentiment-topics. We look at the first period shocks under the IRFs to find topics where leaders initiate discussion and exert quantifiable

influence over the rank-and-file members of their party.

To quantify influence, we employ IRF analyses from a vector-autoregression following a similar strategy to Barbera et al. (2019). Here, we try to quantify the ability of House leaders to drive discussion. We take the average daily propensity to discuss a sentiment-topic by party leadership and by party rank-and-file. The IRF analysis supposes a 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.

5   Results

5.1   Need for Direction by Leadership - Coordination Problem

We expect that topics where the need for policy direction is low will be the topics where House leaders influence discussion for the rank-and-file. By high need for direction, we mean that the electoral, political, or policy costs of coordinating on the "wrong" policy stance are large. In the case of coordinating on the wrong stance, it might be more advantageous to avoid coordination at all. For example, coordinating on the wrong stance could exacerbate a government shutdown crisis, leading to electoral defeat for the party, or catastrophic policy consequences for the public, or both. To test the need for direction hypothesis in the case when the coordination problem dominates, we turn to the micro-level propensities of party leaders and rank-and-file members to discuss each sentiment-topic daily. We find evidence consistent with the theory. The evidence from the IRF analysis provides strong support for the need for direction hypothesis. The IRF analysis suggests leaders can increase the rank-and-file's propensity to discuss these most partisan topics by between 0.1 and 1 percent for each standard deviation increase in the leadership's daily propensity

to discuss a topic. This result is consistent across parties and time periods, even when the party in power changes. We argue that this consistency is evidence that the result is robust across these same dimensions, during the period of 2017 to 2021.

Figure 2: Democratic Topics: Need for Direction
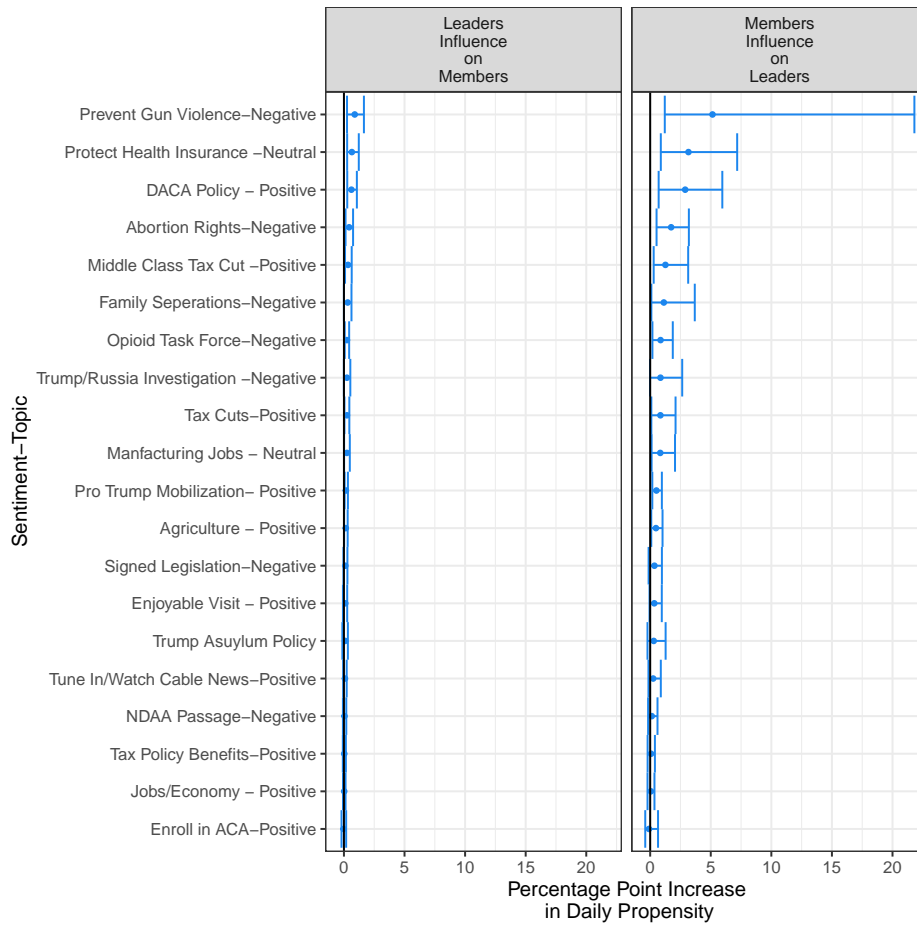Predicted Leader Driven 115th Congress

Figure 2 : Impulse response functions for sentiment-topics predicted to be leader driven for the Democratic Party. Bootstrapped 95-percent confidence intervals are shown.

