



# Network analysis reveals open forums and echo chambers in social media discussions of climate change



Hywel T.P. Williams<sup>a,\*</sup>, James R. McMurray<sup>b</sup>, Tim Kurz<sup>c</sup>, F. Hugo Lambert<sup>d</sup>

<sup>a</sup> College of Life and Environmental Sciences, University of Exeter, Laver Building, North Park Road, Exeter EX4 4QE, UK

<sup>b</sup> Department of Empirical Inference, Max Planck Institute for Intelligent Systems, Tuebingen, Germany

<sup>c</sup> School of Psychology, University of Exeter, Washington Singer Laboratories, Exeter EX4 4QG, UK

<sup>d</sup> College of Engineering, Mathematics & Physical Sciences, University of Exeter, Harrison Building, North Park Road, Exeter EX4 4QF, UK

## ARTICLE INFO

### Article history:

Received 15 November 2014

Received in revised form 11 March 2015

Accepted 19 March 2015

Available online 11 April 2015

### Keywords:

Social network analysis

Climate change

Polarisation

Echo chamber

Opinion leader

## ABSTRACT

Action to tackle the complex and divisive issue of climate change will be strongly influenced by public perception. Online social media and associated social networks are an increasingly important forum for public debate and are known to influence individual attitudes and behaviours – yet online discussions and social networks related to climate change are not well understood. Here we construct several forms of social network for users communicating about climate change on the popular microblogging platform Twitter. We classify user attitudes to climate change based on message content and find that social networks are characterised by strong attitude-based homophily and segregation into polarised “sceptic” and “activist” groups. Most users interact only with like-minded others, in communities dominated by a single view. However, we also find mixed-attitude communities in which sceptics and activists frequently interact. Messages between like-minded users typically carry positive sentiment, while messages between sceptics and activists carry negative sentiment. We identify a number of general patterns in user behaviours relating to engagement with alternative views. Users who express negative sentiment are themselves the target of negativity. Users in mixed-attitude communities are less likely to hold a strongly polarised view, but more likely to express negative sentiment towards other users with differing views. Overall, social media discussions of climate change often occur within polarising “echo chambers”, but also within “open forums”, mixed-attitude communities that reduce polarisation and stimulate debate. Our results have implications for public engagement with this important global challenge.

© 2015 The Authors. Published by Elsevier Ltd. This is an open access article under the CC BY license (<http://creativecommons.org/licenses/by/4.0/>).

## 1. Introduction

Despite broad scientific consensus on the principal mechanisms and causes of climate change (IPCC, 2013), there remains considerable debate and diversity of opinion about these topics in public discourse (Hulme, 2009; O'Neill and Boykoff, 2010; Moser, 2010; Leiserowitz et al., 2014). Such apparent uncertainty may weaken support for political action on strategies of mitigation or adaptation. While broadcast and print media presentation of climate-related issues plays an important role in shaping public opinion (Carvalho, 2010; Moser, 2010), understanding how climate

change is presented and discussed online is rapidly becoming an area of central importance (Schafer, 2012; Auer et al., 2014). Social media are already an important locus for information exchange, debate, and opinion formation on a range of issues, including climate change. The decentralised and participatory nature of online social media offers a novel opportunity to study previously inaccessible aspects of social interaction about climate change (Auer et al., 2014), including the social network structures that link individuals engaged in online debate and that are likely to affect how attitudes evolve over time.

Here we use the popular micro-blogging platform Twitter to examine user communication and social network structures associated with social media discourse on climate change. Twitter, which at the time of its registration on the New York Stock Exchange in October, 2013 had more than 200 million active users worldwide sending approximately 500 million 140-character tweets per day (United States Securities Exchange Commission,

\* Corresponding author. Tel.: +44 01392 723777.

