Explaining the Emergence of Political Fragmentation on Social Media: The Role of Ideology and Extremism

Jonathan Bright

Oxford Internet Institute, University of Oxford, Oxford, UK

This article is a systematic large-scale study of the reasons driving political fragmentation on social media. Making use of a comparative dataset of the Twitter discussion activities of 115 political groups in 26 countries, it shows that groups that are further apart in ideological terms interact less, and that groups that sit at the extremes of the ideological scale are particularly likely to have lower patterns of interaction. Indeed, exchanges between centrists who sit on different sides of the left–right divide are more likely than connections between centrists and extremists who are from the same ideological wing. In light of the results, theory about exposure to different ideological viewpoints online is enhanced.

Introduction

A key strand of contemporary research in online political communication concerns what could be referred to as the "fragmentation" thesis: the idea that online conversations about politics are typically divided into a variety of groups, and that this division takes place along ideological lines with people only talking to those who are ideologically similar. Concerns about this type of fragmentation were voiced in some of the earliest theoretical work on the Internet (e.g., Dahlberg, 2007; Papacharissi, 2002; Sunstein, 2009; Van Alstyne & Brynjolfsson, 1999). Numerous empirical studies in a variety of Internet discussion contexts such as forums (Hill & Hughes, 1998), blogs (Adamic & Glance, 2005; Hargittai, Gallo, & Kane, 2008; Lawrence, Sides, & Farrell, 2010) and most recently on social media (Aragón, Kappler, Kaltenbrunner, Laniado, & Volkovich, 2013; Barberá, 2014; Colleoni, Rozza, & Arvidsson, 2014; Conover, Gonçalves, Ratkiewicz, Flammini, & Menczer, 2011; Conover, Gonçalves, Flammini, & Menczer, 2012; Feller, Kuhnert, Sprenger, & Welpe, 2011; Gaines & Mondak, 2009; Garcia, Abisheva, Schweighofer, Serdült, & Schweitzer, 2015; Gruzd & Roy, 2014; Himelboim,

McCreery, & Smith, 2013; Quattrociocchi, Scala, & Sunstein, 2016; Yardi & Boyd, 2010), have since found evidence that at least some degree of fragmentation exists.[1] These patterns have concerned many theorists of democracy, who have argued that exposure to a diverse range of viewpoints is crucial for developing well informed citizens (Gentzkow & Shapiro, 2010, 2) who are also tolerant of the ideas of others (Nunn, Crockett, & Williams, 1978). By contrast, exposure to only like-minded voices in a kind of "echo chamber" may contribute towards polarization towards the extremes (Sunstein, 2002; Warner, 2010).[2] As social media networks become more important for shaping political viewpoints and exposing people to information (Bakshy, Messing, & Adamic, 2015; Bright, 2016), the social relevance of the fragmentation thesis will continue to grow.

However, while the literature on the subject is rich, empirical studies that focus on the macro level structure of networks as a whole (and the extent to which they are fragmented) have thus far have been largely descriptive (although much analytical research exists at the individual level, which is treated more fully below). The major reason for this is that the vast majority of studies that have looked at the network level have been single country ones (for examples of single country studies see, Aragón et al., 2013; Colleoni et al., 2014; Conover et al., 2011; Garcia et al., 2015; Gruzd & Roy, 2014; Himelboim et al., 2013). This provides little scope for studying variation as each single country network typically contains only a handful of individual groups. Hence the field as a whole has not yet addressed the question of what factors might drive more or less fragmentation in discussion networks. The aim of this study is to take a first step towards remedying this deficit, by seeking to explain variation in levels of political fragmentation through a large scale multi-country study of online political discussion on Twitter.

The article is structured in the following way. Section one offers a clearer definition and conceptualization of political fragmentation, and builds up theory and hypotheses about the factors that might explain its emergence, focusing particularly on ideological variations between different groups in the network in question. Section two outlines the method employed, explaining the collection of data from Twitter concerning political discussions in a variety of different EU countries during the European Parliament elections in 2014, and describing how key concepts are operationalized. Section three sets out the results, with evidence showing that ideology does indeed appear to drive variation in the intensity of fragmentation, with the role of extremism found to be particularly important. The results are then discussed in light of their potential consequences for the circulation of ideas online.

Explaining political fragmentation in online discussion networks

For the purposes of this article, online discussion networks are spaces such as forums, blogs and social media sites where people can engage in discussion and exchange messages. In conceptual terms, political fragmentation occurs in such a network when, during discussions about politics, participants in the network begin to converse more with others who are ideologically similar than they do with others who are ideologically different (Garcia et al. 2015, 47). As patterns of exchange between ideologically similar groups drop (relative to patterns of communication within these groups) fragmentation becomes more severe, and conversations start to resemble what have been called "echo chambers," whereby people simply hear their own viewpoint repeated back to them (Sunstein, 2002).

Political fragmentation could hence be defined as a characteristic of a discussion network as a whole: it is more or less fragmented depending on the amount of division within it. However, a discussion network may contain a number of different ideological groups, each of which communicates to a varying extent with other groups: some pairs of groups may have quite lively exchanges, whilst others may hardly converse at all. As a network level measure might disguise this important local variation,

in this article the focus is placed on explaining fragmentation observed between different *pairs* of ideological groups within a wider discussion network.

Theoretical explanations for political fragmentation on social media have, thus far, focussed on individual level behavioral characteristics (see for example the review of research presented by Colleoni et al., 2014, pp. 318–320). Although the aim of this article is to focus on macro structure rather than individual behavior, it is nevertheless worth exploring these individual factors because the macro structure of discussions on social media are a result of the behavior patterns of individual social media users.

The most commonly identified individual mechanism is the tendency towards "homophily" (Barberá, 2014), an impulse to form social ties with others who are similar to oneself in some way (McPherson, Smith-Lovin, & Cook, 2001). Such a tendency means that political fragmentation will naturally emerge in online discussions as people connect to others with similar views. Closely related to homophily is the concept of "selective exposure" (Knobloch-Westerwick & Meng, 2009), a phenomenon whereby people select information or sources they already agree with whilst filtering out others (Garrett, 2009a). If we consider online discussions as a source of information, then the selective exposure mechanism will produce similar results to the homophily mechanism, as people select themselves into online discussions with which they already agree. A further mechanism identified is the tendency of individuals to moderate their opinions into line with what they perceive the group norm to be (Garcia et al., 2015), or at least to keep quiet if they believe themselves to be outside this norm (Scheufele, 2000). This mechanism means that existing groups of association should become more homogeneous over time (or at least they will appear to).