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

23

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. The costs of coordinating on a sub-optimal message are outweighed by the the costs of coordinating on a "wrong" message. Thus, the coordination problem dominates.

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 GOP benefited politically from raising the salience of the economy. Thus, mis-calibrating the message on the positive economy was less costly than not coordinating at all on a message.

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 for 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 the Impeachment). Second, they discussed topics that are central to the Democratic Party's ideology, such as racial equality and public health. The costs of coordinating on a sub-optimal message are outweighed by the the costs of coordinating on a "wrong" message or no message at all. Thus, the coordination problem dominates.

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 percent increase in the rank-and-file members' daily propensity to discuss that issue. In particular, Republican leaders induced a $\sim$ 1 percentage point increase in their rank-and-file members' propensity to discuss impeachment and 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

24

low unemployment rate. Figure 5 also shows that members induced a ~ 2 percentage point increase in their leadership's propensity to discuss impeachment and humantiarian aid at the border. Members exerted a ~ 1 percentage point increase in their leaders' propensity to discuss crimes at the border and attacking the Democratic Party as socialist. 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.

5.2   Need for Direction by Membership - Information Problem

Next, we examine in close detail the behavior of congressional parties for topics where the information aggregation problem dominates political concerns. 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 for 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

25

Figure 3: Republican Topics: Need for Direction
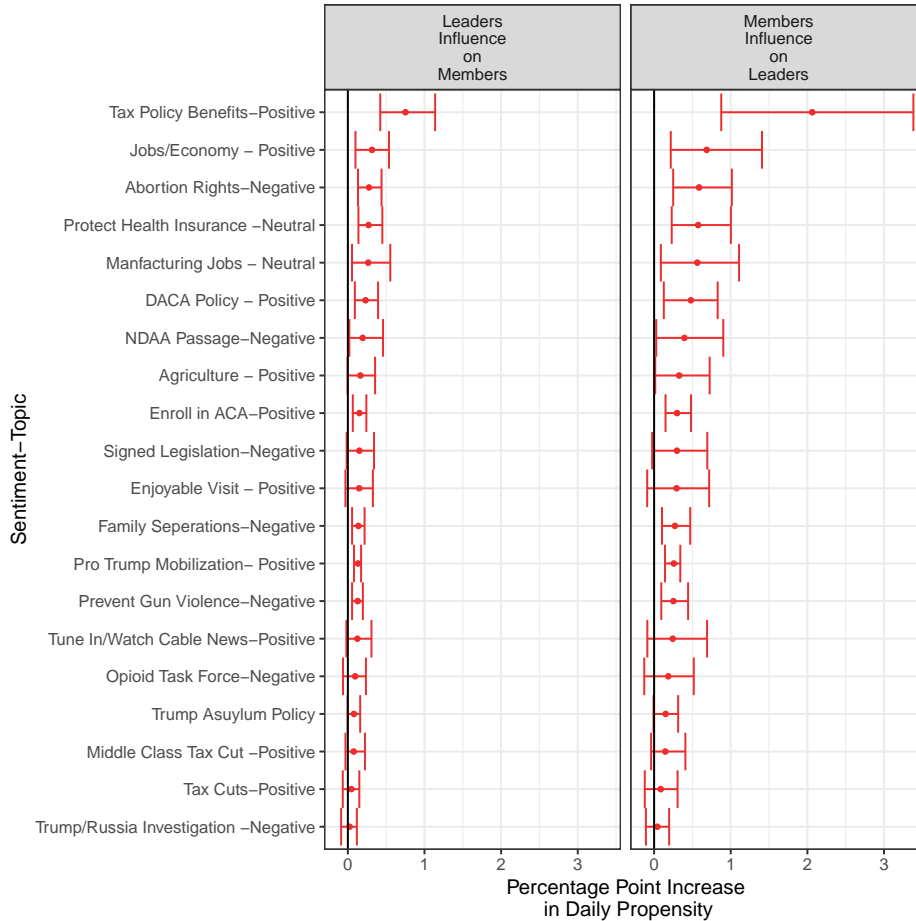Predicted Leader Driven 115th Congress