E-mail addresses: [h.t.p.williams@exeter.ac.uk](mailto:h.t.p.williams@exeter.ac.uk) (Hywel T.P. Williams), [james.mcmurray@tuebingen.mpg.de](mailto:james.mcmurray@tuebingen.mpg.de) (J.R. McMurray), [t.r.kurz@exeter.ac.uk](mailto:t.r.kurz@exeter.ac.uk) (T. Kurz), [f.h.lambert@exeter.ac.uk](mailto:f.h.lambert@exeter.ac.uk) (F. Hugo Lambert).

2013), has been identified as an effective tool for both reflecting (O'Connor et al., 2010) and predicting (Asur and Huberman, 2010) public opinion on a variety of topics. Twitter is also an important medium for political activity (e.g. as an enabling technology for grassroots political organisation (Borge-Holthoefer et al., 2011) or as a predictor of election outcomes (Tumasjan et al., 2010)). Social media in general are fundamentally based in social networks and many-to-many communication; thus their study can reveal social network structures associated with online debates. For example, social network analyses of bloggers (Adamic and Glance, 2005) and Twitter users (Conover et al., 2011, 2012) have revealed that online political debates are often highly polarised.

There is strong evidence to suggest that social network structure will affect opinions and behaviours related to climate change. Peer attitudes are believed to have a strong influence on individual perception of climate change (Kahan et al., 2012), while social networks, including online social networks, are known to heavily influence opinions and behaviour in many other areas of human activity (Centola, 2010; Christakis and Fowler, 2007, 2008; Fowler and Christakis, 2008; Salganik et al., 2006; Sunstein, 2007; Bond et al., 2012). Studies have shown that many characteristics of individuals are clustered on social networks (McPherson et al., 2001), including obesity (Christakis and Fowler, 2007), smoking (Christakis and Fowler, 2008), political opinions (Adamic and Glance, 2005; Sunstein, 2007; Conover et al., 2011, 2012), and happiness (Fowler and Christakis, 2008). Such grouping of like with like on social networks ('homophily') is believed to arise from both preferential connection to similar individuals when forming/breaking links and also from peer influence making linked individuals more similar. Although it can be difficult to distinguish which mechanism has operated to cause homophily in purely observational studies (Shalizi and Thomas, 2011), experimental approaches have been used to demonstrate online peer influence affecting individual attitudes and behaviours, including musical preferences (Salganik et al., 2006), likelihood to vote (Bond et al., 2012), health-related behaviours (Centola, 2010), emotional transfer (Kramer et al., 2014) and rating of news stories (Muchnik et al., 2013). A key implication is that network position will affect the likelihood that an individual will adopt a new attitude/behaviour, an observation that has been used to inform successful network-based interventions (Centola, 2010; Valente, 2012).

The processes by which network structure and user communication interact to produce clustering of individual characteristics in online social networks have been studied from a social psychological perspective. Social identity has been defined as "those aspects of an individual's self-image that derive from the social categories to which he perceives himself as belonging" (Tajfel and Turner, 1979, p. 40). Rather than eroding social identity, the relative anonymity of online communication has been demonstrated to accentuate the relevance of social identity. For example, Postmes et al. (2000) examined social norms in student email communications around a web-delivered university course, finding that norms were established by an iterative process of observation and active (re)negotiation, and that they played an important role in defining newly emergent social groups. With regard to Twitter, Tamburrini et al. (2015) have recently shown that membership of a perceived ingroup can produce observable patterns in how members interact with other users; for instance, consistent with communication accommodation theory (Gallois et al., 2005), members of Twitter communities tended to adjust linguistic features of within-group tweets to be more similar to group norms. These studies illustrate the dynamic interplay between the social network structures that facilitate online interactions, the emergence of group identities of users, and the nature of ingroup/outgroup interactions online.

To date there has been little study of social media discourse or online social networks relating to the important and contentious topic of climate change (Schafer, 2012; Auer et al., 2014). Recent studies have looked at social media discussions around the September 2013 release of the IPCC Fifth Assessment Report (Pearce et al., 2014), at framing of the IPCC reports in legacy and social media (O'Neill et al., 2015), at interconnections between climate sceptic blogs (Sharman, 2014), and at the overall volume of activity and common topics of climate-related discussion worldwide (Kirilenko and Stepchenkova, 2014). However, none has so far performed an in-depth social network analysis of social media debates about climate change.

