Politics, Sharing and Emotion in Microblogs

Tuan-Anh Hoang
Living Analytics Research Centre
Singapore Management University
Email: tahoang.2011@phdis.smu.edu.sg

William W. Cohen
Machine Learning Department
Carnegie Mellon University
Email: wcohen@cs.cmu.edu

Ee-Peng Lim
Living Analytics Research Centre
Singapore Management University
Email: eplim@smu.edu.sg

Doug Pierce
Political Science Department
Rutgers University
Email: drpierce@eden.rutgers.edu

David P. Redlawsk
Political Science Department
Rutgers University
Email: redlawsk@rutgers.edu

***Abstract*—In political contexts, it is known that people act as "motivated reasoners", i.e., information is evaluated first for emotional affect, and this emotional reaction influences later deliberative reasoning steps. As social media becomes a more and more prevalent way of receiving political information, it becomes important to understand more completely the interaction between information, emotion, social community, and information-sharing behavior. In this paper, we describe a high-precision classifier for politically-oriented tweets, and an accurate classifier of a Twitter user's political affiliation. Coupled with existing sentiment-analysis tools for microblogs, these methods enable us to systematically study the interaction of emotion and sharing in a large corpus of politically-oriented microblog messages, collected from just before the 2012 US presidential election. In particular, we seek to understand how information sharing is influenced by the political affiliation of the sender and receiver of a message, and the sentiment associated with the message.**

Keywords. Emotional affect, Sentiment analysis, Information sharing, Microblogs

I. Introduction

It is well-understood that political decision-making is strongly influenced by affective (i.e., emotional) factors [17], [18], [22], [23]; in particular, political psychologists often model people as *motivated reasoners*, where information is evaluated immediately for affect, and affectively incongruent (i.e., cognitively dissonant) information is processed differently from affectively congruent information. One striking instance of apparently non-rational decision making in politics is illustrated by the "affective tipping point" phenomenon [24] in which moderate amounts of negative information about a candidate who is initially liked by a voter increases, rather than decreases, the voter's positive stance toward the candidate.

Due to the rise of sites like Facebook, Twitter, and Weibo, political information is more and more frequently encountered in a social context: even stories published by mainstream media sites are often encountered by users after having been shared by others. Clearly, this social context can influence how information is interpreted and re-shared. As an example, we note that there are rational-actor models which predict the counter-intuitive "affective tipping point" behavior noted above, *in certain social contexts*: for instance, it is rational to behave this way if information comes not from a neutral source, but from a biased source with different preferences, and the information is provided with the intent of manipulating the receiver [21]. Thus, to understand political decision-making in the context of news disseminated by social media, it is necessary to understand the interplay between emotion and the social context of information. It is equally necessary to understand this interplay to predict how information will diffuse through a social-media platform.

In this paper, we study the potentially complex interaction between sharing, community membership, and emotion on the microblogging platform Twitter. We crawl an appropriate subset of Twitter, and develop a high-precision classifier for politically-oriented tweets. We also develop an accurate classifier for the political alignment (liberal or conservative) of Twitter users. Combining these resources with an external sentiment classifier for tweets, we are able to test certain specific hypotheses about sharing and emotion in politics. Additionally, we build a topic model for our corpus, and use this to more precisely characterize sharing behavior for politically-oriented microblog entries—by subject, rather than in the aggregate.

Other than technical contributions on development of high-precision classifiers for politic tweets, we also analyze the frequency of various types of retweet behavior: specifically, we consider the probability of retweets of type T over all users in our sample, for several types T defined based on sentiment and political affiliation of the author of the tweets. (For instance, T might be "positive tweets from a left-leaning user"). For each specific type T, since many users who have few retweets are unlikely to have any retweets of type T, the average probability thefore is very close to zero. Hence, we also compute the average over all user with non-zero probability, i.e., all users that have retweet(s) of type T. Interestingly, these averages are sometimes qualitatively different, perhaps suggesting different information-sharing strategies across different users. In considering tweets from a particular community, we also consider averages over "balanced users" - users who follow some sources outside their own community.

Our key findings are as follows. (1) For users that have some retweets of the corresponding type, emotional tweets are more likely to be retweeted than neutral ones, and negative tweets more likely to be retweeted than positive ones. (2) For balanced users, retweets from within the user's party are more likely to be retweeted; however, for balanced users that have ever retweeted an opposite-party tweet, retweets from the opposite party are more likely. (3) Numerically, most retweets are of intra-party rather than inter-party. Thus, both sentiment and political affiliation have effects on information sharing, though their effects are different for different type of users. We further characterized patterns of topics in political retweets with regard to sentiment and political affiliation.

II. Related Work

Recently, in political science, Pierce *et. al.* reported human-subject experiments on how people handle political information from different sources [21]. The experiment tested two hypotheses related to sharing: (1) the *affective transmission hypothesis*, that people are more likely to share information that engenders an emotional reaction; and (2) the *social transmission hypothesis*, that people are more likely to share information that comes from inside their social group. These experiments were based on small groups of subjects in a mock election setting, and one goal of this paper is to test these hypotheses in a more realistic setting.

Because Twitter has been heavily used for information sharing, product broadcasting, and political campaigning [30], [12], [9], [15], studying the flow of messages in Twitter has been interesting to researchers from a variety of fields. In non-political contexts, for instance, Wu *et. al.* conducted a large scale analysis and found that Twitter users adopt more information from elite users (i.e., celebrities, famous bloggers, etc), than from other users within the same community [27]. Zhiming *et. al.* showed that source features like trustworthiness, source expertise, and source attractiveness have effects on retweeting [29]. Suh *et. al.* reported that the presence of URLs and hashtags in a tweet are useful indicators of whether a tweet will be retweeted [26]. There is also prior work on the correlation between retweetability and network characteristics in general, non-political settings; e.g., Hansen *et. al.* showed that negative news and positive non-news tweets are retweeted more often [14].