Figure 3: Impulse response functions for sentiment-topics predicted to be leader driven for the Republican Party. Bootstrapped 95-percent confidence intervals are shown.

on "the wisdom of the crowd" of the party at large. By aggregating all of their information, the party hopes to coordinate on the "correct" message. In these cases, the costs of coordinating on the wrong message outweigh the costs of failing to coordinate. Sometimes, it is more advantageous to the party to have no message at all than to take risky stances. In these cases, parties will seek to find a party consensus.

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

Figure 4: Democratic Topics: Need for Direction
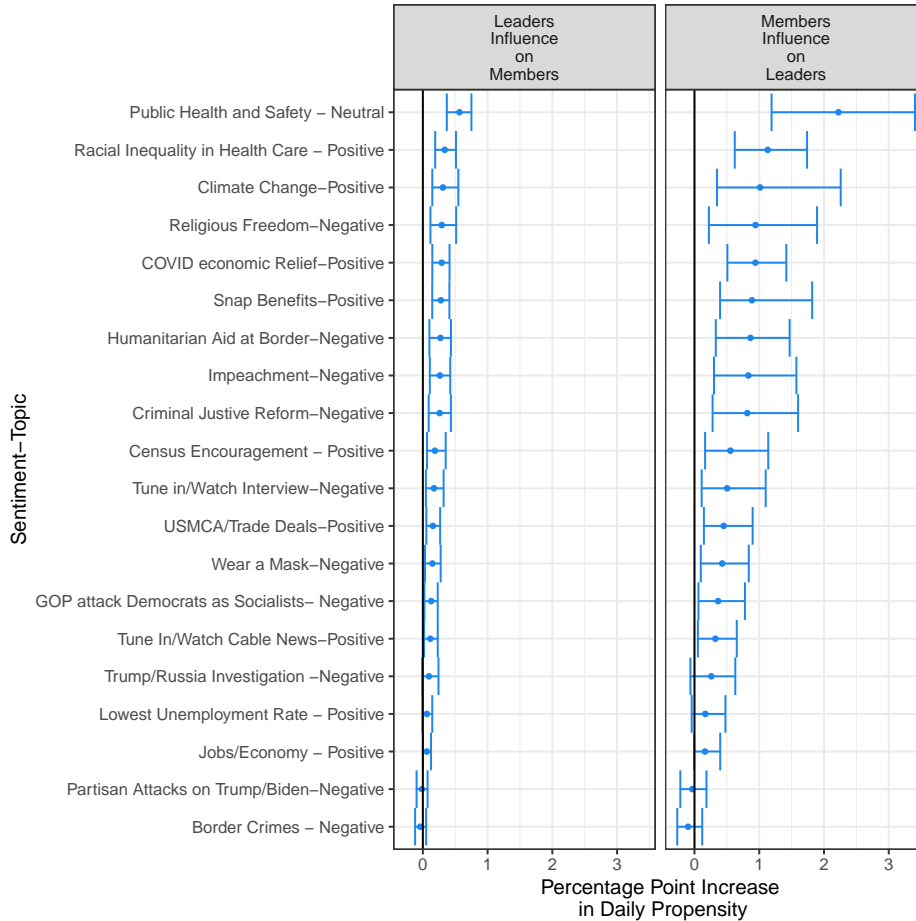Predicted Leader Driven 116th Congress

Figure 4 : Impulse response functions for sentiment-topics predicted to be leader driven for the Democratic Party. Bootstrapped 95-percent confidence intervals are shown.

