Online Fragmentation in Wartime: A Longitudinal Network Analysis of Tweets About Syria, 2011-2013

Deen Freelon,

School of Communication, American University

Marc Lynch

Elliott School of International Affairs and Institute for Middle East Studies, George Washington University

Sean Aday

School of Media & Public Affairs and Elliott School of International Affairs, George Washington University

Abstract

Theorists have long predicted that like-minded individuals will tend to use social media to self-segregate into homophilous enclaves, and that this tendency will increase over time. Many studies have found cross-sectional evidence of such behavior, but very few have been able to directly measure longitudinal changes in the diversity of social media users' information-seeking habits. This is due in part to a lack of appropriate tools and methods. In a step toward remedying this, we propose, justify, and measure a partial method of measuring network homophily in Twitter over time. Drawing on the complete historical record of public retweets posted between January 2011 and August 2013, we demonstrate that Twitter network communities focused on Syria are in general highly fragmented. However, only one of the nine detected network communities clearly exhibits the hypothesized longitudinal increase in fragmentation.

The final version of this preprint manuscript will be published in the *ANNALS of the American Academy of Political and Social Science.*

1. Introduction

One of the most notorious bugbears of the digital age is the possibility that media filtering tools will erect ideological "echo chambers" around their users. Observers disagree on the normative implications of this development—some lament it (Bennett & Iyengar, 2008; Pariser, 2011; Sunstein, 2007) while others celebrate it (Gillan, 2009; Weinberger, 2008). Many predict that ideological fragmentation will grow increasingly extreme over time: Sunstein, for one, believes that at "some time in the future" (2007, p. 1) advanced filtering tools will enable people to perfectly screen out all content not of personal interest" (p. 3). Pariser's (2011) "filter bubble" describes invisible information-presentation algorithms that conform to users' interests over time. And Bennett and Iyengar predict that "over time, avoidance of disagreeable information may become habitual so that users turn to their preferred sources automatically no matter what the subject matter" (2008, p. 724).

Research findings on fragmentation have been mixed: they generally conclude that online political communication is fragmented but also that users are not completely isolated from dissenting views (Adamic & Glance, 2005; Conover et al., 2011; Garrett, 2009b; Hargittai, Gallo, & Kane, 2008). Most of these studies have been cross-sectional, conveying nothing about how people's information consumption habits change over time. To our knowledge, only Hargittai et al. (2008) have even attempted to study the hypothesized longitudinal fragmentation process by directly observing online user behavior. Their data are nearly ten years old as of this writing, and their analysis includes only three data points (most of the study's empirical content is cross-sectional).

Building on Hargittai et al.'s work, our study investigates online fragmentation within retweets about Syria over the course of 32 months. We believe this is the largest-scale digital trace-based study of ideological fragmentation yet attempted, incorporating some 25 million tweets. In addition to our empirical findings, we also contribute several methodological innovations.

1.1. Political fragmentation online

This paper is principally concerned with the proportions of political information people receive that conform to their personal beliefs. The successful use of digital filtering tools to increase this proportion goes by many names, including fragmentation, polarization, selective exposure, and cyberbalkanization. The end-states of this process are known variously as echo chambers, filter bubbles, and "individualized reality construction" (Bennett & Iyengar, 2008, p. 722). Imputed causes include homophily, the general desire of like-minded individuals to congregate; and the development of digital filtering technologies which allow such congregation. The empirical task at hand, therefore, is to measure changes in people's attention patterns over time to ascertain whether fragmentation actually increases.

Existing research in this area falls into three categories: experiments, surveys, and trace-data analyses. We discuss the first two only briefly as they are less directly relevant to our work than the third. Surveys generate subjective impressions of participants' media habits by asking about them. These studies typically find some evidence of selective exposure and fragmentation, although participants are by no means completely isolated from alternative views (Garrett, 2009b; Stroud, 2008). Some results indicate that while people prefer attitude-consistent information, they tend not to avoid attitude-inconsistent information (Garrett, 2009b; Johnson, Zhang, & Bichard, 2011). Experimental research generally supports this conclusion (Garrett, 2009a; Graf & Aday, 2008). Individuals have also reported increasing ideological homogeneity in their media diets over time (Stroud, 2008, 2010).

Analyses of online trace data such as hyperlinks, retweets, and social media “friend” connections offer an important complement to more traditional fragmentation research. Trace-based studies compensate for their lack of fine-grained demographic controls with their ability to unobtrusively document real-world communication behavior. This is especially important given that people tend to be less than truthful about their media consumption habits when asked directly (Thompson, 2014). Many such studies analyze the online networks of the American conservative and progressive movements, and their results generally agree with their traditional counterparts: people form densely interconnected communities with ideological allies but interact only occasionally with adversaries (Adamic & Glance, 2005; Colleoni, Rozza, & Arvidsson, 2014; Hargittai et al., 2008; Smith, Rainie, Shneiderman, & Himelboim, 2014).