These micro level mechanisms are in general well known and have been widely tested (see, for example, the review of research presented in Knobloch-Westerwick, 2015). However, by themselves they offer little reason to expect why levels of observed fragmentation might *vary* at higher levels of observation such as between pairs of groups, which is the major interest of this article. Such variation could only be caused by factors that also vary between pairs of groups, and in so doing contribute to either enhancing or moderating these micro-level mechanisms.

The most obvious factor that varies between pairs of groups is their respective ideological positions. Ideology could have several effects on the micro level mechanisms identified above, and hence the extent to which pairs of groups interact. Firstly and most obviously, it is theoretically plausible that as the ideological distance between groups increases, the fragmentation between them also increases. As groups drift further apart in ideological space, members of each group are likely to perceive those from the other group as ideologically less similar, and are hence less likely to make connections to them (following the homophily mechanism). They are also likely to disagree with their arguments more, and hence more likely to selectively ignore their communication (following the selective exposure mechanism). This line of theory leads to the first hypothesis tested in this article:

H1: As the ideological distance between groups increases, they will interact less

Secondly, it is also possible that the type of ideology has an impact on the tendency of the group to fragment. Groups on the left and right of the political spectrum attract support from different socio-demographic strata; and it is possible that socio-demographic factors have an influence on predisposition towards homophily and selective exposure (Garrett, 2009b). There are also qualitative differences between the two sides of the left–right spectrum which might further accentuate these mechanisms: individuals from groups which sit on either side of the division (even if not very far apart in terms of the scale of the axis itself) are nevertheless likely to perceive each other as dissimilar, and hence may have a tendency to talk to each other less than groups which are a similar distance apart but on the same side of the scale (some evidence for this idea is provided in Feller et al. 2011). These

ideas are supported by descriptive works on political fragmentation which have repeatedly shown that groups on different sides of the left–right political division appear to have different internal patterns of communication, even if no consensus has emerged on the direction of the relationship: some studies have found the right-wing end of the spectrum to be more densely connected (Conover et al. 2011; Hargittai et al., 2008; Warner, 2010), whilst others have found the reverse (Barberá, 2014); one study even found evidence for both conclusions using different measures (Colleoni et al., 2014). These works hence lead to the second hypothesis to be tested:

> H2: When groups are from different sides of the left–right divide, they will interact less

Finally, the "extremism" of a group's ideology may also play a role. As Stroud has shown (2010), individuals with attitudes at more extreme ends of the ideological scale show more pronounced tendencies towards selective exposure than those in the middle, a result attributed to the increased certainty these individuals typically have in their beliefs (see also Johnson, Bichard, & Zhang, 2009; Sunstein, 2009; Wojcieszak, 2009). Elements of this line of thinking can also be seen in "hostile media" theory, where people with strong pre-existing opinions are more likely to perceive alternative viewpoints as biased, and hence ignore or filter them out (K. Kim, 2011; Y. Kim, 2011). This dynamic could also be self-reinforcing, as discussion with like-minded individuals can also lead to the polarization of attitudes towards ideological extremes, which in turn stimulates further fragmentation (Huckfeldt, Morehouse Mendez, & Osborn, 2004; Myers & Lamm, 1975). This branch of theory leads to the final hypothesis tested in this article:

> H3: As the ideology of a group gets more extreme, it will interact less with other groups

While this article places a major focus on ideology as a driver for political fragmentation, there are also a variety of other factors that are worth considering as control variables. First, the overall size of a political grouping might make a difference: larger political groupings might be less likely to communicate with smaller ones as they might be perceived as less worthy of consideration. This idea is supported by Aragón et al. who found diverging communication patterns between small and large political groupings (Aragón et al. 2013). Furthermore, the status of different political groups in the wider political system also differs: some will be related to political parties which are incumbent in government at any given time, whilst others may be in opposition. Previous research has shown that incumbent political forces often make less use of online democratic opportunities (Herrnson, Stokes-Brown, & Hindman, 2007): it may be that online groups which are related to incumbent political forces are hence also less connected to the rest of the discussion network as a whole.

Method

This article aims to collect a sufficiently large sample of pairs of groups within discussion networks in a variety of countries such that the hypotheses identified above about political fragmentation can be tested at large scale and that generalizable conclusions can be drawn. This large scale data collection presents several novel methodological challenges, as the majority of studies thus far have looked at single country networks and have made use of techniques which do not easily scale up to multiple countries. In this section, the methods employed will be described, beginning with a description of the data collection approach, then moving on to introduce a measure of fragmentation between pairs of groups in a discussion network, before finally describing the independent variables used in the study.

Table 1 Example of Edge Formation in a Political Discussion Network on Twitter

User	Tweet	Edges Created
@blueparty	.@stephanie this am delivered a petition calling for prison reform	@blueparty → @stephanie
@john	@blueparty @stephanie didn't we have a vote on this? Less than 6 years ago? And the answer was no?	@john → @blueparty @john → @stephanie
@paul	@john @blueparty @stephanie Nope. We had a vote specifically on reducing sentences not on reform in general	@paul → @john @paul → @blueparty @paul → @stephanie

Data collection

Data for the study is taken from Twitter, a social media platform which allows users to distribute short messages (known as "tweets") to groups of people who have chosen to receive them (known as followers) and to engage in discussion with other users. Twitter is the only major social media network that is both frequently used for political discussion in a variety of countries and that makes its data generally available for research purposes, and it has hence been widely used in previous research on the structure of online discussion networks.