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 GOP 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

Figure 5: Republican Topics: Need for Direction
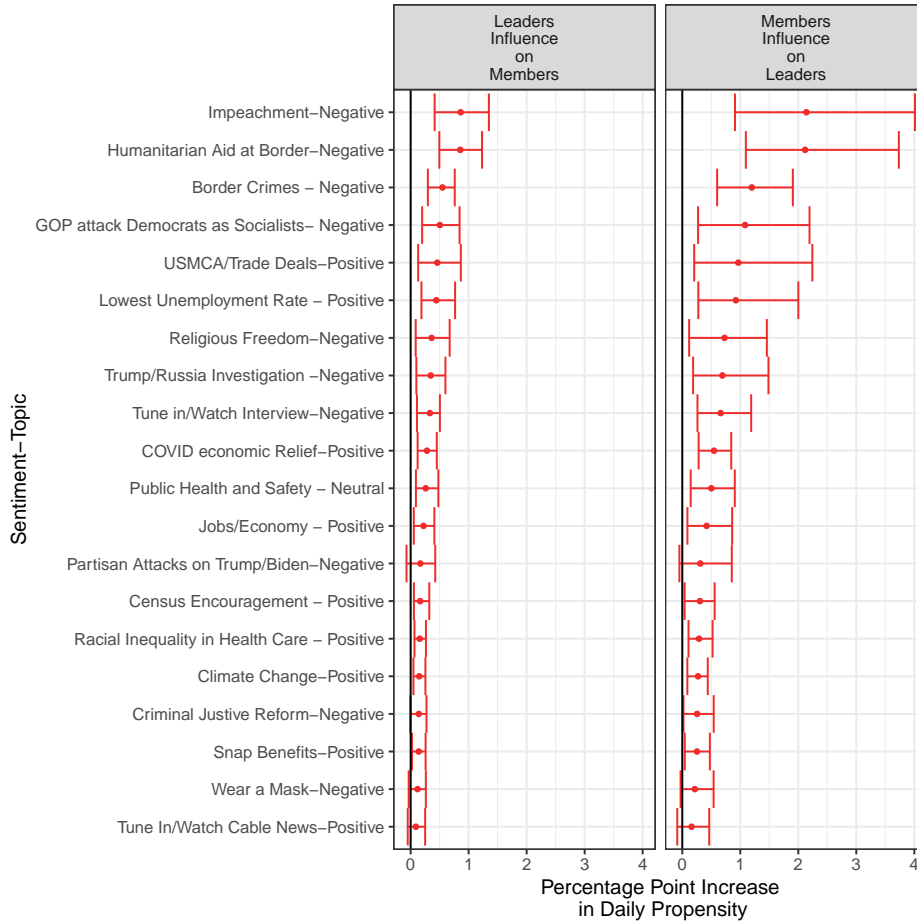Predicted Leader Driven 116th Congress

Figure 5: Impulse response functions for sentiment-topics predicted to be leader driven for the Republican Party. Bootstrapped 95-percent confidence intervals are shown.

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 8 rank-and-file members exert a ~ 1 percent 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 majority parties find it advantageous to engage in measurable amounts of political deflection.