While it is impossible to review all trace-based fragmentation research here, our literature search identified only Hargittai et al. (2008) as attempting a systematic longitudinal analysis. One reason for this lack of empirical attention may be the lack of software packages to preprocess social media traces appropriately. The Python module used to preprocess this study’s data was developed by the first author and is available on an open-source basis (Freelon, 2014). We hope future researchers will use it to build upon the results we present here.

1.2. The Syrian civil war and Twitter

Our data consist of tweets about Syria, which has been embroiled in a civil war since mid-2011. This conflict began as a popular uprising in the midst of the Arab Spring upheavals in Tunisia, Egypt, Libya and Yemen. Peaceful protests in Syria were met with brutal regime reprisals, however, shocking Syrians and the world. Images of devastation of Syrian cities and a massive flow of refugees galvanized international attention. By early 2012, the uprising had evolved into an increasingly militarized insurgency and regional proxy war, with rebels backed by the Gulf states and Turkey and the regime supported by Iran and Hezbollah. Syria became a major focus for Islamist movements across the Middle East, particularly in the Gulf, with massive fundraising and consciousness raising campaigns and a high degree of media attention.

The Syrian civil war has generated vigorous discussion on Twitter. This has been fueled in part by the massive influx of Arabic-speaking users onto Twitter beginning in mid-2011 after the early Arab Spring successes (Arab Social Media Report, 2013; Lynch, Freelon, & Aday, 2014). Though research on the topic is still in its infancy, media reports have already begun to offer some evidence for the use of Twitter not only to discuss Syria but to organize political and humanitarian campaigns and to recruit fighters to the cause (Burke, 2014; Karam, 2013).

1.3. Fragmentation in the Syria Twittersphere

Most fragmentation research has been conducted in and on stable, democratic countries. Because Syria is neither, and because the discourse about it is transnational in nature, it is not clear how well hypotheses derived from such research will apply to it. Nevertheless, they provide the best basis available for research on online political fragmentation in a war zone. But unlike in the US, it is not clear which constituencies will be represented in the Syria Twittersphere. Hence our first research question:

- RQ1: Which communities appear most often in the Syria Twittersphere over time?

Our sole hypothesis derives directly from prior work on how fragmentation is expected to change over time:

- H1: Communities of like-minded individuals will become increasingly fragmented with time.

Although not specifically suggested by prior research, we also take this empirical opportunity to connect fragmentation with several plausible predictors other than time:

- RQ2: What factors other than time will predict changes in fragmentation?

2. Methods

Our methodological approach can be summarized as a series of six steps which we describe in detail due to their novelty. They are:

1. Purchasing Twitter data from vendor
2. Converting raw data into edge list format for network analysis
3. Detecting network communities among Twitter users on a per-month basis
4. Qualitatively labeling network communities
5. Tracking the progress of persistent communities over time
6. Measuring each persistent community's level of fragmentation over time

2.1. Twitter data purchase

We began by purchasing from the certified Twitter data vendor Topsy all public tweets containing the term "Syria" or "سوريا" (the Arabic term for "Syria") posted between January 1, 2011 and August 31, 2013. This query yielded a total of 57,704,424 tweets across 32 months along with the standard metadata the Twitter API offers. While these certainly do not constitute all tweets about Syria, they almost certainly capture a substantial proportion of them using parsimonious, intuitive, and easily replicable acquisition parameters.

2.2. Edge list conversion

After importing the raw data into a MySQL database, we generated a CSV file for each month containing only the full text of each tweet and the author's username. We then used a custom Python script to isolate the retweeter-target pairs from all retweets in the dataset. We analyze retweets here not because they always imply endorsement or authority, but because we consider them to be indicators of attention. In other words, retweets are sufficient but not necessary to imply that users felt that the content was important enough to pass on. Users who retweet the same sources can thus be said to share similar information preferences (Aragón, Kappler, Kaltenbrunner, Laniado, & Volkovich, 2013; Conover et al., 2011; Hanna et al., 2013). Follow/friend networks have also been used to map ideological affinity (e.g. by Golbeck & Hansen, 2011), but these can be manipulated easily and inexpensively (Hockenson, 2012) and are difficult to retrieve through the Twitter API. We considered analyzing Twitter mentions as well but found that we could not consistently classify the resulting network communities.

Our Python script generated a series of text files in edge list format in which each line consists of a retweeter/target pair separated by a comma. This allowed us to create 32 distinct directed networks, one for each month's retweet activity. These networks were based on the 44.1% of tweets in our dataset classified as retweets (N = 25,383,627).