In the domain of politics, Berger *et. at.* investigated a large number of articles from New York Times and found that positive and negative content is more likely to be emailed by readers to their friends [4]. Similarly, Stieglitz *et. al.* examined a set of political tweets and found that there is a positive relationship between the quantity of words indicating sentiment in a tweet and the tweet's retweetability [25]. Another popular topic of investigation is the effect of social-media news transmission on information diversity. Conover *et. al* examined networks among Twitter users that are formed by retweeting relationships (i.e., edges are drawn from a user to other users she retweets) [8]. They found that the network is highly polarized—i.e., users tend to retweet more from other users sharing the same political affiliation. Jisun *et. al.* also investigated how the diversity of information from mass media is affected by retweets and sharing behavior [16]. In social science, Bristor showed that people are more likely to accept information from highly trustworthy sources [7], and Dodele *et. al* showed that information items that engender emotional reaction are more likely to be shared [10].

All of these prior results in the political domain can be viewed as relating to either the social transmission hypothesis, or the affective sharing hypothesis. However, because in each case the effects of sentiment and community membership are considered separately, it is not clear how these two factors interact. In short, while it seems clear that messages are more likely to be shared if they are either within-community or sentiment-bearing, we do not know to what extent these influences can be separated: e.g., it might be that polar messages are shared more readily, on average, simply because more of them come from in-group than average. In contrast, our work considers the effects of sentiment and community membership both independently *and together*. We also make use of topic modeling to obtain more fine-grained descriptions of emotion and sharing.

III. Data Collection

We first selected a set of 56 widely-followed Twitter users which we call the *seed users*. These users include major American politicians, especially candidates for American presidential election in 2012, e.g., Barack Obama, Mitt Romney, and

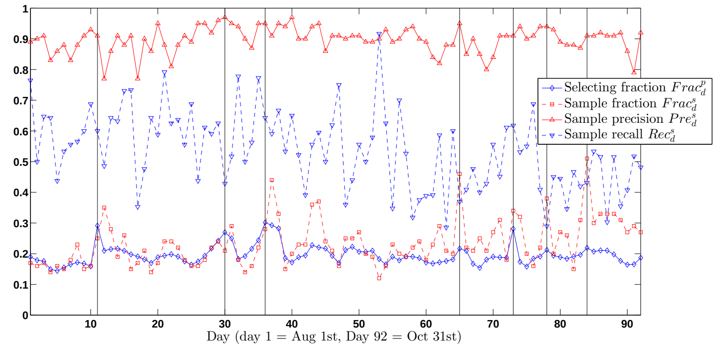

Fig. 1: Performance of proposed political tweet selection method

Newt Gingrich; well known political bloggers, e.g., America Blog, Red State, and Daily Kos; and political sections of mass media sites, e.g., CNN Politics, and Huffington Post Politics. Of the 56 seeds, 20 and 23 seed users are left- and right-leaning respectively, and the remaining 13 seed users are neutral.

We then expanded the set of users by identifying all users following at least three seed users, leading to a set of more than 408K users. A large part of the following network among those users[1], and tweets from all users were then crawled, with the goal of collecting all tweets between August and October 2012. This data thus contains both network and activity information for Twitter users actively following American politics.

The period from August to October 2012 was chosen for two reasons. First, since for each user we can only crawl up to her latest 3200 tweets, it was necessary to restrict the time range to obtain a relatively-complete sample (the crawling was done in fall 2012). Secondly, this period is very politically active, including many events related to the American 2012 presidential election: e.g., the national conventions of both democratic and republican parties, and the debates between presidential candidates. In total, this subset contains 127,812,186 tweets published by 349,976 users.

We found that both the numbers of tweets and retweets per user follow power-law distribution, with many users having very few tweets or retweets. In order to ensure that we have a reasonable amount of activity for each user, we further restricted the dataset to *active* users, who generated at least 45 tweets and/or retweets, and are active (i.e., post a tweet or retweet) for at least 15 days. This yields a smaller dataset of 104,119 users, and 105,167,766 tweets, which will be used in the subsequent analysis.

IV. Methodology

A. Political Tweet Selection

Many of the tweets in our dataset are not political in nature, so the first task was to build a classifier that detects political tweets. Although learning a classifier is a natural approach, the rapidly-changing nature of the dataset makes it difficult to ensure that a learned classifier will properly track political discourse over time; thus, we manually constructed a high-precision classifier.

Hashtag selection. A hashtag is a word prefixed by a # symbol added to a tweet to indicate topic. A tweet can have zero or more hashtags. We manually selected a set of political hashtags recommended by mass media streams, e.g., the hashtag guide from Washington Post[2], and hashtags used

[1] At the time of this writing, we have completely crawled the following network for only some of the users.

[2] http://www.washingtonpost.com/blogs/post-politics/post/hashtag-guide-for-the-2012-election–atmentionmachine/2011/12/05/gIQAOczocO_blog.html

by seed users (e.g, Obama first used *#dontdoublemyrate* to post tweets about his proposal not to double rates on student loans). We manually examined all hashtags of the latter type to filter out those that are not political and those used by fewer than three seed users. In the end, we obtained a list of 608 hashtags.

Keyphrases selection. Since not all tweets contain hashtags, we also developed a set of political keyphrases. We took a set of users that could be manually labeled with respect to political affiliation of users (as presented below, in Section IV-D), and aggregated all the tweets of each user that contain one or more of the political hashtags to form a user-document. We then ran an LDA model [5] with 50 topics, and for each topic, we manually examined the topic's top 20 most frequent words, and filtered out non-political ones. We combined single words into phrases when appropriate, e.g., "white" and "house" are combined into "white house". From this, we obtained an additional 224 political keyphrases.

Performance of the tweet selection method. In the end, we obtained a set of 25,738,776 tweets containing one or more political keyphrase or hashtag. For each day d, we denote by T_d the set of all tweets posted on the day, and denote by T_d^p the set of all tweets posted on the day and selected by our method. Then, the amount of tweets posted in day d and selected as political can be measured by *selected fraction* $Frac_d^p$ which is defined as follows

$$Frac_d^p = \frac{|T_d^p|}{|T_d|}$$

The selected fraction of each day over the three months is shown in Figure 1. On most days, about 20% of tweets are considered political. This fraction peaks on days associated with certain political events, as indicated by the vertical lines: e.g., Romney named Paul Ryan as his teammate on August 11th (day 11); Romney and Obama were nominated by their parties as candidate for the 2012 presidential election on August 30th (day 30) and September 5th (day 36) respectively; and there are also peaks on the days after the debates on October 3rd (day 65), October 11th (day 73), October 16th (day 78), and October 22nd (day 84).