Here we report our analysis of a large dataset of Twitter messages about climate change. We constructed several forms of social network for users sending and receiving climate-related messages. We also analysed message content to assess user attitudes towards climate change and sentiment in user interactions. Our study aimed to characterise social media discussions of climate change by mapping the structure of user social networks, measuring the distribution of user attitudes across those networks, and exploring user interactions and behaviours. Section 2 of the manuscript describes the social media dataset that was collected and gives details of preliminary analysis undertaken to ensure data quality. Section 3 describes the methods used to analyse the data, which included social network construction, qualitative assessment of user attitudes and sentiment in user interactions, and network analysis to quantify homophily and elucidate community structures. Section 4 describes the results from our analysis, including social network structures and user attitudes/behaviour. In Section 5, we summarise our main results and discuss some of their possible implications. Additional results and data can be found in accompanying Supporting Information.

## 2. Dataset

We used Twitter Search API (Twitter, 2013) to collect all messages shared on Twitter between 13th January 2013 and 30th May 2013 that included five representative topic hashtags referencing climate change: #globalwarming, #climatechange, #agw (an acronym for "anthropogenic global warming"), #climate and #climaterrealists. For a subset of Twitter users appearing in the resulting database we also used the API to collect friend/follower connections. Overall we collected 590,608 distinct tweets from 179,180 distinct users. Bulk statistics for this dataset are given in Table S1.

Hashtags are utilised by Twitter users to reference a particular topic or event, enabling users to locate and contribute to related discussion. After a preliminary investigation using a keyword search on the terms "climate change" and "global warming", we chose the three most widely used hashtags (#climate, #climatechange, #globalwarming) for Twitter communication about climate change. We also chose two hashtags (#agw, #climaterrealists) that showed high usage by users expressing sceptic or contrarian views about climate change; these were chosen to ensure representation of a diversity of views in our dataset. To ensure data quality, we sampled 100 tweets from our dataset for each hashtag and manually assessed their relevance to the topic of climate change, finding that 86%, 98%, 97%, 96% and 97% of tweets were relevant for #climate, #climatechange, #globalwarming, #agw and #climaterrealists, respectively. Datasets for different hashtags were largely non-overlapping, with 95% of tweets and 83% of users utilising only one of our study hashtags (Fig. S1).

We used automated text analysis to classify tweets as "retweets" (re-transmitted messages originating from another user, identified by the text string "RT"), "mentions" (messages directly referencing another user, identified by the username-identifier character "@")

and “links” (messages containing hyperlinks to other web resources, identified by the text string “http”). Overall 39% of all tweets were classified as retweets, 22% as mentions, and 73% included links.

Composition of the Twitter user population actively discussing climate change was relatively stable during our study period, showing little turnover amongst the most active users. For each hashtag, we calculated the mean Sorensen similarity (Sørensen, 1957) between user populations for successive 10-day intervals spanning the study period. The mean value across all hashtags, weighted by the total number of active users for each hashtag, was 0.19 considering all users, rising to 0.55 for the top-100 users ranked by volume of activity (number of tweets), and 0.70 for the top-10 users. As Sorensen similarity is measured in a range from 0 (no overlap) to 1 (identical), these values indicate high turnover amongst a peripheral population of low-activity users, but low turnover (high persistence) amongst a core population of high-activity users.

### 3. Methods

#### 3.1. Network construction and visualisation

We constructed social networks for three forms of user interaction for each of the five focal hashtags. “Follower” networks consist of directed links between users who “follow” each other, i.e. a link  $A \rightarrow B$  indicates that user A is followed by user B. On Twitter, users receive all messages transmitted by users that they follow. “Retweet” networks consist of directed links indicating that one user has re-transmitted a message received from another, i.e. a link  $A \rightarrow B$  indicates that a message originally transmitted by user A was retweeted by user B. “Mention” networks consist of directed links indicating that one user has referred to another user in one of their messages, i.e. a link  $A \rightarrow B$  indicates that user A was mentioned by user B in an original tweet. Mentions are sometimes used to draw attention or engage in conversation with a particular user. Follower networks were unweighted. Edges in retweet and mention networks were weighted by number of occurrences of interaction.