2.3. Community detection

Next came the task of identifying the largest network communities (subgraphs) of densely interconnected users. For this we used the Louvain method of community detection (Blondel, Guillaume, Lambiotte, & Lefebvre, 2008), one of the few that can accommodate very large networks. Louvain creates mutually exclusive network communities that maximize internal density and external sparsity, which accords with our criteria for fragmentation. Most retweet networks (including ours) exhibit long-tail distributions, and when applied to these Louvain creates a small number of very large communities and a large number of very small communities (the smallest of which represent one user retweeting another once). For our data it generated over 1,000 communities in the least active month and over 40,000 in the most active month. Only the largest communities map onto identifiable offline constituencies—most are too small and non-representative for quantitative analysis. Therefore we decided to retain only the ten most populous communities in each network. These 320 communities, which represent 0.09% of the 363,706 total communities detected, collectively account for over 70% of all retweeting/retweeted users and over 78% of all retweets.[1]

2.4. Qualitative community labeling

The Louvain method assigns each community a unique number which distinguishes them from one another but conveys nothing about their underlying populations. Therefore the second author, an expert in Middle East politics, inductively categorized each of the 320 network communities. He did so by seeking unifying themes shared between prominent users and their most popular messages as well as verifiable information about their identities. This is comparable to the qualitative labeling methods used in other studies of social media network communities (e.g. Aragón et al., 2013; Etling, Kelly, Faris, & Palfrey, 2010). After the initial inductive labeling process was completed, we decided to place each community into one of three ideological categories: pro-Asad, anti-Asad, or neither/unclear. As the Results section will show, this typology greatly simplifies our fragmentation analysis.

2.5. Tracking communities over time

The labeling process gave each of the 320 communities both an idiosyncratic and an ideological label, but we still needed a way to connect the 32 separate networks. We conjectured that some communities would display a sustained interest in Syria over time while others would appear only once or sporadically. We defined a community as *persistent* if it could be consistently tracked across adjacent months—by definition, changes in fragmentation can only be observed in persistent communities. The goal of our community tracking process was to identify which communities in which months constituted different iterations of the same persistent community.

To do this, we first defined community similarity in terms of membership: the more members two communities share in common, the more similar they are. We further noted that because retweeting behavior exhibits long-tailed distributions, "elite" individuals who garner disproportionate shares of retweets contribute disproportionately to their community's integrity. Removing these elite users would sever the connections between many of their communities' nonelites. Therefore, if we are trying to identify multiple iterations of the same persistent community over time by examining overlaps in community membership, users who are retweeted more should count more toward community similarity than users who are retweeted less. With these insights in mind, we define the similarity of two network communities A and B as the Jaccard index weighted by received retweets (weighted indegree):

$$\text{weighted\_jaccard(A, B)} = \frac{\sum A_i \cap B_j}{\sum A_i \cup B_j}$$

That is, the weighted Jaccard index between A and B is the sum of the weights of the membership intersection of A and B divided by the sum of the weights of the membership union of A and B.[2] The subscript i denotes the retweet weights of community A while the subscript j denotes B's weights. The index ranges between zero and one, with zero indicating that two communities share no users in common and one indicating that their memberships are identical. To make this equation computationally tractable for our largest networks, we applied the weighted Jaccard to only the top 1% of users by received retweets. These "one-percenters" define the communities to which they belong in much the same way that popular media programs define their fandoms (Burwell & Boler, 2008).

The community-tracking process proceeded stepwise using adjacent-month pairs, beginning with January-February 2011 and ending with July-August 2013. Weighted Jaccards were computed between each inter-month pair of communities, and the best match was recorded for each month pair. We set a weighted Jaccard threshold of 0.3, which we found to effectively distinguish between strong and weak community matches. Best matches meeting or exceeding this value were retained as valid, while those below it were dropped. We set an additional condition that all persistent communities had to meet or exceed the 0.3 threshold over at least five consecutive months. These conditions yielded a total of nine persistent communities, which will be described in greater detail in the Results section.

2.6. Measuring changes in community fragmentation

With nine persistent communities identified, the final methodological step was to select a per-community unit of measurement for fragmentation. Previous studies have used the proportion of within-community ties for this purpose (Adamic & Glance, 2005; Conover et al., 2011). We modified this metric slightly to account for the fact that many of our communities support or oppose Asad. So rather than only considering proportions of internal and external ties, for each persistent community in each month we added the proportion of internal ties to the proportion of ties to allied communities and divided that sum by the total number of ties involving the original community. We call this quantity the community's proportion of homophilous ties, or PHT. PHT ranges between zero and one, with the former indicating no internal or allied ties and the latter indicating all the community's ties are internal or allied. Each PHT calculation incorporates all ten communities categorized for the given month (see sections 2.3 and 2.4) For example, if in a given month 60% of persistent anti-Asad community A's ties are internal, 30% are distributed among six other anti-Asad communities, and the remaining 10% are distributed among the remaining three non-allied communities, A's PHT is 0.9. The more allied communities there are in a given network, the higher PHT is likely to be. Tracking each persistent community's PHT over time allows us to answer our research questions and test our hypothesis.