As the groundtruth is not available, we evaluate the performance of our political tweet selection method presented above using daily random samples of tweets that are manually labeled. To do this, for each day d, we randomly select a set S_d of 100 tweets from T_d, and randomly select a set S_d^p of 100 tweets from T_d^p. We then manually labeled these 200 tweets as political or non-political. The fraction of (true) political tweets in day d is now can be approximated by *sample fraction* $Frac_d^s$ which is define as follows

$$Frac_d^s = \frac{|true\_S_d|}{|S_d|}$$

where $true\_S_d$ is the set of tweets in S_d that are hand-labeled as political tweets. Similarly, the precision and recall of our methods with respect to day d can be respectively estimated by *sample precision* Pre_d^s and *sample recall* Rec_d^s which are defined as follows

$$Pre_d^s = \frac{|true\_S_d^p|}{|S_d^p|} \qquad Rec_d^s = \frac{|true\_S_d \cap T_d^p|}{|true\_S_d|}$$

where $true\_S_d^p$ is the set of tweets in S_d^p that are labeled as political tweets

Figure 1 also shows the sample fraction, sample precision, and sample recall of all 92 days in the dataset (August 1st to October 31st, 2012). The figure clearly shows that the fraction of political tweets selected by our method is similar to the sample fraction of true political tweets. The figure also shows that our method achieves high precision (about 90%) and a usefully-high recall (roughly 60%). Looking deeper into political tweets in daily samples S_ds that are not selected by our method, we found that most of these tweets do not have hashtags, or talk about politics using non-political terms.

B. Tweet Sentiment Polarity Detection

To automatically analyze the sentiment expressed in every tweet, we employ the widely used Stanford's sentiment scoring API[3]. This API implements a machine learning method to detect sentiment polarity specifically in tweets [1]. For each tweet, the API turns the tweet content into a score that indicates whether the tweet is positive, negative, or neutral. We use the API to score all extracted political tweets. To evaluate the performance of the API, we manually labeled a set of 1000 randomly selected tweets for sentiment. The overall accuracy of detected sentiment for this set is 77.9%, which is reasonably good considering the small size of the tweets.

C. Tweet Topic Modeling

Since we also want to understand sharing behavior at a more fine-grained level, we built a topic model for our corpus. Direct application of the LDA model [5] to a collection of tweets is not appropriate: as tweets are very short, there is little information in such a corpus on term co-location, which is what drives LDA's grouping of terms into topics. Aggregating tweets from each user to form user-documents (as we previously did in keyphrase extraction) is also not ideal, as this approach may assign multiple topics to each tweet, and for analysis, it is preferable to have a single topic per tweet. Our solution is to employ the TwitterLDA model [28], a variation of LDA that constrains each tweet to have only one topic. The generative process TwitterLDA model is as follows.

- For each topic k, sample a topic-pecific word distribution $\phi_k \sim Dir(\beta)$
- For each user u, sample a user-specific topic distribution $\theta_u \sim Dir(\alpha)$
- For each tweet t by user u:
  - Sample a single topic for the tweet $z_t \sim Multinomial(\theta_u)$
  - For each word position n in tweet t, sample the word $\omega_{t,n} \sim Multinomial(\phi_{z_t})$

To perform inference for the model, we make use of Gibbs sampling. Here, we leave out the derivation of sampling equations due to the limitation on space. We first remove all stop words from tweets. We also remove all infrequent words, i.e., words that appear in fewer than 5 tweets. Then, we filter out all users having less than 10 tweets that are not empty after removing stop words and infrequent words. Next, we divide the set of tweets contributed by each user into two subsets, one with 90% of the tweets from that user, and one with 10%. By combining the larger subsets of all users together, we obtain the training set. The test set is constructed in a similar way using the smaller subsets. Finally, to identify an appropriate number of topics, we run collapsed Gibbs sampling on the training set, varying the number of topics, and evaluate the perplexity on the test set.

Figure 2 shows the perplexity of the TwitterLDA model with respect to the number of topics K. As expected, larger

[3]http://help.sentiment140.com/api

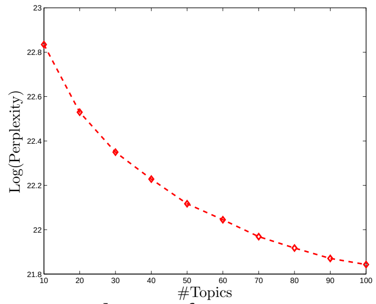

Fig. 2: Perplexity of TwitterLDA model

K gives smaller perplexity, and the amount of improvement decreases as K increases. In consideration of time and space complexity, we set the number of topics to be 80 in the experiments below.

D. User Classification

Perhaps the most challenging problem we encountered was to classify users by their political communities. A natural approach is to use a machine-learned classifier plus a set of gold-standard labels like previous works on the problem e.g, [20] and [6]. However, there are challenges in applying methods proposed in these works in our case. First, it is not immediately obvious how to obtain gold-standard labels for the political affiliation of Twitter users; and second, some features used in these works are computationally expensive, or not available due to the partially crawled following network in our dataset.

To deal with the above challenges, we first used a combination of manual effort and heuristic rules to obtain the training set of labeled users. We then used the SVM classifers trained on tweet-based features, plus a post processing step to increase confidence. These classifiers are chosen since (1) they were shown to have a comparative performance with the best ones [20], and (2) our main objective is not in classifying users but to obtain a “big enough” set of users whose labels are assgined with some level of confidence. We present our method in detail below.

Our first approach was to make use of Twitter *bios* to classify users. These bios are user contributed public statements for self-introduction. Other than background and hobbies, some users indicate political affiliations in their bios. We selected bios containing some term (e.g., “Democrat”, “Conservative”, etc) associated with political affiliation. Unfortunately, this population of users turned out to be both relatively small, and quite unrepresentative: manually classifying all bios in the set produced 1,116 left-leaning users, 10,775 right-leaning users, and 480 neutral users. These 12,371 users represent about 3% of total number of users in our Twitter dataset; and only 4,650 of them are active. One worrisome issue is that in this set of users, right-leaning users outnumber the left-leaning ones by nearly ten to one, an unlikely proportion even in a social media that is believed to be more heavily used by conservatives. In the following, set of all users manually labeled by their bios is denoted by M and called the *manually labeled set* (M).