Figure 6: Democratic Topics: Need for Direction
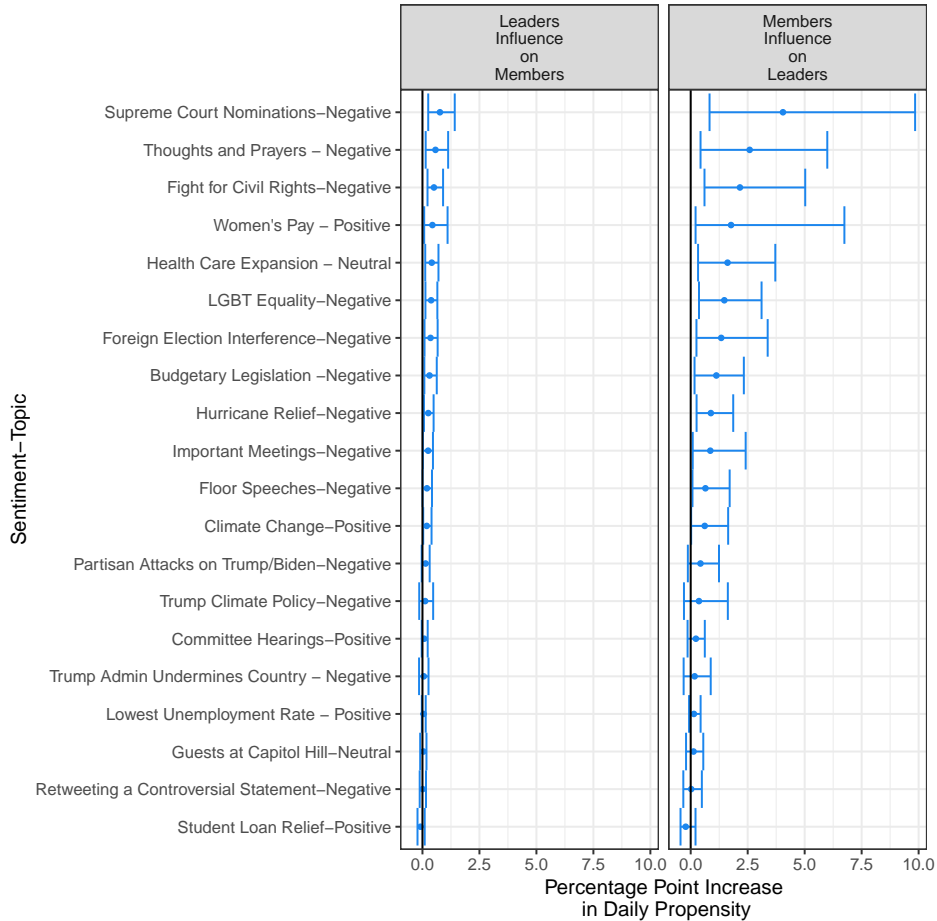Predicted Member Driven 115th Congress

Figure 6 : Impulse response functions for sentiment-topics predicted to be leader driven for the Democratic Party. Bootstrapped 95-percent confidence intervals are shown.

The results for the 116th Congress follow a similar pattern for both parties. Figure 7 shows that the Democratic rank-and-file membership exerts a 2 to 3 percent 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, the party communication behavior on social media is consistent as 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 consistent with

Figure 7: Democratic Topics: Need for Direction
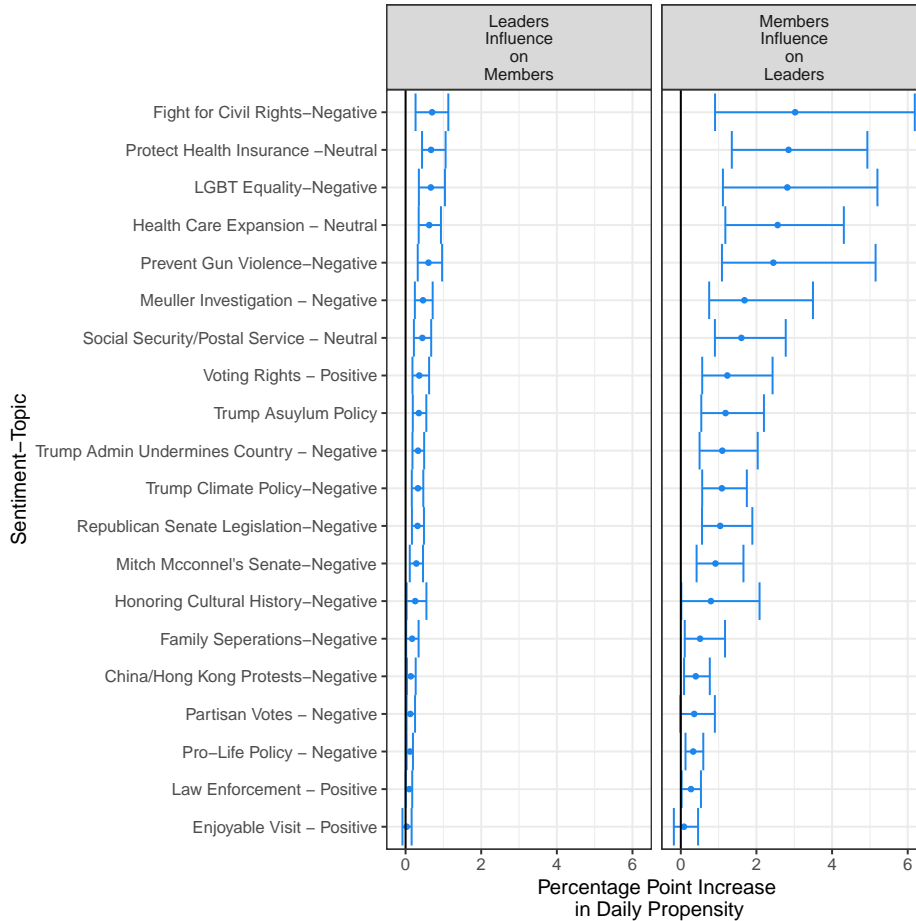Predicted Member Driven 116th Congress

