Is There Anybody Out There?
The Effects of Legislators' Communication with their Constituents.

Pablo Barberá
Richard Bonneau
John T. Jost
Jonathan Nagler
Joshua Tucker

September 21, 2013

Abstract

Are legislators responsive to their constituents in their public communication? To what extent are they able to shape voters' preferences, as expressed by the issues they discuss? We address this twofold question with an analysis of all tweets sent by U.S. Members of Congress and a random sample of their followers from January to August 2013. Using a Latent Dirichlet Allocation model, we extract topics that represent the diversity of issues that legislators and ordinary citizens discuss on this social networking site. Then, we exploit variation in the distribution of topics over time to test whether Members of Congress lead or follow their constituents in their selection of issues to discuss, employing a Granger-causality framework. We find that legislators are more responsive to politically interested constituents and co-partisans, particularly on issues owned by each party.

1 Introduction

An enduring question in the study of democratic polities is the level of responsiveness of government to the preferences of the public. In order for members of a legislature to be responsive to public preferences, they need to be paying attention to the policy views and preferences of the public. In this paper we examine whether or not they are doing that by examining whether the issues discussed by Members of Congress via Twitter are influenced by the issues discussed by their constituents on Twitter. We also examine the ability of Members of Congress to influence the agenda and views of their constituents, by looking at whether or not the content of Tweets by members of the mass public is influenced by the topics that Members of Congress emphasize on Twitter.

We analyze all tweets by Members of Congress over a nine month period. Using a Latent Dirichlet Allocation model, we extract topics that represent the diversity of issues legislators discuss on this social networking site. We find that this method is able to classify legislators' tweets on a limited number of topics, with meaningful variation over time and across parties.

Next, we exploit the interactive nature of Twitter to understand to what extent legislators' expressed political agenda affects what their constituents discuss publicly on this same platform. Using tweets sent by users from different partisanship groups that follow members of Congress, we test whether longitudinal changes in the importance that legislators attribute to different issues affect voters' discussions about this same set of topics.

We are able to show that Members of Congress are surprisingly responsive in their public statements to their constituents. On many issues when co-partisans among the mass public devote more twitter posts to an issue, Members of Congress follow suit. On other issues, the opposite is true: when Members of Congress devote more twitter posts to an issue, co-partisans in the mass public become more likely to tweet about the same issue. Thus we have evidence for both: a necessary condition for representation (members of the legislature are aware of what the mass public is saying), *and* that members of the legislature are able to influence what members of the mass public are saying.

The rest of the paper proceeds as follows. In section 2 we discuss the existing literature on political responsiveness and agenda control, and present our theory and hypotheses. Section 3 describes our dataset of tweets sent by members of Congress and their followers. Section 4 introduces our topic modeling method and how we apply it when using Twitter data. Results of our analysis are shown in section 5. The article concludes in 6 with a summary of findings and a list of possible paths for future research.

2 Theory

There are two distinct questions for democratic polities that our paper relates to. The first important question is that of responsiveness. Does the government respond to what members of the public want? But the second important question is that of agenda-setting. How do issues get on the agenda? If the mass public talks about something, do members of the legislature observe that and react? And do members of the public shift their concerns based on the behavior, or speech, of their legislators? There is also a more narrow question related to United States politics that we consider: who do members of the legislature listen to – co-partisans, or all constituents.

Many scholars have looked at the US legislature to see who individual legislators are responsive to: examining them primarily to see whether legislators served co-partisans, or their district more broadly (see Clinton, 2006 for a recent example). The broader question of overall government responsiveness, rather than that of individual legislator responsiveness, has most recently shifted to discussion of *whom* it is responsive to. Gilens (2012) shows that government is more responsive to the policy preferences of the wealthy than to the policy preferences of the poor.

There are similarly broad and long-standing literatures on how different issues reach the political agenda (see for instance Baumgartner and Jones, 2009).

We approach the current analysis from an agnostic perspective, but with hypotheses to test. First, do members of the legislature observe, and respond, to the views of the mass public. If so, do they respond more to the views of co-partisans than of the general electorate. And, do members of the legislature have the ability to influence the agenda? Does the public follow issues that members of the legislature view as important.

We note several limitations before proceeding. We do not consider the role of the traditional mass media here. Thus it may well be that both the mass public and members of the legislature are following the mass media. We also do not distinguish among the mass public based on wealth, or any measure beyond the number of political actors they follow on Twitter.

3 Data

3.1 Members of Congress on Twitter

To test our hypotheses, we use tweets sent by members of the current House and Senate of the US. Twitter use in Congress has increased steadily over the past few years (Golbeck et al., 2010; Chi and Yang, 2010; Shapiro et al., 2012; Evans et al., 2013). Of all legislators that have served in the current Congress, around 90% of all Representatives (407 of 442)[1] and Senators (92 of 103)[2] have active Twitter accounts[3]. This proportion is similar across parties: 91% of republicans (257 of 281 current members) and 92% of democrats (240 of 261 current members).

The interest that Members of Congress show in using Twitter to communicate with their constituents is illustrated by the high number of tweets they send, a total of 258,096 during our period of study (from January 1st to August 31st, 2013). If we extend this period to the moment each of them created their accounts, this number increases to 650,832 tweets, with an average of 1,515 tweets per member and a median of 903 tweets (see Figure 1). By party, Republicans have sent a higher total of tweets (384,919 tweets vs 258,096), in part because they represent a majority of the current House, but also because they are more active, with a median of 971 tweets vs 906 tweets for Democrats.

[1] We include in our analysis Jason T. Smith, Representative for the 8th district in Missouri, who won a special election in June 2013 after the previous incumbent resigned; as well as the Representatives and Delegates for District of Columbia, Puerto Rico, American Samoa, Guam, Virgin Islands and Northern Mariana

[2] We include in our analysis William Cowan, who substituted John Kerry as junior Senator from Massachusetts; Edward Markey, who substituted William Cowan after he declined to run in a special election; and Jeffrey Chiesa, who substituted Frank Lautenberg as junior senator from New Jersey

[3] The list of Twitter handles of Members of Congress was collected through the New York Times Congress API.

[FIGURE 1 HERE]

There appears to be only mild public interest on what Members of Congress are writing on Twitter. The median Representative or Senator has 5,939 followers, although the average of followers is 17,072 due to the existence of a few outliers, such as Sen. McCain with 1,827,023 followers, Rep. Boehner with 515,208 followers, Rep. Pelosi with 378,812 followers, and Rep. Ryan with 361,132 followers. Figure 2 shows that there is ample variation in their number of followers, which Shapiro et al. (2012) found to be related to ideological positions of the members, with more extreme Members of Congress having larger audiences.

[FIGURE 2 HERE]

Golbeck et al. (2010) argue that Members of Congress use Twitter primarily to advertise their policy positions and to provide information about their activities. However, more recent studies have shown that this online tool can also be a tool for Members of Congress to be responsive to their constituents (Hemphill et al., 2013), to exercise control of the offline and online political agenda, to interact publicly with other Representatives and Senators, and to report on their constituency service (Evans et al., 2013). Figure 5 shows examples of each of these five types of tweets.

These descriptive statistics provide preliminary evidence that tweets sent by Members of Congress can be considered a meaningful representation of how legislators communicate with their constituents: both Members of Congress themselves and their constituents show interest in this platform, and a preliminary analysis of the content of the messages they exchange show a sophisticated use of this tool.

The dataset of tweets we use in this paper consists of all tweets sent by Members of Congress since January 1st, 2013 until August 31st, 2013. These tweets were collected by the Social Media and Political Participation Lab (SMaPP) at New York University using Twitter's Streaming API. As we show in Table 1, our dataset contains a total of 198,352 tweets.

[TABLE 1 HERE]

When we plot the distribution of tweets sent by Members of Congress by day, in Figure 3, we find very strong seasonality, with only 7.2% of all tweets sent during the weekends. Figure 4 shows the same data, but grouped by week, which smoothes the weekend-seasonality component of the figure.

[FIGURE 3 HERE]

[FIGURE 4 HERE]

3.2 Followers of Members of Congress

In addition to tweets sent by Members of Congress, we also collected tweets sent by a random sample of their followers. This will allow us to examine to what extent changes in what ordinary citizens are discussing affect the expressed political agenda of their representatives, and vice versa.