Our second attempt to obtain gold-standard labels was based on the 56 seed users. The well-documented partisan division of social media [11] makes it likely that users who follow primarily right-leaning seeds are themselves right-leaning, and *vice versa*. We thus considered the following heuristic rules:

- Active users who follow Barack Obama but do not follow any right-wing seed users are considered left-leaning users.
- Active users who follow Mitt Romney but do not follow any left-wing seed users are considered right-leaning users.

In this set, called the *heuristically labeled set* (H) below, there are 15,998 left-leaning users, and 30,900 right-leaning users.

TABLE I: Performance of different user classifiers: precision and recall in identifying left-leaning users, and overall accuracy. Please refer to text for meaning of different training/test sets.

Training set	Test set	Precision (%)	Recall (%)	Accuracy (%)
$H - M$	M	73.4	71.0	94.4
M	$H - M$	92.2	47.9	87.2
$90\%(H \cup M)$	$10\%(H \cup M)$	87.7	77.1	92.8
$90\%(H \cup M)$	$10\%(H \cup M) - M$	77.3	73.4	93.8
$90\%(H \cup M)$	$10\%(H \cup M)$ (drop 5%)	91.0	80.7	94.6
$90\%(H \cup M)$	$10\%(H \cup M)$ (drop 10%)	93.2	84.9	96.2

The approximately 2:1 ratio of right- to left-leaning users in H is much more plausible than the ratio in M. The labels in H also coincide quite well with the manual labels: we found that there are 2,204 users in both M and H, and among those users, there are only 35 (about 1.6%) users which are labeled differently by the two methods.

Although the sets M and H mostly agree in their labeling, the distribution of users in these two sets is quite different, so it is not obvious which is the best set for training a classifier. To understand this, we performed some experiments using a SVM[4] trained on different training sets and evaluated on corresponding test sets as shown in Table I. Each user is represented as a document, consisting of all of her political tweets, and the features of each user-documents are TF-IDF scores [19] of the terms contained in the user-document. The results are shown in first rows of Table I. (The precision and recall measure the ability of the classifiers to identify left-leaning users, the minority class, against right-leaning users.) In addition to M, the manual labels, we consider H, the heuristic labels, and $M \cup H$, which denotes the union of these two sets. In $M \cup H$, for the 35 users where M and H disagree on the labels, we take the label from M. We also use $M - H$ and $H - M$ to denote the appropriate set differences.

The first four rows of Table I show that training on $H - M$ and testing on M produces fairly good results (94.4% accuracy), but training on M and testing on $H - M$ performs much less well (only 87.2% accuracy). This suggests that H is a more representative training set than M. A random 90%-10% training-test split of $H \cup M$ gives good accuracy (92.8%), and this result improves slightly (to 93.8%) when the (possibly non-representative) test examples from M are removed.

We concluded from these experiments that training on a dataset that at least includes H is highly desirable, and adopted the classifier from the fourth line of the table (trained on a 90% sample of $H \cup M$).

The classifier we learned is a two-class classifier, which classifies users into left-leaning and right-leaning. Since we are not interested in analysis of neutral users, and since most manually labeled neutral users are generally classified with low confidence, we also evaluated discarding the users associated with the lowest-confidence 5% and 10% of the predictions. By doing this, the accuracy, precision and recall of the classifier for left-leaning users increases slightly on the test set, as shown in the last two rows of Table I. We finally elected to discard the lowest-confidence 10% of the predictions, obtaining 21,948 left-leaning users; 66,118 right-leaning users; and

[4] http://svmlight.joachims.org/

16,053 neutral users. Among those users, there are 9,023 left-leaning users, denoted by U_L, and 25,398 right-leaning users, denoted by U_R, who we have full information abouth their follow relationship, i.e., all of their followers and followees, both included in our dataset or not. Next, our analysis works were performed on these U_L and U_R.

V. Results and Discussion

A. Sentiment and Retweetability

We first examine the effect of sentiment independent of the political affiliation of the sender. Specifically, we look at tweets sent from and received by members of the same political affiliation, and consider how the sentiment of tweets affects the probability of a retweet. More precisely, for each user u in $U_L \cup U_R$, and each type of tweet t, we estimate the probability that u retweets a tweet of type t as the number of tweets of type t that u retweets divided by the number of tweets of type t that u receives from her followees. These numbers are then aggregated over all active users: since our pool of users is large, even small differences in average probability of retweets can often be statistically significant.

Figures 3(a) and (b) show box-plots of retweet probability of tweets among left-leaning users, with respect to four different types of tweets. Figure 3 (a) shows the overall probability over users in U_L, while Figure 3 (b) shows the probability over users in U_L who ever retweeted the corresponding type of tweets. In the figure, *L** denotes all tweets received from left-leaning users; and *L+*, *LN*, and *L-* denote all positive, neutral, and negative tweets received from left-leaning users, respectively. Similarly, Figures 3 (c) and (d) respectively show box-plots of retweet probability for tweets posted by right-leaning users and received by users in U_R and by users in U_R who ever retweeted the corresponding type of tweets, for the analogous four types of tweets.

From Figures 3(a) and (c), we can see that, in general, neutral tweets are more likely be retweeted. This may be the a result of the fact that non-neutral tweets are rare (less than 10%), so most of the retweeting rates of these tweets for individual users are zero. However, the opposite pattern holds if for each type of tweet, we only consider users who ever retweeted tweet(s) of the type (i.e., in plotting *L**, we take all users that ever retweeted any tweet from other left-leaning users; in plotting *L+*, we only take users that ever retweeted any positive tweet from other left-leaning users; and similarly with the other types of tweets). As shown in Figures 3(b) and (d), for users who ever retweeted negative or positive tweets, those emotional tweets are more likely be retweeted for both political affiliations. We conducted two-tailed tests showing that, for those users, negative and positive tweets are statistically significantly more likely be retweeted than neutral tweets; and negative tweets are also statistically significantly more likely be retweeted than positive tweets.