The structure of conversation on Twitter is made visible through "mentions," which represent the inclusion of the "username" of another Twitter user in a tweet. When one user mentions another in the tweets they create, this indicates a directed connection between them (though, of course, the mentioned user may not reciprocate the connection). This practice of mentioning occurs in a variety of situations. A user may write an original message about another user, and hence include their name in it. A user may also respond to a message by another user, in which case this other user is mentioned automatically in the response. Finally, a user may "retweet" a message produced by another user (which involves rebroadcasting this message to all of their followers): again, this retweet will include a mention of the user who originally created the message (as well as any other users referenced in the original message). Collecting a sample of tweets from Twitter, and observing patterns of mentioning found within them, allows observation of the structure of discussion networks which appear on the platform. The "nodes" in these networks are users of Twitter, and directed "edges" are created between users whenever one mentions another.

An illustrative example of this type of discussion network can be found in Table 1. The short series of tweets described in the table is real, though the Twitter usernames have been changed and the text altered slightly to preserve the privacy of the participants. The first tweet is an original tweet by the user "@blueparty" (usernames on Twitter are prefixed with the @ symbol). This tweet mentions another user (@stephanie), hence an edge is created in the network from @blueparty to @stephanie. This tweet attracts a response from @john, who thus mentions @blueparty (and @stephanie) in their response, creating two further edges. Another user (@paul) then responds to @john's tweet, mentioning all the people who had previously been mentioned, and thus creating three further edges.

When studying networks of conversation on Twitter in this way, a key initial question is how to select the tweets to be included in the analysis: clearly, a different subset of tweets may result in a different conversational structure. The strategy most commonly deployed in the literature on political fragmentation on social media (applied in, for example, Aragón et al. 2013; Conover et al. 2011; Feller et al. 2011; Garcia, Mendez, Serdült, & Schweitzer, 2012; Himelboim et al., 2013) is the collection of tweets containing politically relevant "hashtags" (a word preceded by the symbol "#" which is used to label a tweet as being related to a specific topic), often around the time of a key moment in national

politics such as an election. This data collection approach allows the observation of communication activity within what have been described as "ad-hoc publics" (Bruns & Burgess, 2015) that form through the discussion activities of people using these hashtags when they address comments at each other. The fact that the hashtags are politically relevant is designed to ensure that the conversations observed concern politics.

In this article, the 2014 European Parliament elections were chosen as the key event for observation. Making use of the European Parliament elections is useful because it creates the potential to observe political discussion networks in a wide variety of countries (there are 28 member states in the EU) that are all focussed on the same key event. However, the decision was made not to follow a hashtag-based data collection strategy, because it is not something that easily scales up to multiple countries: hashtags are inevitably language specific, and often country specific (referring to particular events within national political life) and hence choosing equivalent hashtags in many different country contexts would be difficult. Instead, the article makes use of a complete list of all available official Twitter usernames of major political parties and party leaders in all 28 EU member states (collected by the *euandi* project—see Garzia, Trechsel, De Sio, & De Angelis, 2015). The data collection involved harvesting all tweets which were either authored by these users or mentioned one of these users (only one username was used per political party, with preference given to the username of the party leader; if that was not available, the username of the party account itself was employed instead).[3] Table 1 again provides an impression of a typical conversation captured by this method: a central party account creates a tweet, and then a discussion forms around that tweet between other Twitter users.

Of course, while this approach is effective in terms of eliminating potential country specific bias, the limitations of the strategy when compared to the hashtag-based approach should also be acknowledged. The data collected represents only a partial account of all the political discussion which took place on Twitter during the moment of the election: tweets which related to the election but which did not mention a party account will not be collected. Of course, hashtag-based collection also suffers from this problem, but arguably to a lesser extent.

The data collection window ran from the 11 May to 10 June, during which time tweets were collected from the Twitter streaming Application Programming Interface (API), which is a web service that allows structured data collection from the Twitter platform. This period was chosen because it provided a sample of data from both before and after the elections took place, which makes it possible to test the sensitivity of results to the election itself (the precise day of the election varied between countries, and indeed some elections took place over several days: however all of them were held between the 22nd and 25th of May). In total 1,426,620 tweets were observed in the time window which were either created by or made reference to at least one of the party usernames in the list.

The focus of this article is explaining variation in the level of fragmentation between different pairs of groups within an overall discussion network. Hence each observation in the dataset is a pair of discussion groups that formed around two party-specific Twitter usernames (the calculation of the extent of fragmentation between the pair is described below). Discussion groups were paired within countries (explaining patterns of cross-country communication would be interesting but is outside the scope of this article), with as many pairs as possible formed in each country (for example, if a country had three party-specific Twitter usernames, X, Y and Z, then three pairs were formed: {X,Y}, {X,Z} and {Y,Z}). However, the extent to which individual party-specific usernames were mentioned on Twitter varied considerably and not all parties that competed in the election had enough data to be included in the final analysis. As the aim was to generate the largest possible number of observations, rules for inclusion were made as permissive as possible: parties were included if they had an official account that was mentioned at least once in the observation window; if they had an available ideology score (as described below) and if they were competing in a country that had at least one other party that also fulfilled these criteria (such that a pair of parties could

be formed). In total 115 parties had enough data to form part of the analysis, from 26 different countries (out of the 28 EU member states, only Croatia and Latvia had no representation in the dataset). These 115 parties were formed into 237 different pairs which are the observations in the dataset.

Measurement of fragmentation between pairs of discussion networks

As described in the theory section, this article seeks to measure (and explain) the extent of fragmentation between these pairs of discussion groups. Measurement of fragmentation was another area that presented a methods challenge. The primary approach in the literature is to label individuals participating in conversation networks as belonging to one ideology or another through some kind of manual or automatic content analysis. Then, various network metrics can be calculated that determine whether the actual structure of conversation observed is related to the way in which users are labeled. However, this is another technique that does not scale well to the context of a large number of observations in a wide selection of countries making use of a variety of different languages.

This problem is addressed through the development of a metric of fragmentation, F, which can be calculated for a pair of discussion groups without the need for a prior fixed assignment of individuals as belonging to one group or another. The measure is inspired by the work of Guerra, Meira, Cardie, and Kleinberg (2013), who developed a measure of network polarization, P, based on comparing the activity of "boundary" nodes (that is, nodes who communicate with members of both groups within the network) to the activity of "internal" nodes who communicate solely within one group; the measure also has considerable similarities with Krackahrdt and Sterns's E-I index (Krackhardt & Stern, 1988), a measure which seeks to capture the extent to which a given group is more or less externally facing in its patterns of communications, and which has previously been used in studies of online political fragmentation (Hargittai et al., 2008).