3. Results

RQ1 asked about which communities would meet our persistence criteria. We label and present a brief description of the nine that did so, including the Twitter handles of some of their most prominent users. Seven of them were anti-Asad, one was pro-Asad, and one was neither.

- **English-speaking journalists.** This community brought together international English-language journalists and news outlets. It recurred in 28 of the 32 months, more often than any other community by far. The community was dominated by media outlets (e.g. @AP, @france24, @nytimesworld, and @finanicaltimes), but a few individual reporters and columnists made appearances as well (e.g. @nickkristof, @DanWilliams, @JakeTapper).

Many of its most-retweeted tweets included headlines and links to Syria-related news stories. It was the only persistent community we classified as neither pro- nor anti-Asad.

- **Syrian revolutionaries.** This community combined Syrian opposition online news sources (such as @shaamnews, @revolutionsyria, @newssyrrev, @syriacampaigns, @freesyria24) and well-known Syrian activists (e.g. @razaniyat, @syria_horra, @anonymousssyria, @bsyria, @alrifai1) along with key Gulf-based Arabic news sources (such as @alarabiya_bk, @ajalive). This should be seen as the core of the Syrian opposition online before the shift towards armed insurgency and the rise of more Islamist and sectarian forces in the spring of 2012.
- **Salafi Islamists.** This community brought together Sunni Salafi Islamists, often but not exclusively from Saudi Arabia, with a strongly sectarian and anti-Asad orientation. This community included prominent Salafi news sites (such as @wesaltv and @ahlalsunna2) and personalities (@aref_12, whose account is entitled "the response to the rejectionists", a derogatory sectarian term for Shi'ites). It included journalists and accounts from across the Gulf, including Qataris (@a_althbah), Kuwaitis (@ahmed_m_alfahad) and Sunni Bahrainis (@bahrainforsyria, @bahraindefence).
- **Kuwait (three communities).** Kuwait emerged as one of the key epicenters of fundraising and support for the Syrian opposition, often with an Islamist flavor. These communities includes well-known Kuwaiti Islamists (such as @nabilalawadhy, @drshafialajmi) along with members of the Kuwaiti parliament (@altabtabie), Syrian relief campaigns (@kuwwithsyria, @kowaitiya), public intellectuals (@dralshayji, @eng_alkandari), and news sites (@akhbaralkuwait, watannews). The attention to Syria in Kuwait was so strong that three distinct communities emerged, centered around

different personalities and media. One concentrated around popular retweet sites, often featuring jokes and sex-focused accounts, which perhaps created high-density clusters when they occasionally touched upon Syria.

- **Free Syrian Army.** In late 2012, another Syrian opposition cluster emerged centered around support for the mainstream Free Syrian Army insurgency. This community included some of the same opposition accounts as the earlier Syrian Revolutionaries community (@revolutionsyria, @shaamnews, @yathalema) but also included new accounts associated with the armed struggle (@free_syrian_army, @syrianfalcon11).
- **Pro-Asad.** The sole pro-Asad community brought together outright supporters of Bashar al-Asad's regime (@syriaman20043, @elpadrino1982, @syrian_truth) with self-proclaimed independent anti-imperialist personalities (@asadabukhalil) and news sources (@almayadeennews, @almanarnews). Lebanese and Bahraini accounts, presumably Shi'a, also frequently appeared in this cluster, suggesting a cross-regional sectarian community.
- **Saudi Arabia**. The extremely active Saudi Twittersphere also paid close attention to Syria. This community features a large number of relatively unknown Saudis, whose discourse tended to support Islamist and sectarian pro-opposition/anti-Asad frames.

To answer H1, which predicted that the communities would grow increasingly fragmented over time, we created longitudinal scatterplots of each community's PHT values (Figure 1). Missing data points on each figure indicate months in which the community did not manifest according to our criteria. Due to the small number of data points for most communities, we formally test the relationship between time and PHT only for the two largest communities, English-Speaking

Journalists (*n* = 28) and Syrian Revolutionaries (*n* = 18). For the rest, we rely on rough visual assessments of the figures.

[Insert figure 1 here]

We tested several linear and nonlinear regression models of the relationship between time and PHT for the two largest communities, using standard error of the estimate as a measure of goodness of fit. For English-Speaking Journalists, the linear model provides the best combination of simplicity and goodness of fit (SE_{est} = .042, p < .001, *B* = .011). As Figure 1 shows, PHT generally increases with time, supporting H1. (The alternative models that fit the data well showed similarly positive slopes.) The linear model also fits the Syrian Revolutionaries community fairly well (SE_{est} = 0.055, p = 0.032, *B* = 0.005), although not as well as it fits English-Speaking Journalists. A cubic model offers slightly improved goodness of fit at the expense of increased complexity (SE_{est} = .045, p = 0.008, B_1 = -0.051, B_2 = 0.005, B_3 = 0.000). The latter model suggests that PHT increases and decreases over time in this community, contradicting the linear model's clear support for H1.