As of August 31st, a total of 3,817,409 unique Twitter users follow at least one member of congress[4]. Of these, 2,858,677 follow only one account, 426,350 follow two accounts, 180,167 follow three accounts, 99,260 follow four accounts, 60,597 follow five accounts, and 192,358 follow more than five accounts. There are relevant differences across parties: Republicans have more unique followers than Democrats (3,052,586 vs 1,040,475 users), but this difference is mostly due to the large audience that Sen. McCain has accumulated (over 1.8 million followers). Excluding Senator McCain's followers, Republicans have 1,467,692 unique followers.

As we show in Table 2, we subdivided the population of followers in six groups, according to the number of Members of Congress they follow (from one to three, "interested" users; more than three, "hardcore" users), and their party (democrats, republicans, and independents)[5]. We excluded from our sample those users who follow only Sen. McCain, since many of them are likely to be outside the US. Additionally, we chose three groups of politically interested Twitter users who do not follow any Member of Congress, but who do follow other politicians, journalists or media outlets. For each of these nine groups, we applied a simple spam and location filter to ensure that all the users in our sample are real citizens located in the US (see Barberá, 2013, p.11 for more details), and took a random sample of 2,000 users from each group. Finally, in order to divide our control groups according to partisanship, we applied the estimation method outlined in Barberá (2013) and classified users in the bottom 40% of the ideology distribution as democrats; those in the top 40% as republicans; and the other 20% as independents.

[TABLE 2 HERE]

Our final data collection step consisted on using the Twitter REST API to capture the tweets sent since January 1st by the 18,000 users in our sample of followers. After excluding users who deleted their accounts or made them private (see last column of Table 2), our effective sample size includes 15,952 users and a total of 14,084,352 tweets[6]. In Table 3 we present summary statistics for this dataset, disaggregated for each of our nine groups. We find that "hardcore" followers tend to be slightly more active on Twitter.

[TABLE 3 HERE]

4 Modeling Legislators' and Followers' Topic Distributions

Our purpose in this paper is to characterize the different issues that Members of Congress and their followers discuss on Twitter, and how their importance varies over time and across groups defined by their partisanship and political interest. To extract these categories, we estimate a probabilistic

[4] We collected the lists of followers for all Representatives and Senators with a Twitter account using Twitter's REST API

[5] We assume that Twitter users who follow Members of Congress of only one party identify with that party. See Barberá (2013) for justification of why following behavior is a strong predictor of partisanship.

[6] Note that this is not the total of tweets that these users sent since January 1st, given the limits of the Twitter API, which only allows to recover around 3,200 of the most recent tweets of each user. However, less than 20% of the users in our sample have sent more than 3,200 tweets since January 1st, so we don't expect this issue to affect our results.

model of word occurrences in documents called Latent Dirichlet Allocation (Blei et al., 2003), which belongs to a general category of latent variable models that infer topics from documents using a "bag-of-words" approach.

This approach is adequate for our application for at least two reasons. First, since it is an unsupervised machine learning technique, it is less subject to the biases inherent to imposing a set of topics pre-defined by the researcher. Second, it allows us to include the entire text of all documents in the analysis, therefore improving the validity and accuracy of our estimates for the topic distributions.

This model considers each document as a sequence of N words, denoted by $\mathbf{w} = (w_1, w_2, \ldots, w_N)$, extracted from a vector of length V containing all possible words in the corpus. (Note that the order of words is irrelevant). LDA treats each document as a random mixture over latent topics, and each topic as a distribution over words. Each document w in the corpus is the result of the following generative model (Blei et al., 2003, p.96):

1. The topic distribution for document w is determined by: $\theta \sim \text{Dirichlet}(\alpha)$
2. The word distribution for topic k is determined by: $\beta \sim \text{Dirichlet}(\delta)$
3. For each of the words in document w
   (a) Choose a topic $z_n \sim \text{Multinomial}(\theta)$
   (b) Choose a word w_n from $p(w_n|z_n, \beta)$, a multinomial probability conditioned on z_n.

This model requires us to fix K, the number of possible topics. There are two main parameters of interest: β, a matrix of dimensions $K \times V$ indicating the distribution of words over topics; and θ, a matrix of dimensions $K \times N$ indicating the distribution of topics over documents.

Our definition of "documents" is the aggregated total of tweets sent by members of Congress each day, by party. There are two reasons for this. First, LDA assumes that each document is a mixture of topics, which is appropriate for our conceptualization of each day's tweets as the political agenda that each party is trying to push for that specific day. Second, conducting an analysis at the tweet level is complex, given its very limited length. The existing literature on topic modeling of tweets has found that applications that aggregate tweets by author or day outperform those that rely on individual tweets (Hong and Davison, 2010).

Note two additional features of our analysis. First, we fit the model at first only for Members of Congress and then, as we describe in Section 5.2, use the estimated parameters to compute the posterior topic distributions of their followers' tweets, also aggregated by day, based on their observed words. We do so to make sure the topics we estimate are political in nature, and because our main focus is language use by Members of Congress.[7] Second, in our estimation we assume that topic distributions are independent over time, and that the number and content of each topic is constant. As we discuss in Section 6, an alternative approach would be to fit a Dynamic Topic Model (Blei and Lafferty, 2006; Quinn et al., 2010).

[7] We note that there is a limitation to this method. If mass followers discuss some political topics that Members of Congress never mention, we will not observe this.

We arbitrarily fix $K = 16$ based on the substantive fit of the topics to different political issues[8]. We fit the model with a collapsed Gibbs sampler (Griffiths and Steyvers, 2004; Phan et al., 2008), implemented in R (Grün and Hornik, 2011). We ran a single chain for 500 iterations, with a warm up period of 200 iterations. We apply the usual pre-processing text techniques (converting all words to lowercase and removing stopwords, all words shorter than 3 characters, and all words that appear less than 3 times in all documents), which gives us a vocabulary of $N = 50,285$ words.

5 Results

5.1 The Political Agenda of Members of Congress

The results of fitting our topic model are summarized in Figure 6. Here, each panel displays the top 10 words for each topic. To facilitate the interpretation, the size of the font for each word is proportional to its normalized score, this is, an index that measures how good a word w is at predicting whether document n belong to topic k. We also computed the specificity for each word, which indicates to what extent a word w_n is specific of a single topic k. The formulas for each of these indicators are:

$$\text{score}_{w,k} = \beta_{w,k}\left(\log\beta_{w,k} - \frac{\sum_{k'}\log\beta_{w,k'}}{K}\right)$$
$$\text{specificity}_{w,k} = \frac{\beta_{w,k}}{\sum_k \beta_{wk}}$$

[FIGURE 6 HERE]

Although there is some variation in the correspondences between each estimated topic and sets of political issues, we can identify the following topics:

1. Health Care (Democrats)
2. National Holidays
3. Health care (Republicans)
4. State of the Union Speech
5. Social Welfare
6. Sequester
7. Student Loans and DOMA
8. Inauguration of Congress
9. Boston Marathon Attack
10. Benghazi Hearing
11. Keystone Pipeline
12. Procedural Words
13. Democrats visiting their constituency
14. Republicans visiting their constituency
15. Budget
16. IRS scandal

[8] In future versions of the paper we will provide additional analysis justifying this choice, such as a comparison of the perplexity of the model using 10-fold cross-validation (see Blei et al., 2003, p.1008)

Based on the top scoring words for each topic, we can classify them in five different categories: topics that are related to Congress but not necessarily political (4, 8, 12), topics related to non-political events (2), topics related to exogenous events that could be politicized (9), and topics that are more related to the agenda of the Democratic (1, 5, 6, 7, 13) or Republican (3, 10, 14, 15, 16) party.

This categorization is helpful because it generates specific hypotheses about what the distribution of the attention devoted to each of these topics over time should be for each party, which will allows us to assess their construct validity[9].

That is precisely the purpose of Figure 7, where we plot the estimated $\hat{\theta}$, this is, the distribution of topics over time separately for each party, which roughly corresponds to the proportion of tweets that are classified in each of the 16 topics for a given day and party.[10]

[FIGURE 7 HERE]

These results are consistent with our theoretical expectations. Topics 5 (social welfare), 6 (sequester), and 7 (student loans and DOMA) are discussed almost exclusively by Democratic Members of Congress. Topics 10 (Benghazi hearing), 15 (budget), and 16 (IRS scandal), on the contrary, seem to be used only by Republicans. At the same time, the distribution for non-partisan topics, such as 2 (national holidays), 8 (inauguration), 9 (Boston marathon attack) and 12 (procedural words), is similar across both parties. Interestingly, Democrats and Republicans appear to be using different language to describe their constituency activities, which we identified based on its distribution over time (only weekends) and words ("community", "school", "center", "town", "hall", "county", "meeting", "great", "thanks"...)[11].