Calculation of F proceeds as follows. To begin, for each of the 237 pairs of political parties, a directed network was created. This pair level network consists of all individuals who mentioned at least one of the two party-specific usernames belonging to that pair of political parties (it also includes the party users themselves). These individuals are the nodes in the network, and directed edges between the nodes were created when one individual mentioned another during the data collection period in the form described in Table 1 above (if individuals mentioned each other multiple times, then multiple edges were created). Figure 1 provides a simplified fictitious example of such a network. Two party nodes, labeled X and Y, sit at the center of the network, with six further nodes (labeled 1 through 6) around them. The direction of
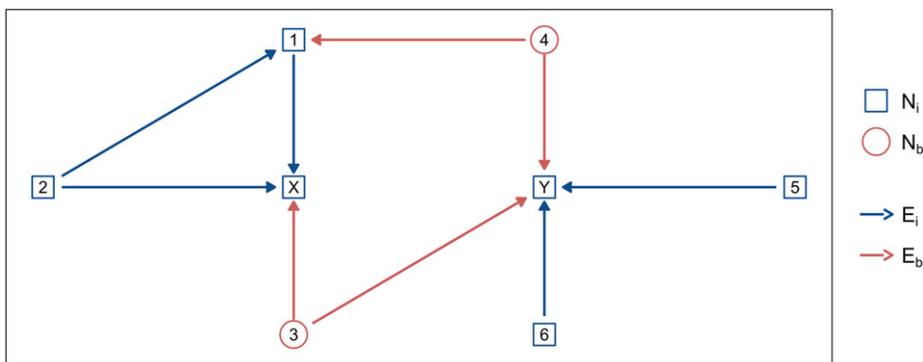

Figure 1 Calculation of F on a toy network.

the arrows indicates one node mentioning another during the observation window. For example, node 2 mentioned node X, hence they are connected by an arrow.

Two sets of party nodes are then created, N_x and N_y, which are the users who mentioned each party (the set also includes the party itself). In the example in Figure 1, $N_x = \{X,1,2,3\}$ and $N_y = \{Y,3,4,5,6\}$. Note that nodes can be members of both sets if they mentioned both parties. These two sets of party nodes allow the definition of a set of boundary nodes, N_b, which includes any nodes in N_x which mentioned a node in N_y and vice versa. In this case: $N_b = \{3,4\}$. This in turn allows the definition of a final set of purely internal nodes, N_i, which represents all nodes which are not members of N_b. In this case: $N_i = \{X,Y,1,2,5,6\}$.

This classification of nodes permits the classification of the edges in the network on the basis of their point of origin: edges which initiate from internal nodes (in N_i) are included in the set of internal edges, E_i, and edges which initiate from boundary nodes (N_b) are included in the set of boundary edges, E_b. In this case: $E_i = \{[1{\rightarrow}X],[2{\rightarrow}X],[2{\rightarrow}1],[5{\rightarrow}Y],[6{\rightarrow}Y]\}$ whilst $E_b = \{[3{\rightarrow}X], [3{\rightarrow}Y],[4{\rightarrow}1],[4{\rightarrow}Y]\}$. As remarked above, in the real networks used in the study, the same edge may appear multiple times if one user mentioned another multiple times during the observation window (this allows F to be weighted for the volume of communication between nodes). The exact metric used to measure fragmentation between a pair of groups, F, is then given by the following formula:

$$F = \frac{|E_i| - |E_b|}{|E_i| + |E_b|}$$

Where $|E_i|$ and $|E_b|$ represent the cardinality (or size) of the sets E_i and E_b. The potential values of F run from -1 to 1, with higher numbers indicating greater degrees of fragmentation. In the example in Figure 1, $|E_i| = 5$ and $|E_b| = 4$, hence F evaluates to $(5 - 4) / (5 + 4)$, that is $1/9$. This low positive number indicates a mild degree of fragmentation (i.e., a slightly higher amount of internal discussion than between group discussion).

While use of F fits well to the task at hand, it is nevertheless worth reflecting on the likely consequences of its key limitation: that it does not have an independent means of assigning users to groups, and is hence based purely on the structure of communications (rather than their content). This is likely to bias the measure in favor of observing more fragmentation than actually exists: a discussion group that does not contain any connections to another group would appear highly insular using this method, however, such a group may well actually contain lively patterns of internal disagreement from individuals of competing ideologies (who have simply not mentioned any other political parties during the observation window). While this is a limitation, it is also one that is, arguably, defensible in this case. First, while the structure of communication patterns does not reveal the whole story, it does show something: if, for example, two party usernames are connected by a strong set of boundary nodes, the conclusion that fragmentation is low ought to be correct, regardless of the content of messages flowing between them. Second, any bias present in the measure is the same for all pairs of parties, and as the main aim is to explain variation between party pairs (rather than offer a point estimate of fragmentation levels), any bias which applies uniformly is not too damaging.

Figure 2 provides an illustrative visualization of how F relates to overall network structures in real data that was collected during the observation window. Four different party pairs were selected for the image at a range of different levels of F. The figure illustrates how the F score increases as the amount of boundary edges (the set E_b) decreases relative to the amount of non-boundary edges (the set E_i). At lower levels of F the boundary starts to dominate the overall graphic and indeed the dual group structure becomes harder to see. The metric, in other words, seems to capture the basic
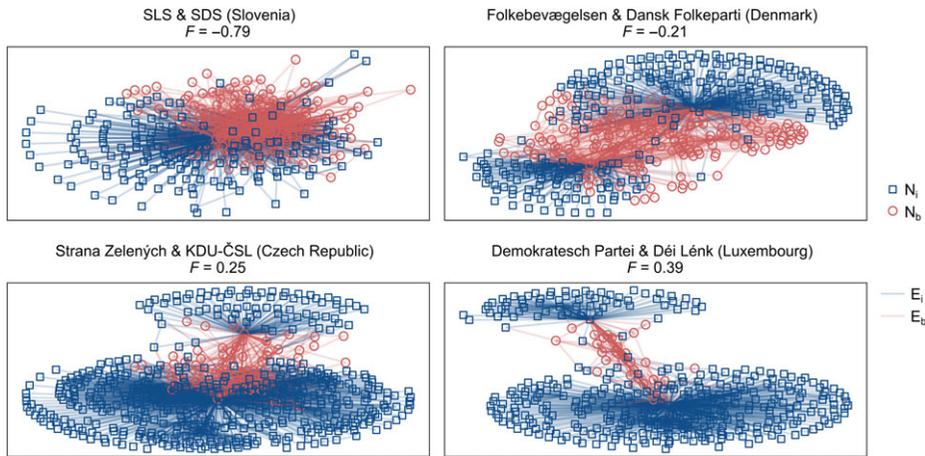
Figure 2 Party pairs at different levels of *F*.