Basic metrics for the 15 networks created in this study are given in Tables S2, S3 and S4. Networks were visualised in Gephi\_v0.8.2 (Bastian et al., 2009) and analysed using the NetworkX\_v1.8.1 module for Python (Hagberg et al., 2008). Networks were filtered for visualisation; follower networks were filtered by removing users with tweet volume below a specified threshold, while retweet/mention networks were filtered by removing edges with weight below a specified threshold. Networks were visualised as directed graphs using the ForceAtlas2 force-directed layout algorithm (Jacomy et al., 2012) provided by Gephi, such that closely connected users are placed near each other on a two-dimensional surface. Network layouts are based solely on topology and are independent of user attitude classifications; node colours in Figs. 2 and 4 were added post-layout to aid interpretation.

#### 3.2. Classification of user attitudes and sentiment

The most active users were classified by a panel of researchers based on their expressed attitude towards climate change, as one of: “activist” (views consistent with the scientific consensus and/or promoting action to prevent climate change), “sceptic” (views in opposition to the scientific consensus and/or opposing action to prevent climate change), “neutral” (clear views but neither activist nor sceptic as defined here), or “unknown” (views could not be identified based on information provided). Users for whom the panel did not reach a unanimous decision were classified as “ambiguous”. These categories were chosen to span a continuum of views and were deliberately few and coarse-grained in order to

avoid ambiguity, notwithstanding the diverse range of more nuanced attitudes towards climate change discussed elsewhere (e.g. (Capstick and Pidgeon, 2013; O’Neill and Boykoff, 2010)). For this analysis, the “scientific consensus” was interpreted broadly as the view that significant climate change is occurring as a result of human activity (IPCC, 2013). The classification sample ( $n = 1545$ ) consisted of the most-active users (by volume of tweets) for each hashtag, chosen using the same activity thresholds as those used to construct the follower networks. Each panel member was shown textual content for each user, consisting of username, personal profile (written by the user for public display on their homepage), and a selection of tweets randomly selected from our dataset; the coder was initially shown up to 10 tweets and could request to see more tweets up to the limit of the number held in our dataset.

Mention tweets in which an activist/sceptic user directly mentioned a single other activist/sceptic user were classified by the same panel according to the sentiment expressed about the target user by the source user, as one of: “positive” (expressing agreement, approval, praise, or other positive sentiment), “negative” (expressing disagreement, disapproval, criticism, or other negative sentiment), “neutral” (expressing a clear sentiment, but neither positive nor negative), or “unknown” (no sentiment could be distinguished from the text). The classification sample consisted of all such mentions using the #climatechange, #globalwarming and #agw hashtags ( $n = 3298$ ), since these showed most evidence of cross-attitudinal interactions. Mentions with multiple targets were omitted to avoid ambiguity. Only instances where all panel members made a unanimous decision were accepted ( $n = 2601$ ). We also recorded the frequency of embedded links in these messages.

All classifications were made by a panel of three climate science researchers. All panel members attended a verbal briefing at which the classification criteria were discussed and clarified. Each panel member then made an independent assessment in an isolated environment; no discussion was allowed. To reduce subjectivity, we only accepted unanimous classifications for further analysis. To assess inter-coder consistency, we measured Cohen’s kappa for pairwise consistency of the three panel members, finding  $\kappa$ -values of 0.60, 0.45 and 0.50 for sentiment classifications (reasonable agreement) and 0.88, 0.90 and 0.91 for user attitude classifications (excellent agreement). This level of inter-coder reliability, together with the conservative condition of unanimity for acceptance of classification, suggests that our analysis is robust.