The remaining communities manifest too rarely to attempt to model. Nevertheless, if we avoid generalizing too strongly, we may tentatively observe that most communities connect overwhelmingly to allies when they do appear, with PHT values mostly over 0.9. So although we cannot identify longitudinal trends here, we can conclude that the data are loosely consistent with H1, as they depict highly fragmented end-states rather than progressive fragmentation over time. The Pro-Asad community is a slight exception, having a lower mean PHT than the others. This is almost certainly a consequence of the fact that pro-Asad voices were extremely rare in our data—in no month was there any more than one such community, and that community was usually relatively small. Pro-Asad users may have engaged with pro-rebel users in an antagonistic manner—future research should explore this possibility.

Our final research question explores the role of two potential predictors of fragmentation in the English-Speaking Journalist and Syrian Revolutionary communities. For each month, we computed the proportion of communities in which the primary language was Arabic as well as the number of news articles about Syria in major global news outlets.[3] We included these two variables as predictors of PHT in two pairs of linear regression models for each community: one model included time as a predictor and one excluded it. For English-Speaking Journalists, time is the only significant predictor when it is included; when excluded, proportion of Arabic communities is positively related to PHT ($B = 0.491$, $p < 0.001$). In the Syrian Revolutionary models, nothing is significant when time is excluded, and when it is included it is marginally significant with a very weak effect size ($B = 0.009$, $p = 0.05$).

4. Discussion

This study makes two major contributions to the study of online fragmentation, one theoretical and the other methodological. Theoretically we have shown that fragmentation within the Syria Twittersphere was indeed high in the aftermath of the Arab Spring. The seven anti-Asad communities shared a mean of 93% of their retweets among themselves, while the sole pro-Asad community shared 81%. Interestingly, the only community to unambiguously exhibit the expected increasing fragmentation pattern was the English-Speaking Journalist community, which began with a PHT of 0.566 but peaked at 0.917 near the end of the time period. Explaining these phenomena proved difficult—only two communities contained sufficient data points for even a basic statistical analysis. Results suggest that the proportion of communities in one’s first language may be related to fragmentation but that volume of news coverage is unrelated.

These findings prompt us to revisit some of the basic assumptions about how fragmentation works. Many theoretical accounts posit vague longitudinal increases in fragmentation, but they are not clear on exactly how researchers should pinpoint the starting point after which fragmentation is expected to increase. While we lack the data to answer this definitively, one possibility is that certain long-term news events (the Arab Spring in this case) are prominent enough to temporarily reconfigure paths of communication between parties who ordinarily would not engage with one another. A common news agenda may bring people together for a time, but eventually they may seek out more congenial news sources. Many of the Arabic-language communities clustered around particular media outlets which carried distinctive ideological messages (for instance, the Lebanese pro-Asad television station @almayadeennews often anchored the Pro-Asad community, while the anti-Shi’ite salafi TV station @wesal_tv anchored the Salafi community). Similarly, the English-Speaking Journalist community attended less and less to Arabic and ideological voices as the Arab Spring faded from Western public consciousness. Of course, Syria represents a very specific case—an autocratic Middle Eastern country destabilized by civil war—and some of these potential explanations may not generalize to other cases.

Another potential explanation stems from the fact that most of our Arabic-speaking users were on the same side of the conflict. Seven of our nine communities were anti-Asad and most differentiated themselves along either national or identity lines. Kuwaitis and Saudis, for instance, had largely similar approaches to Syria but tended to interact with their co-nationals more intensely than with a pan-Arab community. It may be significant that the one major exception, the Syrian Revolutionary community, at first engaged broadly across communities but broke down in early 2012 while the more localized and embedded networks gained traction.

Methodologically, we contribute a novel measure of network fragmentation, proportion of homophilous ties (PHT), which generalizes to multi-community contexts the tie-proportion measure often used in two-community contexts such as the American political left and right. When a given dataset contains more than two network communities, per-community fragmentation cannot be measured simply by the proportion of internal and external ties unless the number of communities and the number of ideological positions are identical. That is, ties between different communities of the same political orientation should be counted as homophilous rather than cross-cutting, especially with algorithmically-detected communities. PHT's major disadvantage is that it cannot be computed automatically like traditional network measures of segregation. Rather, the ideological stances of the network communities must be classified first, a process that usually requires qualitative inspection by subject-matter experts and/or trained coders.

Our study's limitations bear acknowledging. We believe our reliance on retweets as an indicator of informational diversity is justified, but it is only one indicator among many. Other ways people can signal attention on Twitter include "liking" tweets and sharing URLs without retweeting, and these may be less prone to ideological clustering. Moreover, people receive information from many sources, and fragmentation within Twitter does not necessarily imply fragmentation across all information channels. Other online services are more of a technical challenge to study than Twitter, but the goal of fully understanding online fragmentation demands that we make the effort. Finally, despite our empirical timeframe being longer than most Twitter studies, it was not long enough to provide sufficient data points to model most of our communities statistically. More data would have helped, but as data volume increases, so do the computational resources required to obtain, store, and analyze it.