Other than using the whole U_L and U_R, we performed the same analysis on different subsets of users derived from these two sets to make sure that we have confidence on results obtained above. These subsets of users are: (a) U_L^h - set of all users $u \in U_L$ who has a least 50% of her followees in our set of active users, and U_L^l - set of all users $u \in U_L$ who has less than 50% of her followees in our set of active users, and similarly we derived U_R^h and U_R^l from U_R; and (b) U_L^is $(i = 1, \cdots, 5)$ - five disjoint random subsets of U_L (users in U_L are evenly- and randomly-distributed into U_L^is), and similarly, we derived U_R^is (($i = 1, \cdots, 5$) from U_R.

By using U_L^h/U_R^h or U_L^l/U_R^l, we were able to measure the effects in the set of users where we are more or less confident about tweets they received, while by using U_L^is/U_R^is, we were able to measure the effects under randomization. Experiments on both U_L^h and U_L^l, and on all U_L^is use the qualitatively same results as in Figures 3(a) and (b); and, similarly, experiments on both U_R^h and U_R^l, and on all U_R^is give the qualitatively same results as in Figures 3(c) and (d). These further confirms the robustness of the obtained results.

B. Political affiliation and Retweetability

We now evaluate the effect of political affiliation on retweetability, independently from the effect of sentiment. To do this, we consider the set of *balanced users*, defined to be users who follow some left-leaning and some right-leaning users. Within U_L, we have 8,696 balanced left-leaning users, and within U_R, we have 8,962 balanced right-leaning users. For these balanced users, we compare the retweetability of two tweet types corresponding to whether the tweet's sender sharing the same or having different political affiliation.

Figure 4(a)(i) shows box-plots of retweet probability over all balanced left-leaning users in U_L, while Figure 4(b)(i) shows box-plots of retweet probability over all balanced left-leaning users in U_L who ever retweeted tweet(s) of the corresponding type. In these figures, again, *L** and *R** denote all tweets received from left-leaning and right-leaning users respectively. Similarly, Figure 4(c)(i) shows box-plots of retweet probability over all balanced right-leaning users in U_R, while Figure 4(d)(i) shows box-plots of retweet probability over all balanced right-leaning users in U_R who ever retweeted tweet(s) of the corresponding type (computed similarly to Section V. A), for the two types of tweets.

The box-plots in Figures 4(a)(i) and (c)(i) show that, in general, balanced users are more likely to retweet tweets from other users sharing the same political affiliation. This agrees with the social transmission hypothesis of Pierce *et. al.* [21]. However, this is not fully true if, for each type of tweets, we consider balanced users who ever retweeted tweet(s) of the type. The box-plots in Figure 4(b)(i) show that balanced left-leaning users who ever retweeted tweets from right-leaning users are still more likely to retweet tweets from other users sharing the same political affiliation; while box-plots in Figure 4(d)(i) show that balanced right-leaning users who ever retweeted tweets from left-leaning users are more likely to retweet tweets from other users having different political affiliation. Though this surprisingly contradicts the social transmission hypothesis for the community of right-leaning users, our result agrees with previous works showing that weak ties are more helpful in dissemination of information [13], [2].

Similarly to the previous analysis, we performed the same analysis on subsets of balance users in U_L^h/U_L^h/U_L^is and U_R^h/U_R^h/U_R^is and obtained results qualitatively the same to the above ones, confirming the robustness of the laters.

C. Sentiment & Political affiliation and Retweetability

Next, we examine the effects of political affiliation together with sentiment. Figures 4(a)(ii), (iii), and (iv) show box-plots of retweet probability of different type of tweets over all balanced left-leaning users in U_L, while Figures 4(b)(ii), (iii), and (iv) show box-plots of the probabilities over all balanced left-leaning users in U_L who ever retweeted tweet(s) of the corresponding type. In these figures, again, *L+*, *LN*, and *L-* denote all positive, neutral, and negative tweets received from by left-leaning users respectively; and *R+*, *RN*, and *R-* denote all positive, neutral, and negative tweets posted by right-leaning users respectively. Similarly, Figures 4(c)(ii), (iii), and

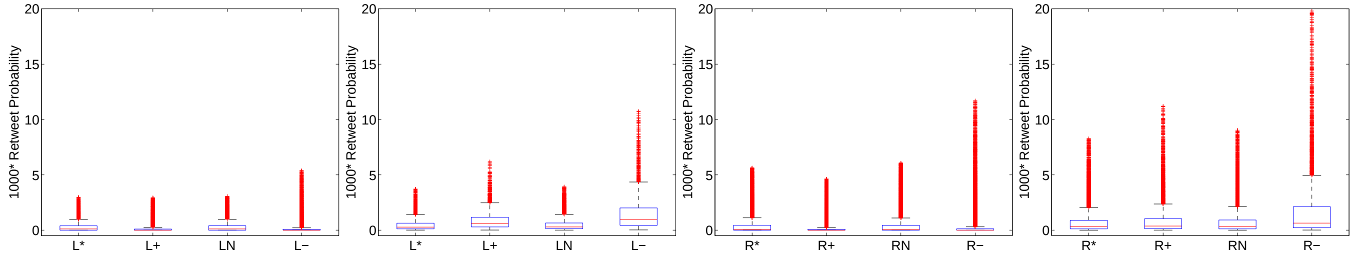

(a) All left-leaning users (b) Left-leaning users who ever retweeted tweets of the corresponding type (c) All right-leaning users (d) Right-leaning users who ever retweeted tweets of the corresponding type

Fig. 3: Intra-political affiliation retweets

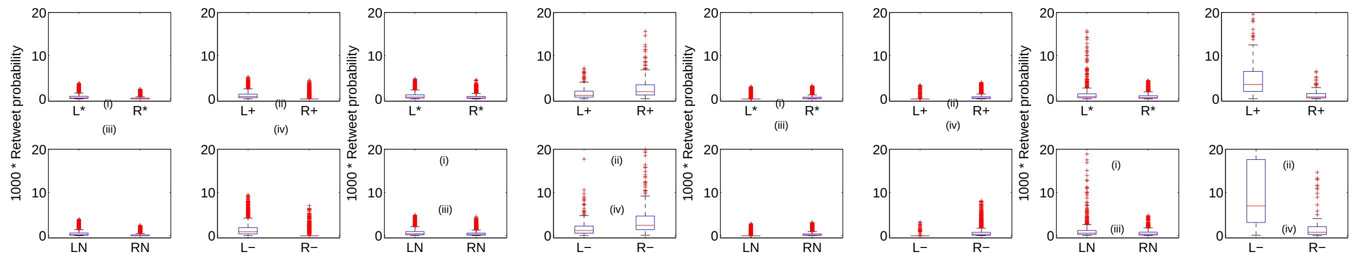

(a) All balanced left-leaning users (b) Balanced left-leaning users who ever retweeted tweets of the corresponding type (c) All balanced right-leaning users (d) Balanced right-leaning users who ever retweeted tweets of the corresponding type

Fig. 4: Inter-political affiliations retweets.