intuition behind fragmentation. However, the figure also suggests that *F* may be sensitive to the ratio of sizes of the two groups (for example, the highest *F* score in the figure relates to a pair comprising a relatively large and relatively small group). The sensitivity of models to this factor is hence tested below.

Independent variables

The operationalization of independent variables will now be described. The main independent variable in the study is the ideology of political groups. In order to operationalize ideology, parties in the dataset were classified on a left–right scale (see e.g., Castles & Mair, 1984). Left-wing parties tend to favor policies relating to economic redistribution and equality; whilst right-wing parties favor policies relating to individual liberty and the free market. Parties on the extremes of either scale favor more radical and far-reaching visions of these policies. Such scales are certainly simplifications of the complexity of political life. Yet they are also widely used in both popular discourse and academic research as a way of understanding politics, and represent the simplest way of classifying a wide variety of groups. The particular classifications for each party were taken from the Parlgov dataset (Döring, 2012), which uses a 10-point scale ranging from 0 (extreme left) to 10 (extreme right).

Positioning parties on an ideological scale allowed calculation of the three main independent variables that relate to each of one of the three hypotheses. First, the ideological distance between parties could be measured, which is the absolute value of the difference between their two ideology scores. Second, parties were coded as belonging to either the left or right of the political scale. Finally, the "extremism" of a party's ideology could be measured, which is simply the absolute value of the distance of their ideology score from the center-ground score of 5. This variable, obviously, takes a value between 0 and 5 at the level of the individual party; the total extremism within a pair of parties can therefore range from 0 to 10.

Three control variables were used. The size of a political grouping was given by the number of votes won in the 2014 election, measured as a percentage of total votes. This allows specification of the size difference between groups, with an expectation that a greater difference in size would result in more fragmentation. The political status of the party in question was also recorded (i.e., whether it was incumbent in government or in opposition). This allows a specification of the relationship

Table 2 Descriptive Statistics for Pairs of Parties

Numerical Variables	Min	Max	Mean	SD
Observed Nodes	5	175,299	10,629	24,777
Observed Edges	14	1,152,127	77,539	169,947
F	−1.0	1.0	−.25	.56
Ideological Distance	.00	8.30	2.98	1.86
Extremism	.90	8.30	4.13	1.42
Size Difference	.00	40.6	9.85	8.87

Categorical Variables		N	%
Incumbency Status	Incumbent–Incumbent Pair	40	17%
	Incumbent–Opposition Pair	130	55%
	Opposition–Opposition Pair	67	28%
Left–Right Status	Left–Right Pair	137	58%
	Left–Left/Right–Right Pair	100	42%
Total Observations		237	

between parties: whether they are both incumbent in government, both in opposition, or whether it is an incumbent–opposition pair. Finally, the ratio of tweets observed for the pair of parties was measured, which is simply the number of tweets observed for the smallest party divided by the number observed for the largest party. This controls for potential effects on the F score caused by imbalances in size between the two groups in question, as described above.

Analysis

Initial descriptive statistics on the data collected are presented in Table 2. There is wide variation in the degree of fragmentation (F) between different groups, with the observed data covering the full range of possible values (though the data is skewed to the right). There is also considerable variety in the absolute sizes of pairs in terms of number of participants (nodes), with some having less than 100 and some having tens of thousands. The smaller groups present a concern in that, when the number of contributions is small in absolute terms, the extent to which the observed level of fragmentation is affected by individual nodes and edges is increased: hence there may be potential for small group F scores to be noisy. This point will be revisited in the analytical section. It is worth highlighting that there are some correlations between the independent variables: in particular as might be expected pairs of groups that are from different sides of the left right axis, or that consist of an incumbent party and an opposition party, are typically also further apart in ideological terms. The measure of ideological distance is also strongly correlated with the extremism score.

The analysis makes use of a series of multilevel models, whereby party pair observations are nested in countries and country differences themselves are modeled as random effects.[4] This type of model is appropriate in this circumstance considering that the observations are drawn from different countries and that variation in the overall extent of fragmentation might well be expected between different political systems. A potential concern about this modeling strategy is of course the fact that the same party may appear in multiple pairs of observations, hence creating a dependence across some data points. In this regard, however, it is worth emphasizing that as nodes are not assigned strictly to one group or another, what is an "internal" node when considering one pair of parties may instead become

Table 3 Multilevel Models Explaining Fragmentation (*F*) for Party Pairs. Observations are Nested in Country Groups.

	1.1	1.2	1.3	2.1	2.2	2.3
Ideology	.17**	.10†	.29***			
Extremism				.14**	.10*	.22***
Left–Right Pair	−.26*	−.16	−.45**	−.01	−.01	−.02
Size Difference		.01	.13*		.01	.14*
I–O Pair		.14	.20		.13	.16
O–O Pair		.10	.20		.05	.10
Tweet Ratio		−.28***	−.34***		−.28***	−.32***
Observations	237	237	161	237	237	161
Marginal R^2	.01	.07	.21	.01	.07	.21
Conditional R^2	.77	.81	.75	.78	.81	.76

All numerical variables are standardized

†p < .10; *p < .05; **p < .01; ***p < .001

I-O: Incumbent–Opposition Pair (as compared to Incumbent–Incumbent Pair)

O-O: Opposition–Opposition Pair (as compared to Incumbent–Incumbent Pair)

a "boundary" node when a different pair of parties is considered. Hence, as parties do not make a fixed contribution to either the internal or boundary node sets in each pair they are part of, the problem of dependence amongst observations is not too great.[5]

Model fit was assessed through the use of marginal and conditional R^2 values (Nakagawa & Schielzeth, 2013). The marginal R^2 gives a measure of the amount of variance explained by the individual level predictors while the conditional R^2 gives the total amount of variance explained by the full model: comparing the two also reveals the amount of variance explained by the country level differences. The main results are presented in Table 3. This table contains two sets of models, one addressing the effect of ideological distance between parties (models 1.1 to 1.3), and one addressing the effect of extremism (models 2.1 to 2.3). These variables are analyzed in separate models because, as highlighted above, they are highly correlated.