Limitations notwithstanding, this study contributes meaningfully to the online fragmentation literature. By directly measuring online fragmentation over time, it helps bring empirical practice into accord with theory. Just as importantly, it offers a replicable metric of fragmentation that can be applied and compared across a wide variety of contexts.

Notes

[1] Modularity values for these Louvain partitions ranged between 0.55 and 0.88, where “values greater than about 0.3 appear to indicate significant community structure” (Newman, 2004, n.p.).
[2] Our method is similar to Greene, Doyle, and Cunningham’s (2010), though we developed ours independently. The chief differences are that our method is designed for directed networks, uses a weighted Jaccard index, and analyzes only each network’s top users.
[3] To obtain this information, we searched the Factiva “Major News and Business Publications” list for articles containing “Syria” in the headline or lead. The full list of included publications is available at Factiva.com.

References

Adamic, L. A., & Glance, N. (2005). The political blogosphere and the 2004 US election: divided they blog. In *Proceedings of the 3rd international workshop on Link discovery* (pp. 36–43). New York, NY: ACM.

Arab Social Media Report. (2013). *Twitter in the Arab Region*. Dubai, United Arab Emirates: Dubai School of Government. Retrieved from http://www.arabsocialmediareport.com/Twitter/LineChart.aspx?&PriMenuID=18&CatID=25&mnu=Cat

Aragón, P., Kappler, K. E., Kaltenbrunner, A., Laniado, D., & Volkovich, Y. (2013). Communication dynamics in twitter during political campaigns: The case of the 2011 Spanish national election. *Policy & Internet*, *5*(2), 183–206. doi:10.1002/1944-2866.POI327

Bennett, W. L., & Iyengar, S. (2008). A new era of minimal effects? The changing foundations of political communication. *Journal of Communication*, *58*(4), 707–731.

Blondel, V. D., Guillaume, J.-L., Lambiotte, R., & Lefebvre, E. (2008). Fast unfolding of communities in large networks. *Journal of Statistical Mechanics: Theory and Experiment*, *2008*(10), P10008.

Burke, S. (2014, April 22). Syria's Foreign Fighters Use Social Media As Recruitment Tool. Retrieved from http://www.nbcnews.com/news/world/syrias-foreign-fighters-use-social-media-recruitment-tool-n86901

Burwell, C., & Boler, M. (2008). Calling on the Colbert Nation: Fandom, politics and parody in an age of media convergence. *Electronic Journal of Communication, 18*(2). Retrieved from http://www.cios.org/www/ejc/EJCPUBLIC/018/2/01845.html

Colleoni, E., Rozza, A., & Arvidsson, A. (2014). Echo Chamber or Public Sphere? Predicting Political Orientation and Measuring Political Homophily in Twitter Using Big Data. *Journal of Communication*, *64*(2), 317–332. doi:10.1111/jcom.12084

Conover, M. D., Ratkiewicz, J., Francisco, M., Goncalves, B., Flammini, A., & Menczer, F. (2011). Political polarization on twitter. In *Proc. 5th Intl. Conference on Weblogs and Social Media* (pp. 89–96). Barcelona, Spain: AAAI.

Etling, B., Kelly, J., Faris, R., & Palfrey, J. (2010). Mapping the Arabic Blogosphere: Politics and Dissent Online. *New Media & Society*, *12*(8), 1225–1243. doi:10.1177/1461444810385096

Freelon, D. (2014). Twitter Subgraph Manipulator (Version 3). Python. Retrieved from https://github.com/dfreelon/TSM

Garrett, R. K. (2009a). Echo chambers online?: Politically motivated selective exposure among Internet news users. *Journal of Computer-Mediated Communication*, *14*(2), 265–285.

Garrett, R. K. (2009b). Politically motivated reinforcement seeking: Reframing the selective exposure debate. *Journal of Communication*, *59*(4), 676–699.

Gillan, K. (2009). The UK anti-war movement online. *Information, Communication & Society*, *12*(1), 25–43.

Golbeck, J., & Hansen, D. (2011). Computing political preference among twitter followers. In *Proceedings of the 2011 annual conference on Human factors in computing systems* (pp. 1105–1108). New York, NY, USA: ACM. doi:10.1145/1978942.1979106

Graf, J., & Aday, S. (2008). Selective Attention to Online Political Information. *Journal of Broadcasting & Electronic Media*, *52*(1), 86–100. doi:10.1080/08838150701820874

Greene, D., Doyle, D., & Cunningham, P. (2010). Tracking the Evolution of Communities in Dynamic Social Networks. In *2010 International Conference on Advances in Social Networks Analysis and Mining (ASONAM)* (pp. 176–183). doi:10.1109/ASONAM.2010.17

Hanna, A., Wells, C., Maurer, P., Friedland, L., Shah, D., & Matthes, J. (2013). Partisan Alignments and Political Polarization Online: A Computational Approach to Understanding the French and US Presidential Elections. In *Proceedings of the 2Nd Workshop on Politics, Elections and Data* (pp. 15–22). New York, NY, USA: ACM. doi:10.1145/2508436.2508438

Hargittai, E., Gallo, J., & Kane, M. (2008). Cross-ideological discussions among conservative and liberal bloggers. *Public Choice*, *134*(1), 67–86.