(iv) show box-plots of retweet probability over all balanced right-leaning users in U_R, while Figures 4(d)(ii), (iii), and (iv) show box-plots of the probabilities over all balanced right-leaning users in U_R who ever retweeted tweet(s) of the corresponding type (computed similarly to Section V. A), for the six types of tweets.

Figures 4 (a) (ii), (iii), and (iv), and Figures 4 (c) (ii), (iii), and (iv) show that, in general, the social transmission hypothesis holds regardless of sentiment type of the tweets: both balanced left- and balanced right-leaning users are more likely to retweet tweets from other sharing the same political affiliation. However, the opposite pattern holds if, for each type of tweet, we consider users ever retweeted tweet(s) of the type. Figures 4 (b) (ii), (iii), and (iv), and Figures 4 (d) (ii), (iii), and (iv) show that, for both balanced left- and balanced right-leaning users who ever retweeted emotional tweet(s) from the other political affiliation are more likely to retweet tweets of the sentiment type from other users of the opposite affiliation. It is interesting that sentiment seems to have a stronger effect than affiliation: in particular, users are more likely to retweet an emotional tweet posted by a user from a different political affiliation than they are to retweet a neutral tweet from the same affiliation.

Again, we obtained results qualitatively same to the above ones when performed the same analysis on subsets of balance users in $U_L^h/U_L^h/U_L^i$s and $U_R^h/U_R^h/U_R^i$s. This further confirms the robustness of the obtained results.

D. Topic of Retweets

Next, we investigate the patterns of retweets with regards to political affiliation and topic. Table II shows the *number* of retweets of users across all 80 topics—here each subfigure shows a particular community relationship, e.g., the top-left figure shows the topic distribution for tweets both received and sent by left-leaning users. Note that all the figures are plotted at the same scale, hence bar-plots in these figures therefore show both actual number and topic distribution of retweets. These figures clearly show that despite the relatively high *probability* of retweeting an inter-affiliations tweet (at least on the part of "balanced" users), numerically, most retweets are of intra-affiliation tweets. This observation holds across topics, and agrees with previous findings by Conover *et. al.* [8].

TABLE II: Number of retweets by different topics and by different political affiliations of the sender and the retweeter

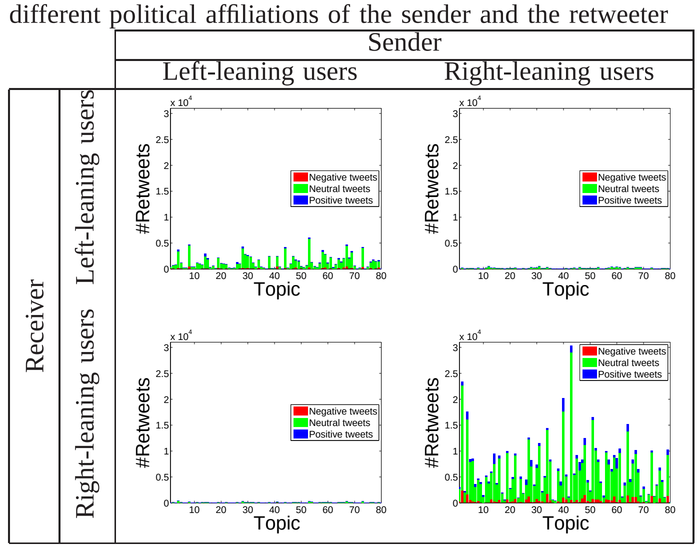

Table III shows some representative topics, each together with representative words. Note that we have manually assigned labels for those topics. The topics in the table contain the most popular topics in the intra-affiliation retweets. It is interesting that *Paul Ryan's Medica plan* (topic 52), *Economic*

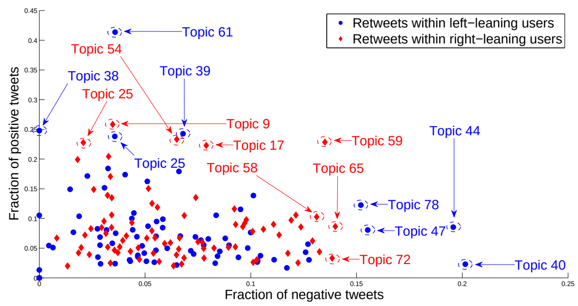

Fig. 5: Fraction of positive and negative of on-topic tweets in intra-political affiliation retweets

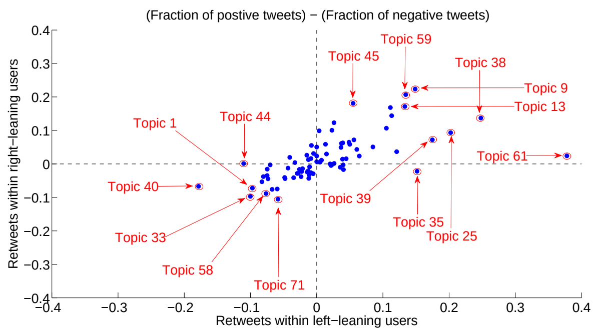

Fig. 6: Difference between fraction of positive tweets and fraction of negative tweets in intra-political affiliation retweets by topics

foreign policies (topic 66), and *Romney's taxes* (topic 7) are most popular for in retweets within left-leaning users; while *Facts and news about Republican* (topic 42), *U.S. consulate attacked in Benghazi* (topic 1) and *Jokes about politics*(topic 39) are the most popular ones among right-leaning users.