Overall, these results justify our topic modeling approach and show that the emerging categories represent coherent sets of issues that can be useful for our analysis of how Members of Congress and their constituents interact on Twitter.

5.2 The Political Agenda of Their Followers

In this section we present a descriptive analysis of our results after applying the estimated parameters for Members of Congress to the observed distributions of the words used by their followers and our control groups[12]. Using simulation, we compute $p(\hat{\theta}_F|\hat{\alpha}, \hat{\beta}, \mathbf{w}_F)$, the posterior distribution of

[9] This table also suggests that some topics encompass events that take place simultaneously (e.g. "Boston" and "tax" in Topic 9, since the so-called "tax day" is April 15, the same day of the attack). This is a problem for interpretation, for not for our analysis, since our main focus is to characterize how Members of Congress communicate on Twitter to understand if their followers cause (or respond to) variation on language use, independently of the substantive interpretation of the estimated topics.

[10] Note that in our estimation we pooled tweets from both parties into the same corpus, which allows us to extract topics that are common to both of them.

[11] We made this inference based on the distribution of these topics over time (see Figure 7), which increase during weekends and in August, when legislators visit their constituencies. While some words ("great", "thanks"), are common across both topics (see Figure 6), it appears that the rest differ across parties, even if they are semantically similar ("community", "school", "work", "center" for Democrats; "town", "county", "meeting", "hall" for Republicans), and that might explain why constituency activities emerge as two different topics in our analysis.

[12] We apply identical pre-processing techniques to those described in Section 4, and include only words that appear in the vocabulary of words used by Members of Congress.

topics over documents, $\hat{\theta}_F$ (i.e. tweets, aggregated by day and group), given the Dirichlet prior $\hat{\alpha}$, the estimated distribution of words over topics $\hat{\beta}$, and the observed word counts $\mathbf{w}_F$. This process gives us θ_F, this is, an estimate of how much each topic was discussed by our groups of ordinary Twitter users over the period of analysis.

Figures 8 and 9 display a selection of our results[13]. The first plot compares the relative importance of topics that we identified as being "owned" by Democratic Members of Congress in the political agenda of each of our three groups of Democratic users. This plot confirms our intuition that users who follow more Members of Congress write more tweets with a political content. The differences are particularly evident for Topic 6 (sequester), as we would expect given that it's very related to Congress' activities.

[FIGURE 8 HERE]

The second plot makes the complementary comparison, by showing the distribution of three political topics for hardcore democrats and hardcore republicans. As we expected, Topics 3 and 15 (health care, for Republicans; budget) are more frequently discussed by hardcore republicans, while the result is the opposite for Topic 7 (student loans and DOMA). However, note that these differences are not as wide as for Members of Congress, which is consistent with some of the insights of the literature on mass and elite political polarization (McCarty et al., 2006; Fiorina and Abrams, 2008). Overall, these results suggest that our measure of the political agenda of ordinary Twitter users has face validity, which allows us to draw meaningful comparisons across groups and also with Members of Congress.

[FIGURE 9 HERE]

5.3 Members of Congress and their Followers: Who's Listening to Whom?

The key substantive question we want to answer is whether the distribution of topics discussed by Members of Congress leads or follows that of their constituents. Are members reacting to their constituents, or vice-versa?

As a preliminary analysis of our results, in Figure 10 we show contemporaneous correlations of the aggregated distributions of topics over time[14], which demonstrate that topics discussed by Members and topics discussed by constituents are related. As we hypothesized, the average correlations are higher for groups of hardcore users than for interested users; and also higher for interested users than for uninterested users. In other words, ordinary Twitter users tend to use language that is more similar to that of Members of Congress the higher the number of Representatives and Senators they follow. There are also substantive differences across parties. If we focus on the three groups of hardcore users, we find that distribution of topics for hardcore Democrats is much more correlated with Democratic Members of Congress than with Republican Members of Congress. The result is the opposite for hardcore Republicans, with Hardcore Independents being in the middle. Interestingly, while these differences are consistent across groups with different levels of political

[13] Full results for all topics and groups are available upon request

[14] In other words, we stack the topic distributions over time for all 16 topics in 11 different variables, one for each group of Members of Congress or followers, and then compute pairwise Pearson's correlations.

interest, their magnitude decreases for uninterested users, which is also an indication of low political polarization at the mass level.

[FIGURE 10 HERE]

Figure 10 shows that topic distributions across groups are similar, and that the differences are consistent with our theoretical expectations, but does not address whether longitudinal changes in these distributions are also related. To try to determine causality, we now turn to the standard Granger-causality framework (Granger, 1969). To limit the scale of our analysis, we focus our attention on six groups: Democratic and Republican Members of Congress, hardcore Democrat and Republican constituents, and uninterested Democrat and Republican constituents.

For each of our six groups, we regressed the proportion of tweets on topic k at time t sent by that group[15] on the proportions of tweets for each of the six groups (including itself) in each of the preceding three days[16]. In other words, we computed 6 different regressions for each topic, each of them with 18 independent variables (six groups times three lags).

We then performed a standard F-test for the joint significance of each set of three lagged variables corresponding to our group of interest. When the F statistic is significant, we can reject the null that the topic distribution of the group to which the lagged variables correspond has an effect on the topic distribution of the group to which the dependent variable corresponds. When the test suggests that one group's tweets help predict the other group's tweets, but not vice-versa, then we conclude that one group's tweets granger-causes the other group's tweets. If we conclude that both groups' previous tweets help predict contemporaneous tweets of the other group, then we conclude that both groups are influenced by previous behavior of the other.

In Figure 11 we report the main results of our analysis. Each cell corresponds to one test; for example, the first cell on the top left-hand corner indicates whether changes in the topic distribution for hardcore Democrats have a significant effect on the topic distribution of Democratic Members of Congress; and the second cell displays the same result when the analysis is conducted in the opposite direction. The color of the cells indicates whether the F statistic[17] leads us to reject the null hypothesis (black cells) or not (grey cells). To facilitate the interpretation, the color of each topic indicates how we classified it (see Section 5.1).

[FIGURE 11 HERE]

These results only partially meet our theoretical expectations. On one hand, we find that Members and Congress and their followers do not influence each other with respect to non-political topics such as national holidays, procedural words or constituency activities. This is consistent with the intuition that, if there exists any effect, we shouldn't find it with regards to this type

[15] We exclude weekends from the analysis to avoid measurement error induced by the strong weekly seasonality in our data, as we described in Section 3

[16] We chose 3 days as the appropriate lag structure after examining the fit of models of Members' tweets as a function of different lag lengths of Members' tweets. From a theoretical perspective, it is also unlikely that any effects take place with a lag of more than 3 days

[17] Note that, given that we are conducting multiple F-tests, around 5% of them will have p-values lower than 0.05 by chance only.

of topics. However, these results do not provide a common answer to the question of who follows whom: the direction of the Granger causality varies across topics. This variation fits with our theoretical expectations in some cases, but not all: for example, Members of Congress have an impact on their followers for topics that are more closely related to their activities (State of the Union, social welfare, sequester). The same applies to differences across parties: for example, Republican Members of Congress and their constituents influence mutually in a clearly Republican topic (the Benghazi hearing).

We complement this analysis with a placebo test, consisting on replicating our analysis for cross-party effects. Figure 12 summarizes our results. As before, each cell indicates the p-value of F-tests on the joint significance of the lagged variables, but in this case we assess whether variation in the topic distribution of hardcore Republicans (Democrats) has a significant effect on the topic distribution of Democratic (Republican) Members of Congress, and vice versa.

[FIGURE 12 HERE]

In this case, we would expect only around 5% the p-values for the F-tests to be higher than 0.05, if we assume that hardcore users and Members of Congress do not influence each other when they belong to different parties. The results of this placebo test provide only partial evidence in support of the validity of our approach. We find that Democratic Members of Congress have no significant influence on hardcore Republicans, but the same analysis for Republican Members of Congress and hardcore Democrats show significant effects in both directions.