Hockenson, L. (2012, December 15). Fake Twitter Followers: An Easy and Dangerous Game. Retrieved from http://thenextweb.com/twitter/2012/12/15/fake-followers-an-easy-game-but-not-worth-the-risk/

Johnson, T. J., Zhang, W., & Bichard, S. L. (2011). Voices of Convergence or Conflict? A Path Analysis Investigation of Selective Exposure to Political Websites. *Social Science Computer Review*, *29*(4), 449 –469. doi:10.1177/0894439310379962

Karam, Z. (2013, October 19). Syria's Civil War Plays Out On Social Media. Retrieved from http://www.huffingtonpost.com/2013/10/19/syria-social-media_n_4128360.html

Lynch, M., Freelon, D., & Aday, S. (2014). *Blogs and Bullets III: Syria's Socially Mediated Civil War*. Washington, DC: United States Institute of Peace.

Newman, M. E. J. (2004). Fast algorithm for detecting community structure in networks. *Physical Review E*, *69*(6). Retrieved from http://pre.aps.org/abstract/PRE/v69/i6/e066133

Pariser, E. (2011). *The filter bubble: What the Internet is hiding from you*. New York: Penguin.

Smith, M. A., Rainie, L., Shneiderman, B., & Himelboim, I. (2014). *Mapping Twitter Topic Networks: From Polarized Crowds to Community Clusters*. Pew Internet & American Life Project. Retrieved from http://www.pewinternet.org/2014/02/20/mapping-twitter-topic-networks-from-polarized-crowds-to-community-clusters/

Stroud, N. J. (2008). Media Use and Political Predispositions: Revisiting the Concept of Selective Exposure - Springer. *Political Behavior*, *30*(3), 341–366. doi:10.1007/s11109-007-9050-9

Stroud, N. J. (2010). Polarization and Partisan Selective Exposure. *Journal of Communication*, *60*(3), 556–576. doi:10.1111/j.1460-2466.2010.01497.x

Sunstein, C. R. (2007). *Republic. com 2.0*. Princeton, NJ: Princeton University Press.

Thompson, D. (2014, June 17). Why Audiences Hate Hard News—and Love Pretending Otherwise. Retrieved June 19, 2014, from http://www.theatlantic.com/business/archive/2014/06/news-kim-kardashian-kanye-west-benghazi/372906/

Weinberger, D. (2008). Echo Chambers= Democracy. In A. Fine, M. Sifry, A. Rasiej, & J. Levy (Eds.), *Rebooting America* (pp. 32–37). New York: Personal Democracy Press.