We now look into topic distribution of retweets within the two political parties. From upper-left and lower-right figures in Table II, we can easily see that intra-political affiliation retweets in the two affiliations are distributed quite differently. The Pearson rank correlation coefficient of topic "popularities" in the two sets of retweets is only 0.37. (Here, popularity of a topic in a set of retweets is measured by the fraction of retweets about the topic.) The low coefficient implies that left-leaning and right-leaning users have different interests and focuses. This can be seen quite clearly in Figures 5 where each point is plotted so the x-position is the fraction of negative retweets for a political affiliation, the y position is the fraction of positive retweets, and the color indicates the affiliation in question. The figures clearly show that, regardless of the affiliation, neutral retweets are dominant in within-affiliation retweets, while a number of topics are strongly positive or negative for one affiliation or the other, but never for both. (In the figures we exclude topics having fewer than 30 retweets.) This suggests that, within a political affiliation, there are few controversial topics.

With respect to Table III, for retweets within left-leaning users, positive tweets have largest fractions in *Olympic and historic events* (topic 61 - where most tweets are bout victories of U.S teams in Olympic 2012 or about memorial events), *Patriotism and national issues* (topic 38), *Jokes about politics* (topic 39), and *Special days* (topic 25 - where tweets are about reminding events going to happen); while negative tweets have largest fractions in *Previous U.S. presidents* (topic 40 - where tweets are about works by previous U.S. presidents), *Religion and political ideologies* (topic 44), *Liberalism vs conservatism* (topic 47 - where tweets are about statements on the opposite affiliation), and *Personal opinions about the candidates* (topic 78 - where tweets are personal statements on Obama and Romney). On the other hand, for retweets within right-leaning users, positive tweets have largest fractions in *The use of technology in campaign* (topic 9), *Politics as a football game* (topic 17 - where tweets are about talking politics using football terms), *Special days* (topic 25) and *Campaign slogans and report* (topic 59); while negative tweets have largest fractions in *Tax related issues* (topic 65 - where tweets are most about tax policies and job creation), *Romney's tax cut plan* (topic 72), *Campaign slogans and report* (topic 59). This emphasizes that users having different political affiliations are not only interested in different political topics, but also demonstrate different emotions toward these topics.

In Figure 6, we spot out, in within-parties retweets, topics that have fraction of one type of sentiment tweets more than fraction of the other type of sentiment tweets. In the figure, each point is plotted so the x-position is the difference between the fractions of positivite and negative retweets for left-leaning users, and the y position is the similar difference for the right-leaning users. From the figure and with respect to Table III, we can see that most of topics having much more positive tweets than negative tweets in within-affiliation retweets of both the parties (e.g., *The use of technology in campaign* (topic 9), *Republican National Convention* (topic 13), *Special days* (topic 25), *News in big cities & Personal "checkin"* (topic 35), *Patriotism and national issues* (topic 38), *Campaigning among Republicans* (topic 45), *Olympic and historic events* (topic 61)) are also directly about the campaign between the two candidates. On the other hand, most of topics having much more negative tweets than positive tweets in within-affiliation retweets of both the parties are related to the campaining arguments, e.g., *Unemployment rate* (topic 33), *Previous U.S. presidents* (topic 40), *Religion and political ideologies* (topic 44), *Congress meeting on bills* (topic 58), or disaster and crisis, e.g, *U.S. consulate attacked in Benghazi* (topic 1), and *Issues on Syria* (topic 71). Deeper insights from content of tweets about these topics shows that, though sharing the same patterns of emotions toward these topics, users having different political affiliations demostrate different emotions on different aspects of the topics.

VI. Conclusion

In this paper, we examined the effects of sentiment and political affiliation on retweetability of political tweets in Twitter. Our analysis is performed on a large dataset of tweets collected from politics oriented users in U.S. during a long politically active period. Our key findings in this papers confirm that both sentiment and political affiliation have effects on retweetability of political tweets. Moreover, we found that these effects are very different in different type of users: who ever retweeted tweet(s) of certain sentiment type or not. We also obtained the previous work's results about the polarization of political retweets, and further characterized patterns of topics in the retweets with regard to sentiment and political affiliation.

A possible opportunity for future work would be to couple our analyses with more domain-specific and community-specific models of sentiment [3]. Also, while in this paper we analyze user behavior in the aggregate, it would also be of interest to construct models that predict behavior for individual users, and also that jointly predict sharing behavior, network structure, and tweet and retweet polarity.

VII. Acknowledgments

This research is supported by the Singapore National Research Foundation under its International Research Centre

TABLE III: Topic top words

1	U.S. consulate attacked in Benghazi	obama,libya,benghazi,#tcot,americans,president,#obama,american,ambassador,security
9	The use of technology in campaign	million,price,jobs,apple,facebook,campaign,bill,steve,people,iphone
13	Republican National Convention	#gop2012,#dnc2012,#rnc2012,#rnc,convention,speech,#dnc,#tcot,romney,#romneyryan2012
17	Politics as a football game	game,job,ryan,#mlb,patriots,football,#sec,bills,team,#rolltide
25	Special days	day,election,job,labor,politics,days,love,night,people,country
33	Unemployment rate	jobs,unemployment,obama,rate,job,million,economy,numbers,labor,#tcot
35	News in big cities & Personal "Check in"	#nyc,#usa,#chicago,#news,#sandy #us,#dc,#politics,#business,#tcot
38	Patriotism and national issues	#tcot,american,flag,god,#neverforget,country,patriot,obama,#obama,patriots
39	Jokes about politics	#tcot,obama,#p2,#romneyryan2012,#teaparty,#gop,#obama,#tlot,#lnyhbt,romney
40	Previous U.S. presidents	bush,obama,reagan,president,bin,laden,ronald,george,#tcot,romney
44	Religion and political ideologies	obama,muslim,barack,communist,president,america,socialist,american,country,hussein
45	Campaigning among Republicans	senate,candidate,election,republican,#tcot,gop,sen,party,senator,#politics
47	Liberalism vs conservatism	liberal,party,liberals,republican,conservative,republicans,people democrats,women,conservatives
52	Paul Ryan's Medica plan	ryan,paul,medicare,romney,obama,budget,cuts,mitt,obamacare,tax
54	"American dream"	american,paul,country,ryan,rick,bill,video,music,job,love
58	Congress meeting on bills	congress,obama,house,senate,bill,gop,jobs,budget,republicans,president
59	Campaign slogans and report	obama,romney,campaign,#romneyryan2012,president,ryan,#tcot,paul,mitt,ohio
61	Olympic and historic events	#usa,american,gold,olympic,usa,medal,#olympics,olympics,china,country
65	Tax related issues	jobs,people,government,money,obama,pay,tax,taxes,job,govt
71	Issues on Syria	#syria,syria,syrian,#iran,#israel,rebels,turkey,#egypt,war,iran
72	Romney's tax cut plan	tax,taxes,class,romney,middle,cuts,pay,obama,income,rich
78	Personal opinions about the candidates	obama,romney,media,liberal,lies,msnbc,dems,job,ryan,campaign