6 Discussion and Conclusions

In this paper we have characterized the political agenda of Members of Congress and their constituents using latent topic modeling applied to the text of their tweets. While further work is necessary to validate our approach, our analysis demonstrates that meaningful insights can be extracted from an analysis of how Representatives and Senators communicate through social media.

We plan to extend and improve our present study in two directions. First, one important advantage of using social media messages as a source of information about political discussion is that it allows us to increase the granularity of our observations. So far our unit of analysis was tweets aggregated by day, but it would be easy to reduce it half days or even hours in order to adapt to the fast pace of Twitter; and also aggregate it to weeks or months to study long-term changes in language use. Similarly, we aggregated all Democratic and Republican Members of Congress, but we could also focus on individual legislators and their followers. In doing so, we need to be aware of the trade-offs that these choices present: more granularity can also be accompanied by more measurement error in the estimation of topic distributions.

A second possible improvement would lie in our methodology. As we discussed in Section 4, our model assumes independence over time and that topics are stable in number and content. Other latent topic models such as the dynamic topic model introduced by Blei and Lafferty (2006), whose most prominent application in Political Science can be found in Quinn et al. (2010), would allow us to relax these assumption and thus could be a more appropriate choice that might improve the validity of our estimates. Similarly, we could also complement our empirical strategy with (semi-)supervised approaches, such as focusing on the use of specific words that belong to a set of pre-defined issues, or manual classification of tweets in topics. Along the same lines, our estimates require further validation, which could be achieved by splitting Members of Congress and their followers into different groups (e.g. House vs Senate, senior vs junior Members of Congress, male vs. female, by geographic location...) and then examining whether their posterior topic distributions are consistent with our theoretical expectations. It is also important to note that we did not exclude retweets from our analysis, which could be an important source of common variation in language use across groups.

The scope of this paper is limited to an analysis of Members of Congress and their constituents, but the method we have developed could be applied to any group of Twitter users, including media outlets and journalists, politicians and ordinary users in different countries, celebrities, etc. Our analysis could also be replicated in the context of a study of social influence: to what extent do media outlets follow or set the public agenda? what types of political issues spread across borders? what characteristics make specific groups of users influential in patterns of language use? These are some of the questions that we hope to address in our future work.

References

Barberá, P. (2013): "Birds of the Same Feather Tweet Together. Bayesian Ideal Point Estimation Using Twitter Data." *Unpublished manuscript, New York University.*

Baumgartner, F. and B. D. Jones (2009): *Agendas and Instability in American Politics*, Chicago, Illinois: University of Chicago Press, 2nd ed.

Blei, D. M. and J. D. Lafferty (2006): "Dynamic topic models," in *Proceedings of the 23rd international conference on Machine learning*, ACM, 113–120.

Blei, D. M., A. Y. Ng, and M. I. Jordan (2003): "Latent dirichlet allocation," *the Journal of machine Learning research*, 3, 993–1022.

Chi, F. and N. Yang (2010): "Twitter adoption in congress," *Review of Network Economics, Vol. 10, No. 1.*

Clinton, J. D. (2006): "Representation in Congress: Constituents and Roll Calls in the 106th House," *Journal of Politics*, 68, 397–409.

Evans, H. K., V. Cordova, and S. Sipole (2013): "Twitter Style: An Analysis of How House Candidates used Twitter in their 2012 Campaigns," in *Paper Presented at the 2013 Midwest Political Science Association Conference.*

Fiorina, M. P. and S. J. Abrams (2008): "Political polarization in the American public," *Annu. Rev. Polit. Sci.*, 11, 563–588.

Gilens, M. (2012): *Affluence and Influence: Economic Inequality and Political Power in America*, Princeton, New Jersey: Princeton University Press.

Golbeck, J., J. M. Grimes, and A. Rogers (2010): "Twitter use by the US Congress," *Journal of the American Society for Information Science and Technology*, 61, 1612–1621.

Granger, C. W. (1969): "Investigating causal relations by econometric models and cross-spectral methods," *Econometrica: Journal of the Econometric Society*, 424–438.

Griffiths, T. L. and M. Steyvers (2004): "Finding scientific topics," *Proceedings of the National academy of Sciences of the United States of America*, 101, 5228–5235.

Grün, B. and K. Hornik (2011): "topicmodels: An R package for fitting topic models," *Journal of Statistical Software*, 40, 1–30.

Hemphill, L., J. Otterbacher, and M. Shapiro (2013): "What's congress doing on twitter?" in *Proceedings of the 2013 conference on Computer supported cooperative work*, ACM, 877–886.

Hong, L. and B. D. Davison (2010): "Empirical study of topic modeling in twitter," in *Proceedings of the First Workshop on Social Media Analytics*, ACM, 80–88.

McCarty, N. M., K. T. Poole, and H. Rosenthal (2006): *Polarized America: The dance of ideology and unequal riches*, MIT Press Cambridge.

Phan, X.-H., L.-M. Nguyen, and S. Horiguchi (2008): "Learning to classify short and sparse text & web with hidden topics from large-scale data collections," in *Proceedings of the 17th international conference on World Wide Web*, ACM, 91–100.

Quinn, K. M., B. L. Monroe, M. Colaresi, M. H. Crespin, and D. R. Radev (2010): "How to analyze political attention with minimal assumptions and costs," *American Journal of Political Science*, 54, 209–228.

Shapiro, M., L. Hemphill, and J. Otterbacher (2012): "Doing What I Say: Connecting Congressional Social Media Behavior and Congressional Voting," .

Tables and Figures

Table 1: Number of tweets in dataset (Members of Congress)

Members of Congress	N	Average	Std.Dev.	Min	Max	Tweets
House Republicans	212	424.0	437.8	0	3528	89,879
Senate Republicans	42	487.7	464.8	50	2592	20,482
House Democrats	186	363.3	320.6	0	1814	67,572
Senate Democrats	48	425.4	258.6	59	1458	20,419
Total	488	406.5	384.9	0	3528	198,352

Table 2: Distribution of followers

Treatment groups	N	%	n_s	n_a
Hardcore democrats (follow >3 Dem. MCs and 0 Rep. MCs)	33,945	0.9	2,000	1,861
Hardcore independents (follow >3 MCs of either party)	178,333	4.8	2,000	1,820
Hardcore republicans (follow >3 Rep. MCs and 0 Rep MCs)	129,959	3.5	2,000	1,853
Interested democrats (follow $1-3$ Dem. MCs and 0 Rep. MCs)	657,071	17.6	2,000	1,829
Interested independents (follow $1-3$ MCs of either party)	171,126	4.6	2,000	1,801
Interested republicans (follow $1-3$ Rep. MCs and 0 Dem. MCs)	1,020,658	27.3	2,000	1,788
Follow only Sen. McCain	1,552,510	41.5	0	
Control groups				
Uninterested republicans (follow no MCs; ideology on top 40%)	–	–	2,000	1,563
Uninterested democrats (follow no MCs; ideology on bottom 40%)	–	–	2,000	1,683
Uninterested independents (follow no MCs; ideology on middle 20%)	–	–	2,000	1,754

n_s = intended sample. n_a = actual sample. Numbers differ because private and deleted accounts were excluded.