Figure 7 : Impulse response functions for sentiment-topics predicted to be leader driven for the Democratic Party. Bootstrapped 95-percent confidence intervals are shown.

the obfuscation tactics we saw among the GOP rank-and-file when they were in the majority in the 115th Congress. This is consistent with 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 Party rank-and-file exhibits behavior consistent with their behavior in the 115th and consistent with their contemporaneous colleagues in the Democratic Party 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 a approximately 0.5 to

1 percent increase in the rank-and-file members' daily 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. Notable, the magnitudes derived for Republicans 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, 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. We hypothesize these last two facets of communication behavior are related to the interactions of the presidency.

5.3 Discussion

Although these point estimates may seem substantively small, in fact, shocks of 3 or 4 standard deviations (40 to 60 percent) on the daily propensity to discuss a topic are common. This reflects the nature of conversation on Twitter, which tends to react to the daily news cycle.

We highlight the consistency of these findings across the parties: on issues where House rank-and-file influence discussion, their effect on leaders is larger in 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. Substantively, these observations speak to the nature of discussion on social media, where exogenous events can drive conversation – they also highlight that leaders are more sensitive to changes in the topics they discuss than the average member of the party.

Finally, in Table SI 3 we show that institutional 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 finding highlights the relative strength

31

Figure 8: Republican Topics: Need for Direction
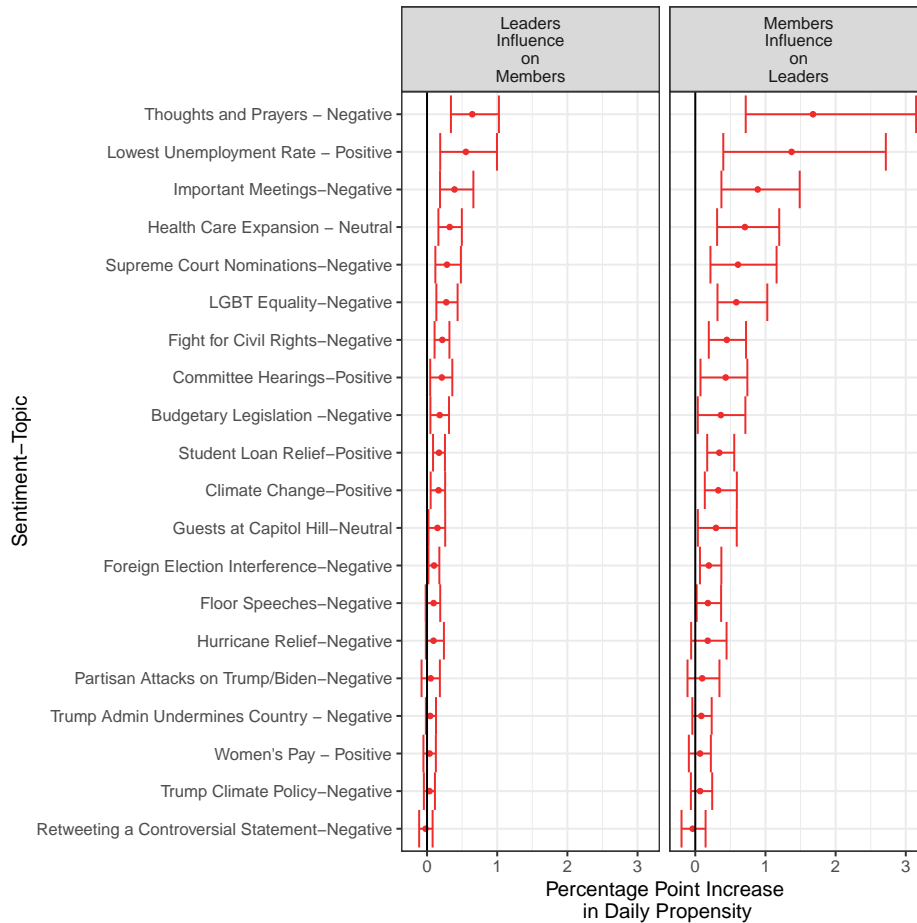Predicted Member Driven 115th Congress