Model 1.1 investigates the effect of two key variables of interest, ideological distance and mismatching left–right pairs, on the fragmentation index. Model 1.2 includes a number of theoretically relevant control variables, whilst model 1.3 limits the observations to only party pairs with at least 1,000 nodes, to check whether small pairs are influencing the overall results. The results for ideological distance in model 1.1 go in the expected direction and are statistically significant: as parties get further apart in ideological space, communication between them decreases. This result becomes only borderline significant in model 1.2, but is significant again in model 1.3. Hence overall there is reasonable support for the idea that increasing ideological distance between parties causes increasing fragmentation (Hypothesis 1). The term for a mismatching left–right pair is also significant in models 1.1 and 1.3. However, it is also unexpectedly positive, indicating that pairs from different sides of the left–right divide communicate more than those on the same side (directly contradicting Hypothesis 2). This suggests, for example, that (holding ideological distance constant) a given center-left party would communicate more with a center-right party than it would with an extreme-left party.

One way of explaining this counterintuitive finding is that it supports Hypothesis 3, which states that as the ideology of a group gets more extreme, it will interact less with other groups. This idea is

tested in models 2.1 to 2.3, which look at the impact of extremism rather than ideological distance. The extremism term (which is the sum of both extremism scores for the two parties in the pair) is statistically significant in all models. This shows that as one or both parties in a party pair tend towards the ideological extremes, communication patterns between them drop and fragmentation increases. This provides strong support for Hypothesis 3. The term for a left–right mismatch is, meanwhile, statistically insignificant in these models. This means that overall there is little support for Hypothesis 2 (which specified that parties from different sides of the left–right divide ought to interact less).

It is worth briefly commenting on the control variables in models 1.1–2.3. There is some evidence that an increasing size difference between parties leads to weaker patterns of communication, though this result was only found in the models restricted to at least 1,000 nodes (1.3 and 2.3). There was no evidence that incumbency or opposition status made a difference in any model. Meanwhile, the term for the ratio of tweets between parties was significant in all models, confirming the sensitivity of F to this factor. It is also worth highlighting the R^2 scores. Conditional R^2 is quite high in all cases, whilst marginal R^2 is quite low. This indicates a large amount of variation at the country level. Marginal R^2 rises to its highest levels in the models which operate on the subset of party pairs with at least 1,000 nodes. As suggested above, this may indicate that F scores for very small party pairs are more unreliable.

Although the analysis above shows that extremism matters, it does not reveal precisely what effect it has. In particular, it does not show whether it is extremist parties that connect less, or whether centrist parties are less willing to connect to their extremist counterparts. One way of addressing this is by looking at how edges within the boundary between pairs of centrist and extremist parties are distributed between the two party node sets (N_x and N_y). If a greater proportion of boundary edges initiates in N_x than in N_y, then it seems reasonable to conclude that the group around party X is making more external connections than party Y.

Of the 237 observations in the data sets, 99 represent a pair of parties containing one centrist and one extremist (defining centrists as parties with extremism scores of less than or equal to 2.5, which is the middle of the scale, and extremists as parties with scores above 2.5). This set of 99 parties was unfolded into 198 party-in-pair observations, and the proportional contribution of each one to the boundary of their pair was measured. The level of contribution was found to be negatively correlated with a party's extremism score (R = −.20), providing evidence that it is indeed extremist parties which communicate less. This difference is statistically significant when estimated within a multilevel model with parties nested in countries, and robust to the inclusion of all the control variables in Table 2 above. However, there is of course an interdependence problem in the data used to estimate this model (as each pair of parties generates two observations). Hence the statistical significance of the results should be considered suggestive rather than conclusive. Interestingly, the same exercise was conducted for parties involved in left–right pairs, and no statistically significant differences were found. This undermines the idea that left- and right-wing parties have fundamentally different communication patterns.

A potential objection to these findings is that it is not clear how sensitive they are to alternative means of sampling data from Twitter. The networks employed make use of both normal conversational tweets and retweets: yet retweets have been shown to exhibit different network structures in previous research (Conover et al., 2011). Furthermore, the data collection window encompasses both the before and after election period, yet conversation patterns have been shown to be different before and after this type of critical juncture (Garcia et al., 2015). Finally, making use of "weighted" networks creates the potential for small numbers of very active users, who may send a lot of messages and hence create very strong connections, to have a disproportionate influence on the results. To test robustness of the results to these alternative means of sampling data, duplications of models 1.1–2.3 were estimated with five different sampling methods: using no retweet data; using only retweet data; using only data from before the election; using only data from after the election; and using an unweighted

network. The key results concerning ideological distance and extremism were robust to all of these different specifications with the exception of the retweet data, where terms still pointed in the same direction but were no longer statistically significant. This supports the idea that conversational tweets and retweets may produce quite different network structures (though it is also worth noting that retweets constituted only 27% of the full dataset: the fact that the data was more limited could also explain the discrepancy found here). Overall, however, these further models offer good support for the idea that the results are robust to different means of sampling data.