Table 3: Number of tweets in dataset (Followers)

Members of Congress	Average	Std.Dev.	Min	Max	Sum
Hardcore democrats	1,023.5	1,061.9	0	3,293	1,936,428
Hardcore independents	996.3	1,045.5	0	3,406	1,862,065
Hardcore republicans	887.4	1,049.2	0	3,275	1,677,191
Interested democrats	849.5	973.5	0	3,269	1,581,846
Interested independents	844.7	976.8	0	3,271	1,562,614
Interested republicans	718.2	900.4	0	3,255	1,337,209
Uninterested democrats	808.8	959.3	0	3,287	1,448,479
Uninterested independents	834.6	933.0	0	3,253	1,433,025
Uninterested republicans	765.5	921.0	0	3,253	1,245,495
Total	860.9	987.6	0	3,406	14,084,352

Figure 1: Distribution of number of tweets sent by Members of Congress

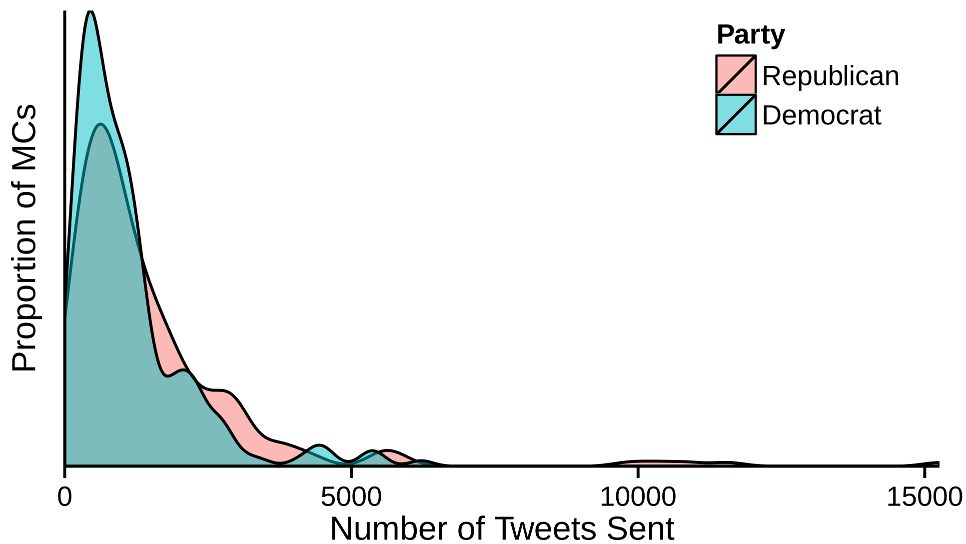

Figure 2: Distribution of number of followers of Members of Congress

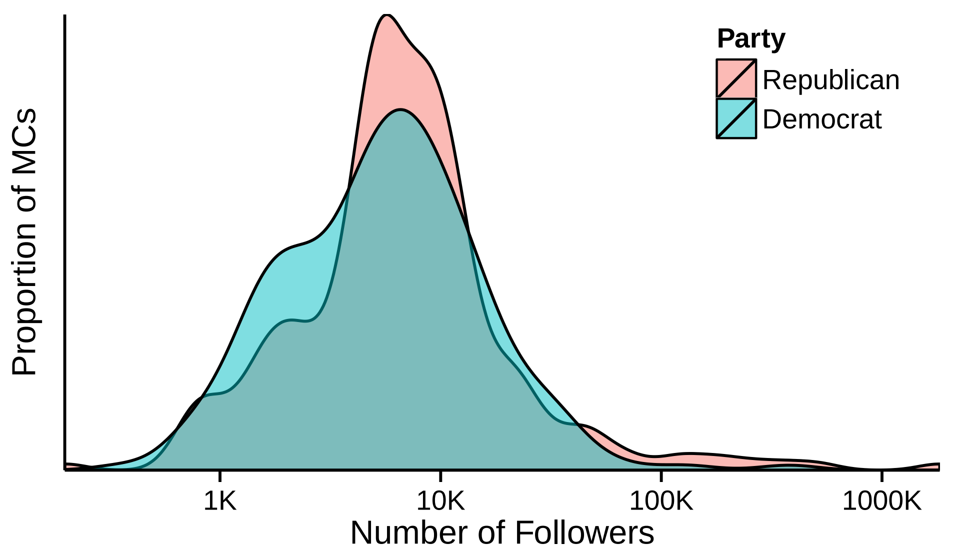

Figure 3: Number of tweets in dataset, by day

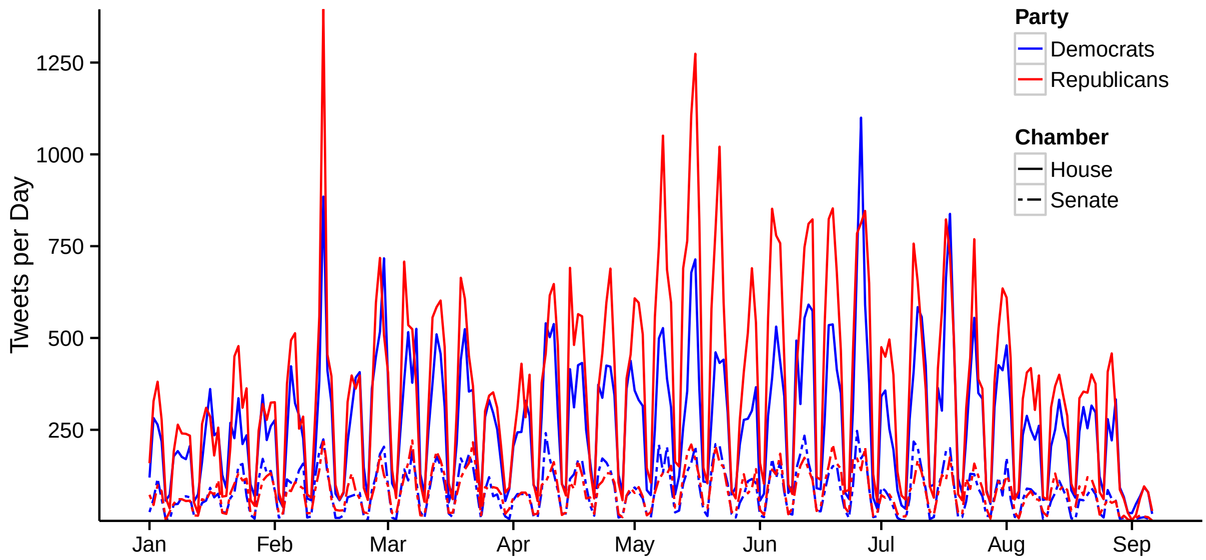

Figure 4: Number of tweets in dataset, by week

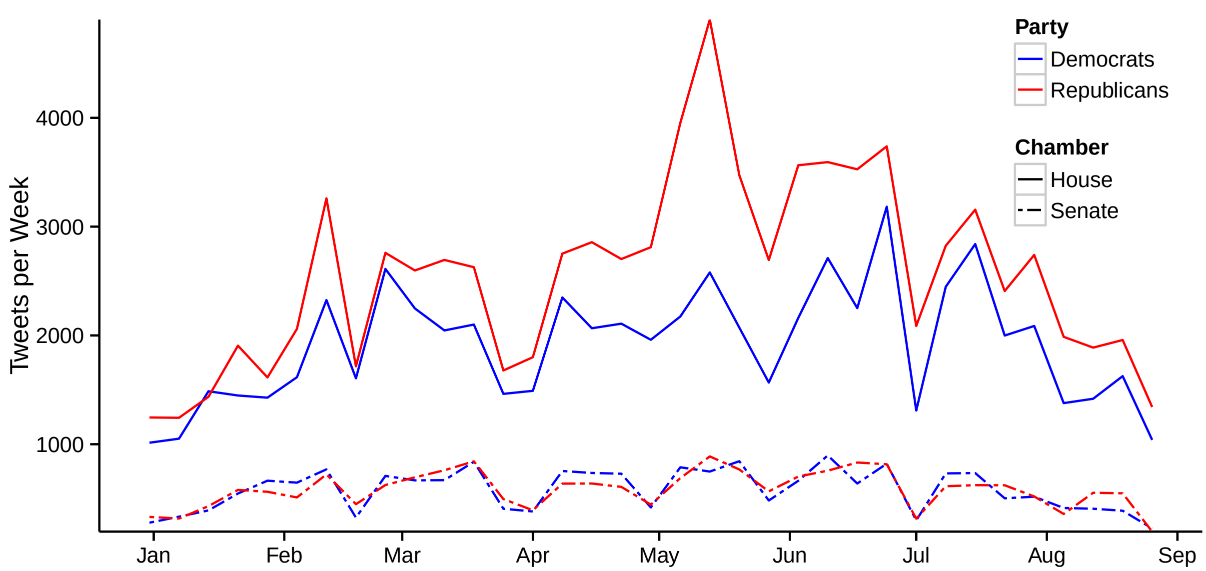

Figure 5: Examples of Tweets Sent by or to Members of Congress

(a) Position-Taking Tweet

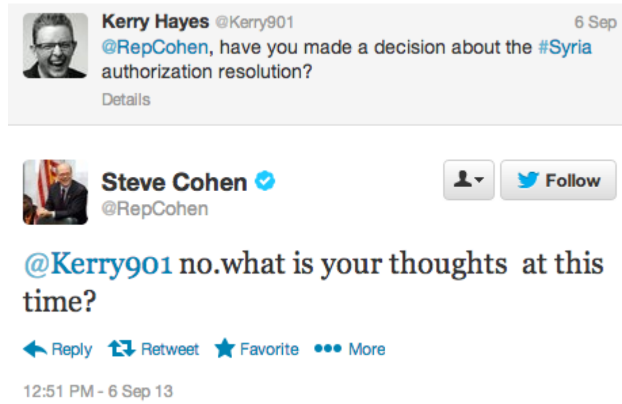

(b) Responsiveness Tweet

(c) Agenda-Setting Tweet

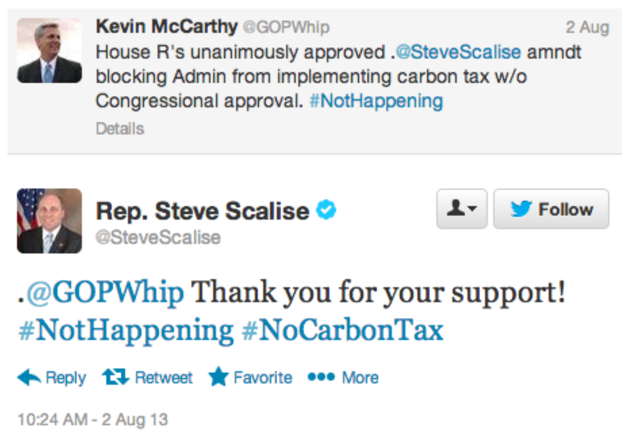

(d) Interaction Tweet

(e) Constituency Service Tweet