Figure 8: Impulse response functions for sentiment-topics predicted to be leader driven for the Republican Party. Bootstrapped 95-percent confidence intervals are shown.

of institutional leadership within the party caucus relative to the influence of members of the party who are popular with the public social media [12].

---

[12] We also show in Table SI 3 that Congressional 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 any particular secular trend in the discussion on social media. Instead, this result suggests something particular about the institutional role of Congressional

Figure 9: Republican Topics: Need for Direction
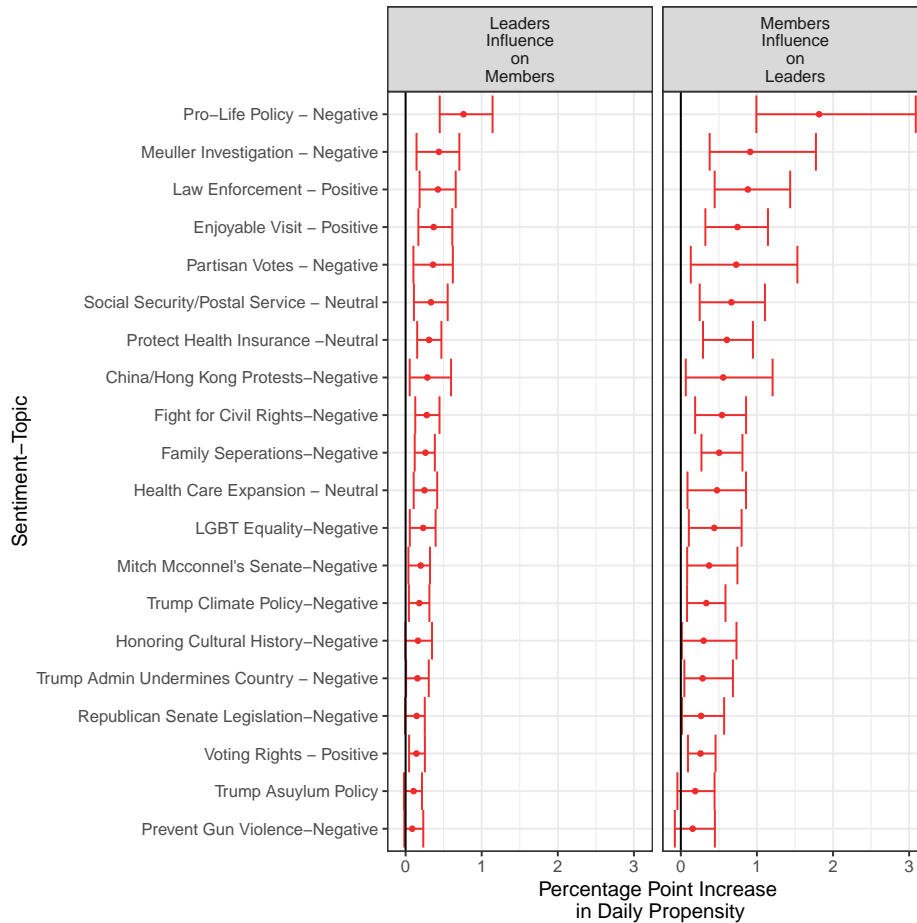Predicted Member Driven 116th Congress

Figure 5: Impulse response functions for sentiment-topics predicted to be leader driven for the Republican Party. Bootstrapped 95-percent confidence intervals are shown.

6 Conclusion

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 stance and the uncertain nature of politics. We present

leaders within their own party explains the result

empirical support for our hypothesis that House party leaders initiate 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 House leadership's propensity to discuss a sentiment-topic where the coordination problem dominates, with leaders exerting a statistically significant influence in the short-run over their rank-and-file member's propensity to discuss that sentiment-topic. Notably, this effect also operates in 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. For a standard deviation ($\sim$10 percentage point) shock to leadership's propensity to discuss, we might observe 0.5 percent to 2 percent increases in rank-and-file's propensity to discuss. For the reverse, we see a standard deviation ($\sim$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

34

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). 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.