Discussion and conclusion

This article has sought to provide a first large scale analytical treatment of the reasons behind the emergence of political fragmentation on social media, defined as the widely observed phenomenon of online networks self-organizing into groups along ideological lines. Based on a novel dataset drawn from Twitter, and working at the party-pair level, it has shown that fragmentation varies within networks, with some pairs of groups communicating more than others. It found good evidence supporting the idea that the ideological distance between groups plays a role in explaining this variation, with particular evidence that extremist parties communicate less. Indeed, centrist parties (even from different sides of the left–right divide) were more likely to communicate than a centrist and extremist party from the same side of the left–right divide. Furthermore, when individual party dynamics were considered in the second step of the analysis, no evidence was found that left- and right-wing parties have different communication patterns, contradicting a range of pieces of previous research. These results were robust to a variety of alternative model specifications.

These findings have implications for the theory of online political fragmentation and indeed online discussion more generally. Several points can be made. One of the major worries advanced by theorists of political fragmentation has concerned exposure to alternative viewpoints, with worries that the development of online "echo chambers" will lead to people hearing their own views repeated again and again. The evidence presented here, by contrast, shows that lots of communication does occur between different ideologies, especially across the left–right divide but within the center-ground. In fact, the real area of separation appears to occur with people who hold extreme ideologies, who become separated both from people of other viewpoints and even people who hold more moderate versions of their viewpoint. This may indicate that the most important factor is, as Stroud has suggested (2010), the certainty with which people hold beliefs, rather than ideological differences between individuals.

Another factor concerns some of the other determinants of fragmentation. There was some evidence that the position of a political party within the political system changes the way they interact with technology: larger parties seemed to generally communicate less with other groups. It is intriguing to find that parties which are more politically successful offline are also typically more disconnected online, and it is significant because it shows that online fragmentation is not purely a result of decisions made by individuals online; the offline context has an impact. However the evidence also undermines the idea that left- and right-wing parties have fundamentally different communication patterns, at least on Twitter, and hence that qualitative differences between these two types of ideology generate more or less fragmentation. In the future, any single country studies which report such differences should be aware that they may not reflect a general trend.

It is worth concluding by highlighting issues that could not be addressed in this study, thus indicating potential directions for further research. Two points stand out. First, the measures used were "naïve" in that they paid no attention to the content of messages (or indeed the type of person sending the message). Second, as the data in the article is a snapshot, it is unable to address temporal

dynamics, which research has suggested might play a role in fragmentation. As research continues in these areas, we will continue to understand more about what drives the emergence of political fragmentation in online discussion networks.

Acknowledgments

Notes

1 Though many studies have also shown at least some evidence of cross ideological exposure (Wojcieszak & Mutz, 2009), whilst others have argued that the Internet makes a positive overall contribution to the heterogeneity of political networks (Brundidge, 2010; K. Kim, 2011; Y. Kim, 2011).

2 Though others have countered that too much cross ideological exposure can lead to political ambivalence, and hence at least a certain amount of fragmentation is necessary for political action (Dahlberg, 2007; Mutz, 2002).

3 The analyses reported below were also repeated without the tweets created by the official party accounts themselves: the results produced were virtually identical.

4 Standard diagnostic tests were applied to all reported models: residual plots were checked for normality, models were re-estimated without high leverage observations, the linearity of relationships between predictor and response variables was checked, and models were also estimated with a normalized dependent variable. These diagnostics did not provide any reason to doubt the results of the models.

5 One way of measuring the magnitude of interdependence is to calculate, for all pairs of observations in the data sets which share a common party, the proportion of nodes which are in the same position between the pairs (e.g., whether the same nodes are internal in both pairs). The results of this exercise show that on average 46% of nodes change position when a given pair of observations is considered. There is hence considerable variation between observations even when they share a common party.

References

Adamic, L., & Glance, N. (2005). *The political blogosphere and the 2004 US election: divided they blog.* Proceedings of the 3rd Workshop on Link Discovery, 36–43.

Aragón, P., Kappler, K., 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.

Bakshy, E., Messing, S., & Adamic, L. (2015). Exposure to ideologically diverse news and opinion on Facebook. *Science*, 348(6239), 1130–1132. doi:10.1126/science.aaa1160

Barberá, P. (2014). Birds of the same feather tweet together: Bayesian ideal point estimation using Twitter data. *Political Analysis*, 23(1), 76–91.

Bright, J. (2016). The social news gap: How news reading and news sharing diverge. *Journal of Communication*, 66(3), 343–365. doi:10.1111/jcom.12232

Brundidge, J. (2010). Encountering "difference" in the contemporary public sphere. *Journal of Communication*, 60(4), 680–700. doi:10.1111/j.1460-2466.2010.01509.x

Bruns, A., & Burgess, J. (2015). Twitter hashtags from ad hoc to calculated publics. In N. Rambukkana (Ed.), *Hashtag publics: The power and politics of discursive networks*. New York: Peter Lang. doi:10.3726/978-1-4539-1672-8

Castles, F., & Mair, P. (1984). Left–right political scales: Some "expert" judgments. *European Journal of Political Research*, 12(1), 73–88.

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., Gonçalves, B., Flammini, A., & Menczer, F. (2012). Partisan asymmetries in online political activity. *EPJ Data Science*, 1(6). doi:10.1140/epjds6

Conover, M., Gonçalves, B., Ratkiewicz, J., Flammini, A., Menczer, F. (2011). *Predicting the political alignment of twitter users.* Proceedings of the 2011 International Conference on Privacy, Security, Risk and Trust, Boston, MA (USA), 192–199.

Dahlberg, L. (2007). Rethinking the fragmentation of the cyberpublic: From consensus to contestation. *New Media & Society*, 9(5), 827–847. doi:10.1177/1461444807081228

Döring, H. (2012). The collective action of data collection: A data infrastructure on parties, elections and cabinets. *European Union Politics*, 14(1), 161–178.

Feller, A., Kuhnert, M., Sprenger, T., Welpe, I. (2011). *Divided they tweet: The network structure of political microbloggers and discussion topics.* Proceedings of the Fifth International AAAI Conference on Weblogs and Social Media, Barcelona, Spain, 474–477.

Gaines, B., & Mondak, J. (2009). Typing together? Clustering of ideological types in online social networks. *Journal of Information Technology & Politics*, 6(3), 216–231.