Figure 6: Top Scoring Words in each Topic

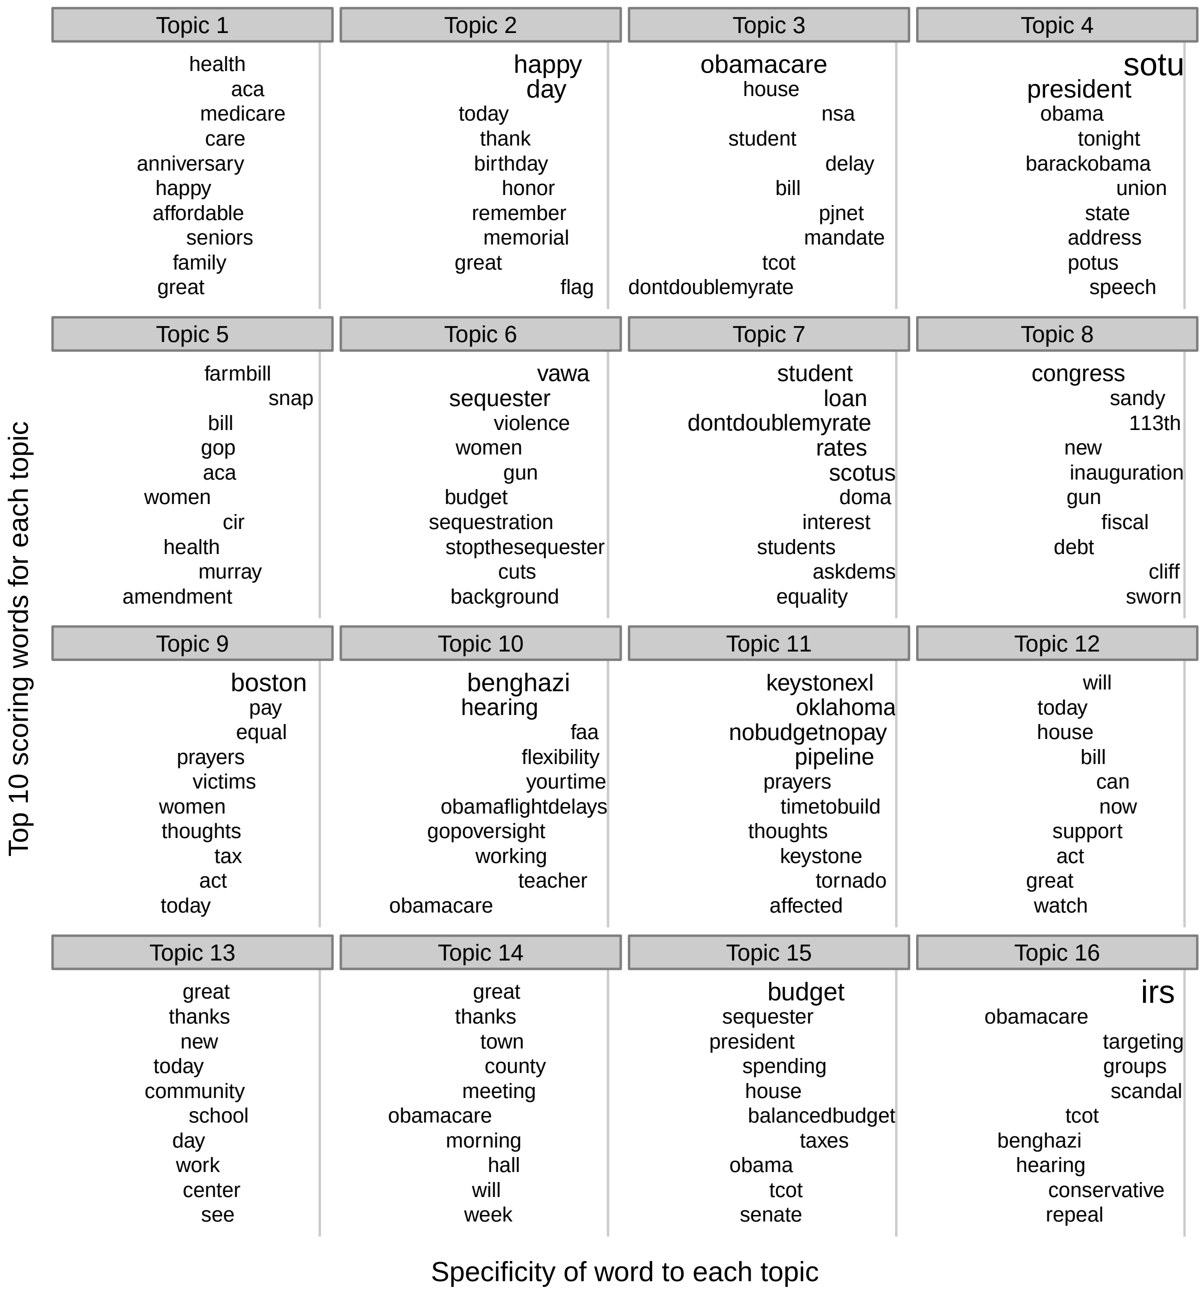

Figure 7: Distribution of Topics over Time, by Party

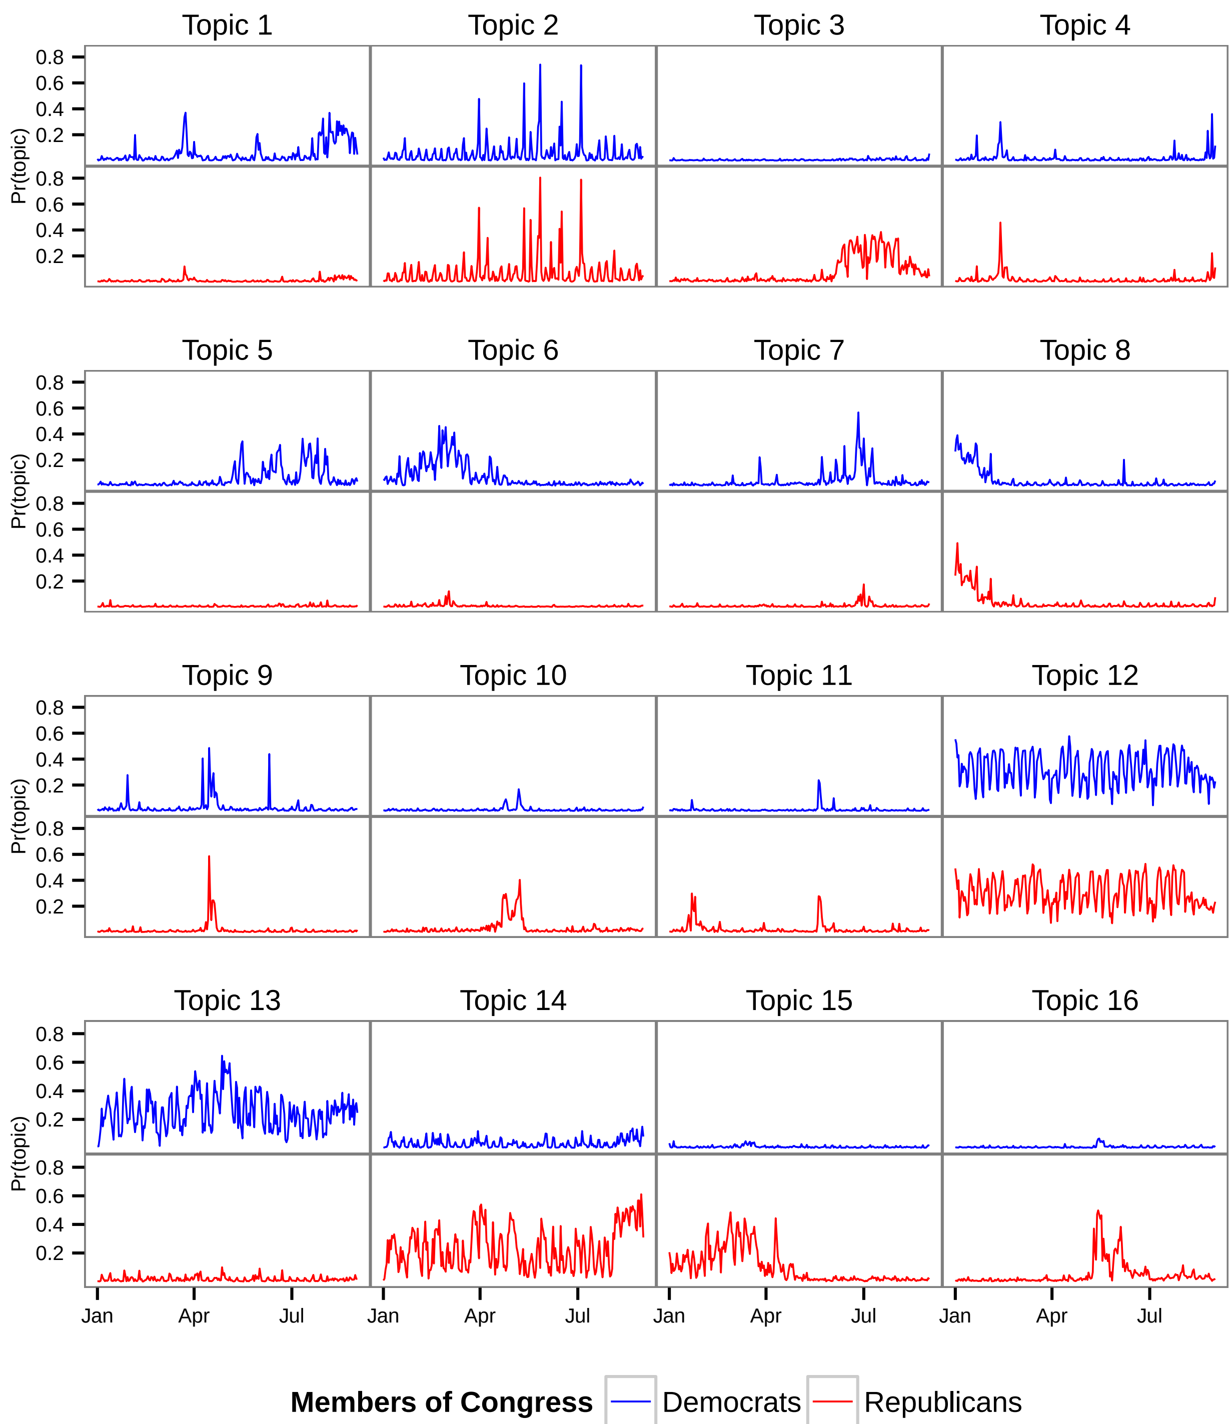

Figure 8: Distribution of Topics (Selection), for Democratic Followers

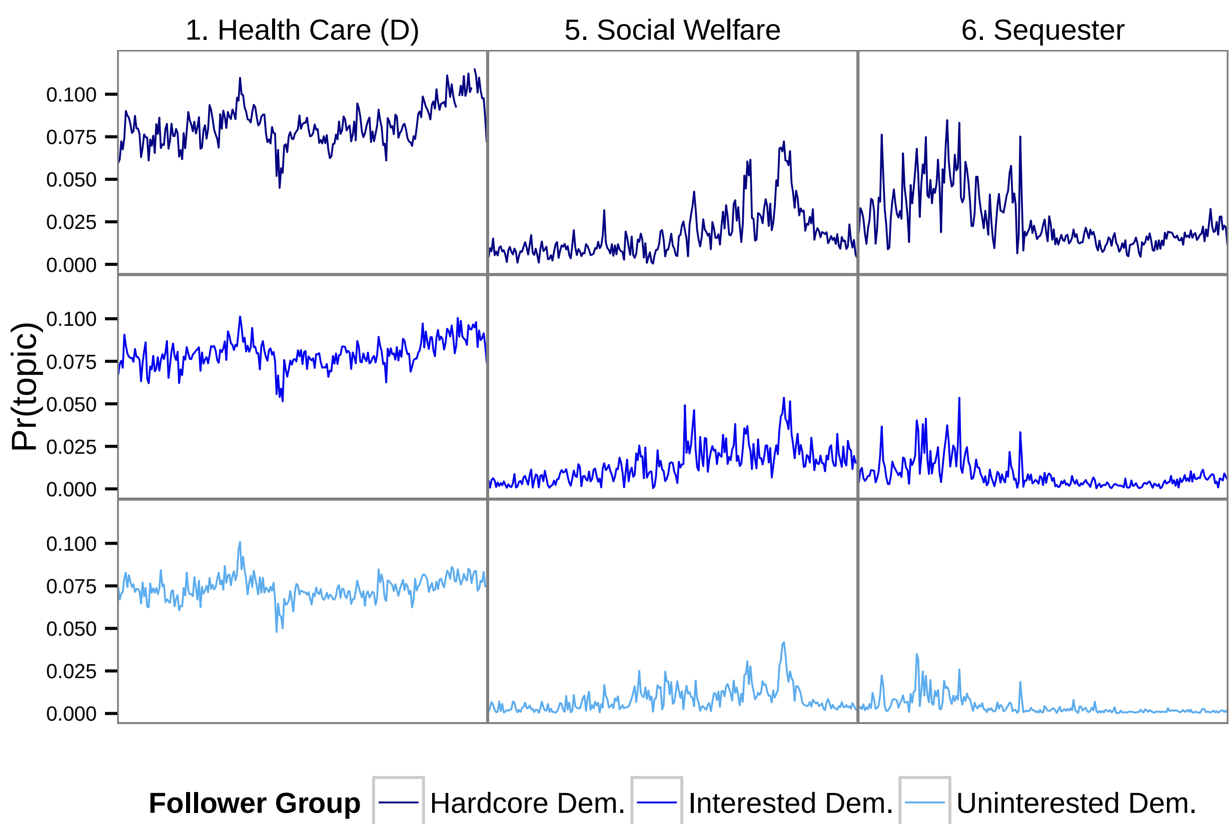

Figure 9: Distribution of Topics (Selection), for Hardcore Followers

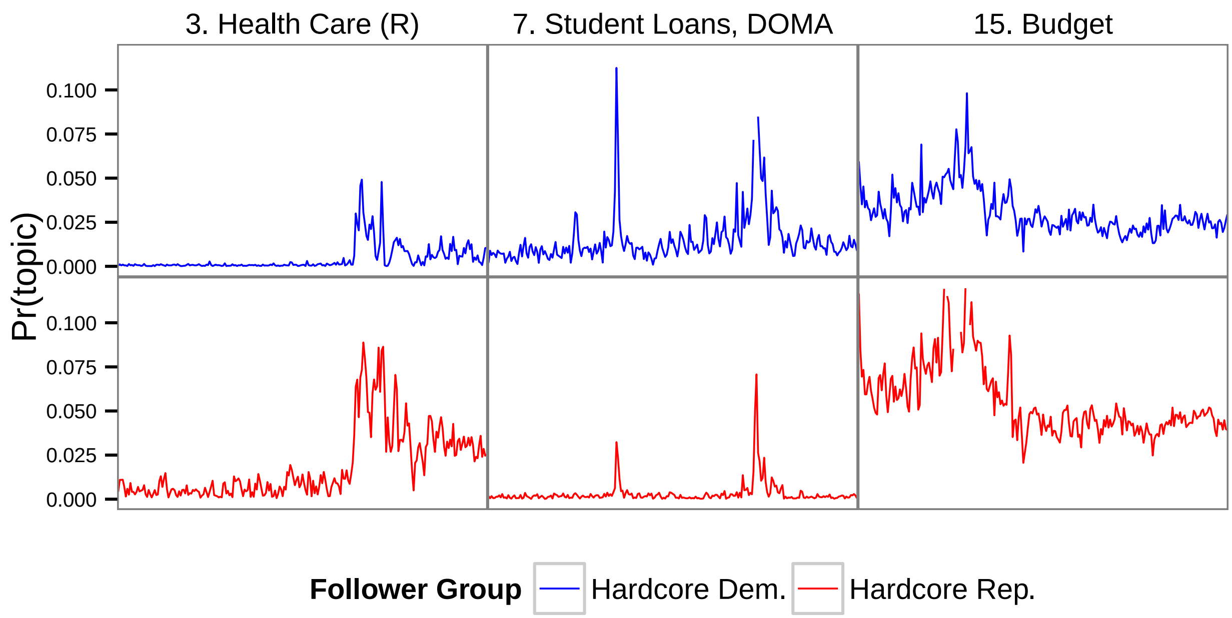

Figure 10: Correlations in Contemporaneous Topic Distributions of Members of Congress and their Followers

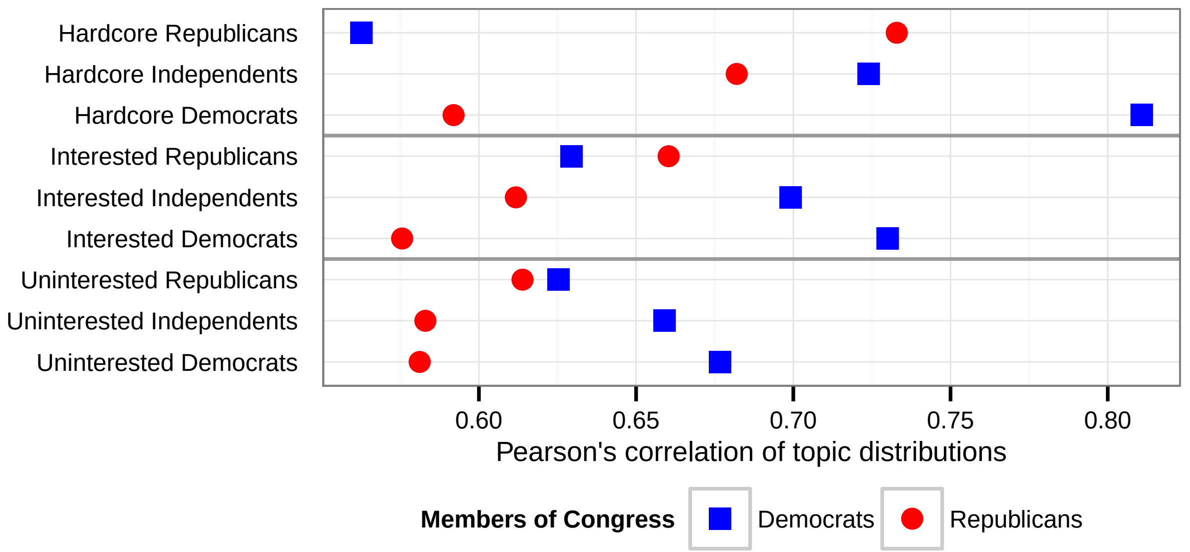

Figure 11: Results of Granger Causality Tests

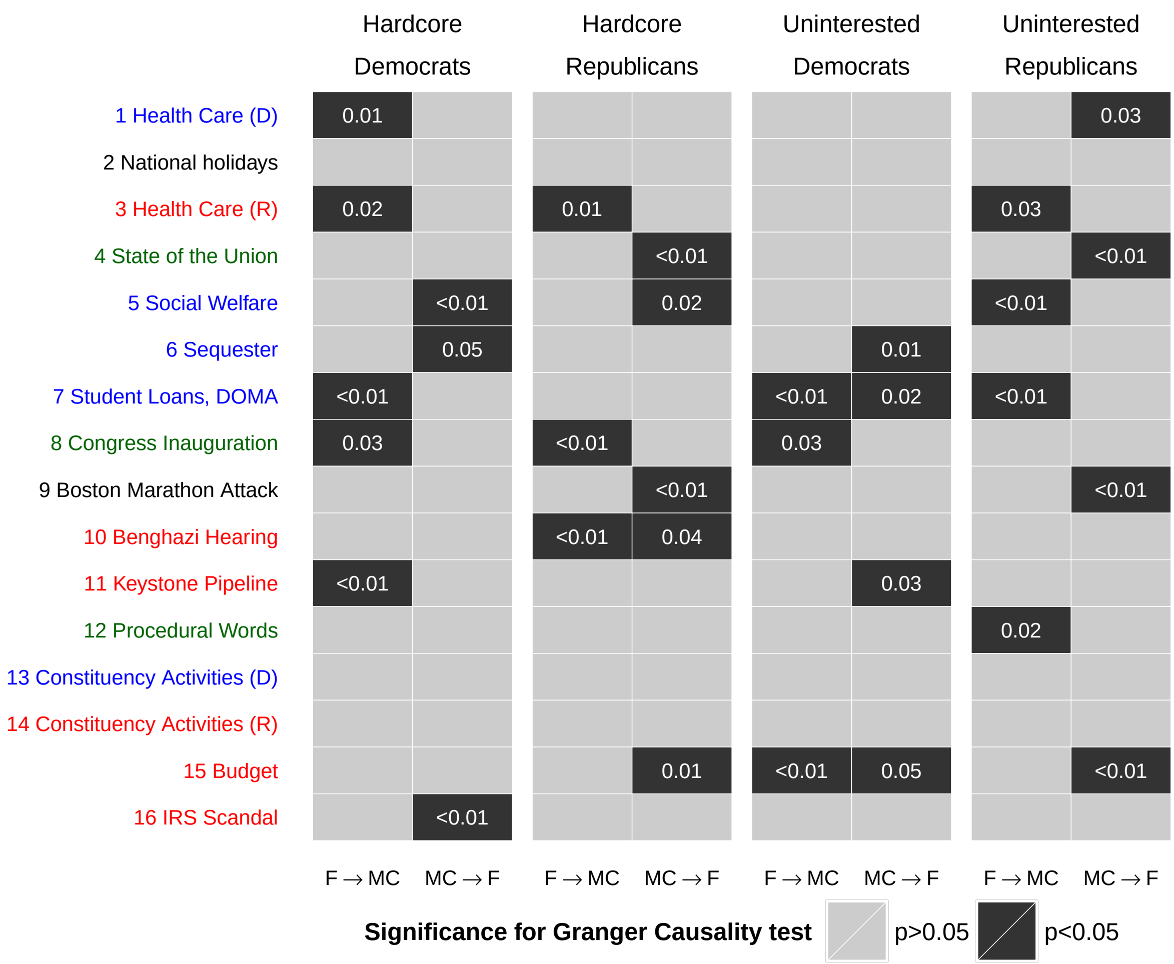

Figure 12: Results of Granger Causality Tests (Placebos)

	$HC_R \rightarrow MC_D$	$MC_D \rightarrow HC_R$	$HC_D \rightarrow MC_R$	$MC_R \rightarrow HC_D$
1 Health Care (D)				
2 National holidays				
3 Health Care (R)			0.03	0.02
4 State of the Union				0.04
5 Social Welfare				
6 Sequester	0.04			0.01
7 Student Loans, DOMA			<0.01	
8 Congress Inauguration				
9 Boston Marathon Attack			0.01	<0.01
10 Benghazi Hearing	0.01			
11 Keystone Pipeline			0.02	<0.01
12 Procedural Words	0.03		0.03	
13 Constituency Activities (D)				
14 Constituency Activities (R)				
15 Budget	<0.01			<0.01
16 IRS Scandal	0.04			

Significance for Granger Causality test p>0.05 p<0.05