@ Singapore Funding Initiative and administered by the IDM Programme Office. This research is also supported in part by National Science Foundation Grants: SES-1200154, SES-1022551, IIS-0968295.

References

[1] G. Alec, B. Richa, and H. Lei. Twitter sentiment classication using distant supervision. *Technical report, Stanford University*, 2009.

[2] E. Bakshy, I. Rosenn, C. Marlow, and L. Adamic. The role of social networks in information diffusion. In *WWW '12*, 2012.

[3] R. Balasubramanyan, W. W. Cohen, D. Pierce, and D. P. Redlawsk. Modeling polarizing topics: When do different political communities respond differently to the same news? In *6th ICWSM*, 2012.

[4] J. A. Berger and K. L. Milkman. What makes online content viral? *Social Science Research Network*, 2009.

[5] D. M. Blei, A. Y. Ng, and M. I. Jordan. Latent dirichlet allocation. *J. Mach. Learn. Res.*, 2003.

[6] A. Boutet, H. Kim, and E. Yoneki. What's in your tweets? i know who you supported in the uk 2010 general election. In *ICWSM*, 2012.

[7] J. M. Bristor. Enhanced explanations of word-of-mouth communication: the power of relationships. *Research in Consumer Behavior*, 1990.

[8] M. Conover, J. Ratkiewicz, M. Francisco, B. Gonçalves, A. Flammini, and F. Menczer. Political polarization on twitter. In *5th ICWSM*, 2011.

[9] N. A. Diakopoulos and D. A. Shamma. Characterizing debate performance via aggregated twitter sentiment. In *CHI '10*, 2010.

[10] A. Dobele, A. Lindgreen, M. Beverland, J. Vanhamme, and R. van Wijk. Why pass on viral messages? because they connect emotionally. *Business Horizons*, 2007.

[11] A. Feller, M. Kuhnert, T. O. Sprenger, and I. M. Welpe. Divided they tweet: The network structure of political microbloggers and discussion topics. In *5th ICWSM*, 2011.

[12] J. Golbeck, J. M. Grimes, and A. Rogers. Twitter use by the u.s. congress. *J. Am. Soc. Inf. Sci. Technol.*, 2010.

[13] M. S. Granovetter. The Strength of Weak Ties. *The American Journal of Sociology*, (6):1360–1380, 1973.

[14] L. K. Hansen, A. Arvidsson, F. A. Nielsen, E. Colleoni, and M. Etter. Good friends, bad news - affect and virality in twitter. In J. J. Park, L. T. Yang, and C. Lee, editors, *Future Information Technology*, Communications in Computer and Information Science. Springer Berlin Heidelberg, 2011.

[15] J. A. Hendricks and R. E. D. Jr. *Communicator-In-Chief: How Barack Obama Used New Media Technology to Win the White House*. Lexington Books, 2010.

[16] A. Jisun, C. Meeyoung, G. P. Krishna, and C. Jon. Media landscape in twitter: A world of new conventions and political diversity. In *5th ICWSM*, 2011.

[17] Z. Kunda. Motivated inference: Self-serving generation and evaluation of causal theories. *Journal of Personality and Social Psychology*, 53(4):636, 1987.

[18] M. Lodge and C. Taber. Three steps toward a theory of motivated political reasoning. In *Elements of Political Reason*. Cambridge University Press, London, 2000.

[19] C. D. Manning, P. Raghavan, and H. Schutze. *Introduction to Information Retrieval*. Cambridge University Press, 2008.

[20] M. Pennacchiotti and A.-M. Popescu. Democrats, republicans and starbucks afficionados: user classification in twitter. In *KDD*, 2011.

[21] D. Pierce, D. P. Redlawsk, W. W. Cohen, T. Yano, and R. Balasubramanyan. Social and affective responses to political information. *APSA*, 2012.

[22] D. Redlawsk. Hot cognition or cool consideration? testing the effects of motivated reasoning on political decision making. *The Journal of Politics*, 64(04):1021–1044, 2002.

[23] D. Redlawsk. Motivated reasoning, affect, and the role of memory in voter decision-making. In D. Redlawsk, editor, *Feeling Politics: Emotion in Political Information Processing*. Palgrave Macmillan, 2006.

[24] D. P. Redlawsk, A. J. W. Civettini, and K. M. Emmerson. The affective tipping point: Do motivated reasoners ever get it? *Political Psychology*, 31(4):563–593, 2010.

[25] S. Stieglitz and L. Dang-Xuan. Political communication and influence through microblogging–an empirical analysis of sentiment in twitter messages and retweet behavior. In *45th HICSS*, 2012.

[26] B. Suh, L. Hong, P. Pirolli, and E. H. Chi. Want to be retweeted? large scale analytics on factors impacting retweet in twitter network. In *SocialCom*, 2010.

[27] S. Wu, J. M. Hofman, W. A. Mason, and D. J. Watts. Who says what to whom on twitter. In *WWW '11*, 2011.

[28] W. X. Zhao, J. Jiang, J. Weng, J. He, E.-P. Lim, H. Yan, and X. Li. Comparing twitter and traditional media using topic models. In *ECIR*, 2011.

[29] L. Zhiming, L. Lu, and L. Hong. Determinants of information retweeting in microblogging. *Internet Research*, 2012.

[30] Z. Zhou, R. Bandari, J. Kong, H. Qian, and V. Roychowdhury. Information resonance on twitter: watching iran. In *1st SOMA*, 2010.