Garcia, D., Abisheva, A., Schweighofer, S., Serdült, U., & Schweitzer, F. (2015). Ideological and temporal components of network polarization in online political participatory media. *Policy & Internet*, 7(1), 46–79. doi:10.1002/poi3.82

Garcia, D., Mendez, F., Serdült, U., & Schweitzer, F. (2012). *Political polarization and popularity in online participatory media.* In Proceedings of the First Edition Workshop on Politics, Elections and Data—Plead'12, Maui, Hawaii, USA.

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

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

Garzia, D., Trechsel, A., De Sio, L., & De Angelis, A. (2015). Euandi: Project description and datasets documentation. EUI RSCAS. Available from http://cadmus.eui.eu/handle/1814/34147

Gentzkow, M., & Shapiro, J. (2010). "Ideological Segregation Online and Offline." NBER Working Papers. NBER. Available from: http://www.nber.org/papers/w15916

Gruzd, A., & Roy, J. (2014). Investigating political polarization on Twitter: A Canadian perspective. *Policy & Internet*, 6(1), 28–45. doi:10.1002/1944-2866.POI354

Guerra, P., Meira, W., Cardie, C., & Kleinberg, R. (2013). A measure of polarization on social media networks based on community boundaries. Available from https://humancomputation.com/blog/?p=7046

Hargittai, E., Gallo, J., & Kane, M. (2008). Cross-ideological discussions among conservative and liberal bloggers. *Public Choice*, 134(1–2), 67–86. doi:10.1007/s11127-007-9201-x

Herrnson, P., Stokes-Brown, A., & Hindman, M. (2007). Campaign politics and the digital divide: Constituency characteristics, strategic considerations, and candidate internet use in state legislative elections. *Political Research Quarterly*, 60(1), 31–42.

Hill, K., & Hughes, J. (1998). *Cyberpolitics: Citizen activism in the age of the Internet*. Lanham, MD: Rowman & Littlefield.

Himelboim, I., McCreery, S., & Smith, M. (2013). Birds of a feather tweet together: Integrating network and content analyses to examine cross-ideology exposure on Twitter. *Journal of Computer-Mediated Communication*, 18(2), 40–60.

Huckfeldt, R., Morehouse Mendez, J., & Osborn, T. (2004). Disagreement, ambivalence, and engagement: The political consequences of heterogeneous networks. *Political Psychology*, 25(1), 65–95. doi:10.1111/j.1467-9221.2004.00357.x

Johnson, T., Bichard, S., & Zhang, W. (2009). Communication communities or "Cyber-Ghettos?": A path analysis model examining factors that explain selective exposure to blogs. *Journal of Computer-Mediated Communication*, 15(1), 60–82.

Kim, K. (2011). Public understanding of the politics of global warming in the news media: The hostile media approach. *Public Understanding of Science*, 20(5), 690–705.

Kim, Y. (2011). The contribution of social network sites to exposure to political difference: The relationships among SNSs, online political messaging, and exposure to cross-cutting perspectives. *Computers in Human Behavior*, 27(2), 971–977.

Knobloch-Westerwick, S. (2015). *Choice and preference in media use: Advances in selective exposure theory and research*. London: Routledge.

Knobloch-Westerwick, S., & Meng, J. (2009). Looking the other way: Selective exposure to attitude-consistent and counterattitudinal political information. *Communication Research*, 36(3), 426–448. doi:10.1177/0093650209333030

Krackhardt, D., & Stern, R. (1988). Informal networks and organizational crises: An experimental simulation. *Social Psychology Quarterly*, 51(2), 123–140.

Lawrence, E., Sides, J., & Farrell, H. (2010). Self-segregation or deliberation? Blog readership, participation, and polarization in American politics. *Perspectives on Politics*, 8(1), 141–157.

McPherson, M., Smith-Lovin, L., & Cook, J. (2001). Birds of a feather: Homophily in social networks. *Annual Review of Sociology*, 27, 415–444.

Mutz, D. (2002). The consequences of cross-cutting networks for political participation. *American Journal of Political Science*, 46(4), 838–855.

Myers, D., & Lamm, H. (1975). The polarizing effect of group discussion. *American Scientist*, 63(3), 297–303.

Nakagawa, S., & Schielzeth, H. (2013). A general and simple method for obtaining R^2 from generalized mixed-effects models. *Methods in Ecology and Evolution*, 4(2), 133–142.

Nunn, C., Crockett, H., & Williams, J. (1978). *Tolerance for nonconformity*. San Francisco, CA: Jossey-Bass.

Papacharissi, Z. (2002). The virtual sphere: The internet as a public sphere. *New Media & Society*, 4(1), 9–27. doi:10.1177/14614440222226244

Quattrociocchi, W., Scala, A., & Sunstein, C. (2016). Echo chambers on Facebook. Available from https://ssrn.com/abstract=2795110

Scheufele, D. (2000). Twenty-five years of the spiral of silence: A conceptual review and empirical outlook. *International Journal of Public Opinion Research*, 12(1), 3–28.

Stroud, N. (2010). Polarization and partisan selective exposure. *Journal of Communication*, 60(3), 556–576. doi:10.1111/j.1460-2466.2010.01497.x

Sunstein, C. (2002). The law of group polarization. *Journal of Political Philosophy*, 10(2), 175–195. doi:10.1111/1467-9760.00148

Sunstein, C. (2009). *Republic.com 2.0*. Princeton, NJ: Princeton University Press.

Van Alstyne, M., & Brynjolfsson, E. (1999). Global village or cyber-Balkans? *Management Science*, 51(6), 851–868. doi:10.1287/mnsc.1050.0363

Warner, B. R. 2010. Segmenting the electorate: The effects of exposure to political extremism online. *Communication Studies*, 61(4), 430–444.

Wojcieszak, M. (2009). Carrying online participation offline: Mobilization by radical online groups and politically dissimilar offline ties. *Journal of Communication*, 59(3), 564–586.

Wojcieszak, M., & Mutz, D. (2009). Online groups and political discourse: Do online discussion spaces facilitate exposure to political disagreement? *Journal of Communication*, 59(1), 40–56. doi:10.1111/j.1460-2466.2008.01403.x

Yardi, S., & Boyd, D. (2010). Dynamic debates: An analysis of group polarization over time on Twitter. *Bulletin of Science, Technology & Society*, 30(5), 316–327.