References

Aldrich, John and David Rohde. 1998. Measuring Conditional Party Government. Working paper Duke University.

URL: *https://www.researchgate.net/publication/251842516_Measuring_Conditional_Party_Government*

Aldrich, John and David Rohde. 2001. The logic of conditional party government. In *Congress Reconsidered*, ed. LC Dodd and BI Oppenheimer. Washington, DC: CQ Press p. 269–92.

Barbera, Pablo, Andreau 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.

Bendix, William. 2016. "Bypassing Congressional Committees: Parties, Panel Rosters, and Deliberative Processes." *Legislative Studies Quarterly* 41(3):687–714.

Bueno de Mesquita, Ethan and Anthony Fowler. 2021. *Thinking Clearly with Data: A Guide to Quantitative Reasoning and Analysis*. Princeton University Press.

Curry, James M. 2015. *Legislating in the Dark: Information and Power in the House of Representatives*. Chicago: The University of Chicago Press.

Dewan, TORUN and DAVID P. Myatt. 2007. "Leading the Party: Coordination, Direction, and Communication." *American Political Science Review* 101(4):827–845.

Gamm, Gerald and Steven Smith. 2020. The dynamics of party government in Congress. In *Congress Reconsidered, 11th Edition*, ed. LC Dodd and BI Oppenheimer. Washington, DC: CQ Press p. 197–224.

Hall, Thad E. and Betsy Sinclair. 2018. *A Connected America: Politics in the Era of Social Media*. Oxford University Press.

Harbridge, Laurel. 2015. *Is Bipartisanship Dead? Policy Agreement and Agenda-Setting in the House of Representatives*. New York: Cambridge University Press.

Howard, Nicholas O. and Mark E. Owens. 2020. "Circumventing Legislative Committees: The US Senate." *Legislative Studies Quarterly* 45(3):495–526.

Kang, T., E.F. Fowler, M.M. Franz and T.N. Ridout. 2018. "Issue consistency? Comparing television advertising, tweets, and e-mail in the 2014 Senate Campaigns." *Political Communication* 35:32–49.

Kim, In Song, John Londregan and Marc Ratkovic. 2018. "Estimating Spatial Preferences from Votes and Text." *Political Analysis* 26(2):210–229.

Koop, Gary, M. Hashem Pesaran and Simon M. Potter. 1996. "Impulse response analysis in non-linear multivariate models." *Journal of Econometrics* 74(1):119–147.

Lin, C., Y. He, R. Everson and S. Ruger. 2012. "Weakly Supervised Joint Sentiment-Topic Detection from Text." *IEEE Transactions on Knowledge and Data Engineering* 24(6):1134–1145.

Lin, Chenghua and Yulan He. 2009. "Joint Sentiment/Topic Model for Sentiment Analysis." *Proceedings of ACM International Conference on Information and Knowledge Management (CIKM-2009)* .

Roberts, Margaret E., Brandon M. Stewart, Dustin Tingley, Christopher Lucas, Jetson Leder-Luis, Shana Kushner Gadarian, Bethany Albertson and David G. Rand. 2014. "Structural Topic Models for Open-Ended Survey Responses." *American Journal of Political Science* 58(4):1064–1082.

Röder, Michael, Andreas Both and Alexander Hinneburg. 2015. 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

p. 399–408.

URL: *https://doi.org/10.1145/2684822.2685324*

Tiefer, Charles. 2016. *The Polarized Congress: The Post-Traditional Procedure of Its Current Struggles*. Lanham, MD: University Press of America.

Wallis, Kenneth F. 1987. "Time Series Analysis of Bounded Economic Variables." *Journal of Time Series Analysis* 8(1):115–123.

Wallner, James I. 2013. *The Death of Deliberation: Partisanship and Polarization in the United States Senate*. New York: Lexington Books.

Yan, Hao, Sanmay Das, Allen Lavoie, Sirui Li and Betsy Sinclair. 2019. 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.