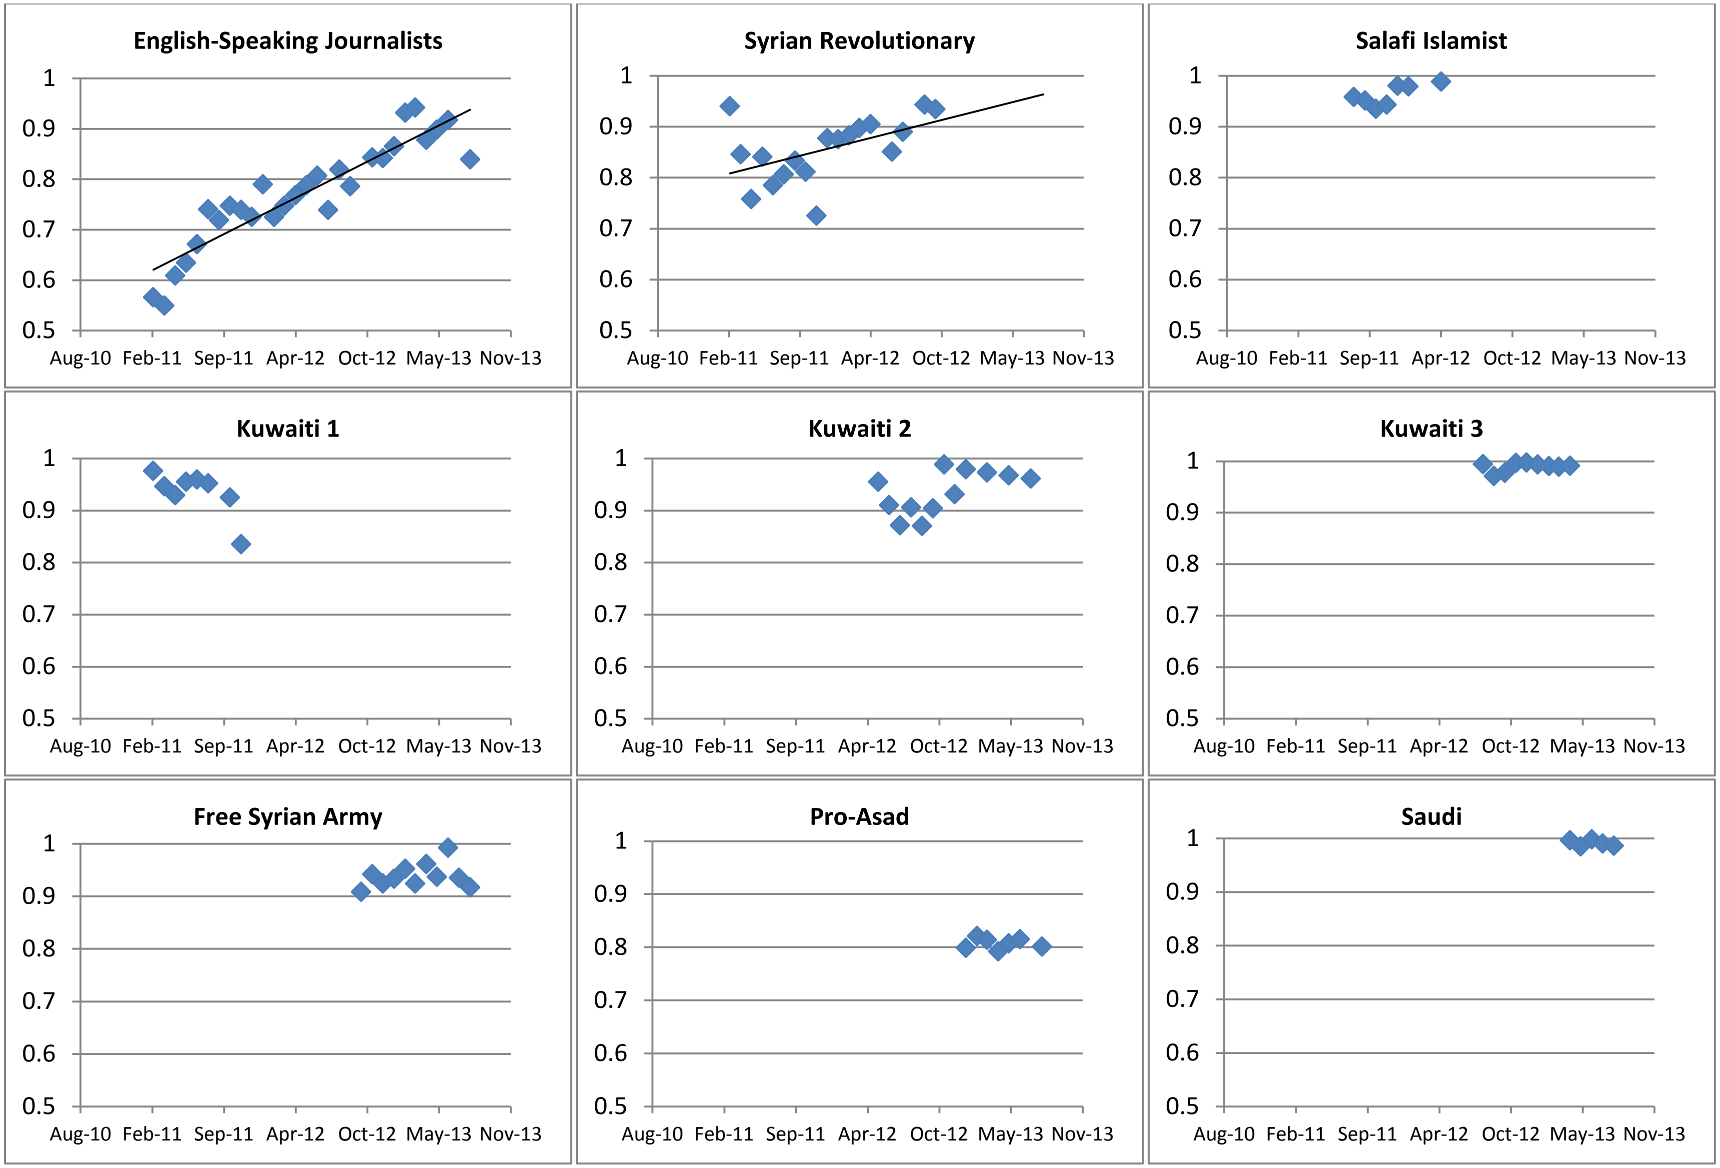

Figure 1. Proportions of homophilous ties over time for nine persistent network communities. All y-axes represent PHT.