#### 3.3. Measuring homophily

We measured homophily based on observed frequency of edges connecting users with similar or different views; high frequency of edges between similar users and/or low frequency of edges between dissimilar users was considered evidence of homophily. We restrict the analysis to users unanimously classified as activist or sceptic. The observed frequency of activist–activist, sceptic–sceptic, activist–sceptic and sceptic–activist edges was compared to the expected frequency based on the number of nodes in each class. The deviation of observed from expected frequency was used to measure homophily effect size in each network, given in units of standard deviations of edge frequencies across a bootstrap ensemble of 10,000 null networks ( $z$ -scores). Null networks were created by random re-wiring of the original network. Re-wiring replaced all activist–activist, sceptic–sceptic, activist–sceptic and sceptic–activist edges with an equal number of edges for which source and target nodes were stochastically chosen from the set  $\{a, s\}$  of all activist and sceptic nodes. For each edge, the probabilities of selecting node  $i$  as the source or target node were given by  $P_{source}(i) = \frac{k_{out}(i)}{\sum_{j \in a,s} k_{out}(j)}$  and  $P_{target}(i) = \frac{k_{in}(i)}{\sum_{j \in a,s} k_{in}(j)}$ , where  $k_{in}$  is node in-degree and  $k_{out}$  is node out-degree. This method creates null

networks with similar degree distributions as the original network and controls for variation in node degree distributions between different categories of user, which would otherwise affect the frequency of within-group and between-group edges and bias the measurement of homophily. See (Jackson, 2010) for similar reasoning. Since the bootstrap distribution of values is approximately normal, significance can be estimated as  $p < 0.05$  for  $z > 2$  and  $p < 0.003$  for  $z > 3$ . Detailed results (including exact  $p$ -values calculated from bootstrap distributions) are given in Table S6.

### 3.4. Community analysis

We used the Louvain method (Blondel et al., 2008) to algorithmically detect user communities based on social network topology. By this method, a “community” is defined as a group of users who interact more frequently with each other than they do with others. Each user can only belong to a single community. This method searches for a partition of all nodes into distinct sets (communities) such that network modularity is maximised. Modularity  $Q$  is defined (Girvan and Newman, 2002; Newman, 2006) as  $Q = \sum_{ij} \left[ \frac{A_{ij}}{2m} - \frac{k_i k_j}{(2m)^2} \right] \delta(c_i, c_j)$  for all pairs of node  $i$  and node  $j$ , where  $A_{ij} = 1$  if there is an edge between  $i$  and  $j$  and  $A_{ij} = 0$  otherwise,  $k_i$  is the degree of node  $i$ ,  $m$  is the total number of edges,  $c_i$  is the community of node  $i$ , and  $\delta(c_i, c_j)$  is the Kronecker delta. That is, modularity is increased by edges falling within communities and lowered by edges falling between communities. High modularity indicates well-defined community structure; the best partition found (i.e. the partition that maximises modularity) then gives a robust estimation of user community structure.

Community-level interaction networks were visualised using nodes to represent user communities and edges to represent the residual user interactions after within-community interactions are removed. Nodes were coloured based on community composition, calculated from the relative frequency of unanimously classified users having activist, neutral, sceptic or unknown views. We use the frequency of activists as a proxy for community composition; since neutral and unknown users are very rare in our dataset, this metric also determines the frequency of sceptic users. Counts of community sizes and member attitude classifications for each identified community in each network are given in Supporting Datasets S1, S2 and S3.

We also analysed community heterogeneity, defined as the balance between members holding sceptic and activist views, and measured as  $H = 1 - \left| \frac{a-s}{a+s} \right|$  where  $a$  is the observed frequency of activist members and  $s$  is the observed frequency of sceptic members. This measure gives values on a linear scale from perfect homogeneity ( $H = 0$ , only activist or only sceptic members) to perfect heterogeneity ( $H = 1$ , equal proportions of activist and sceptic members).

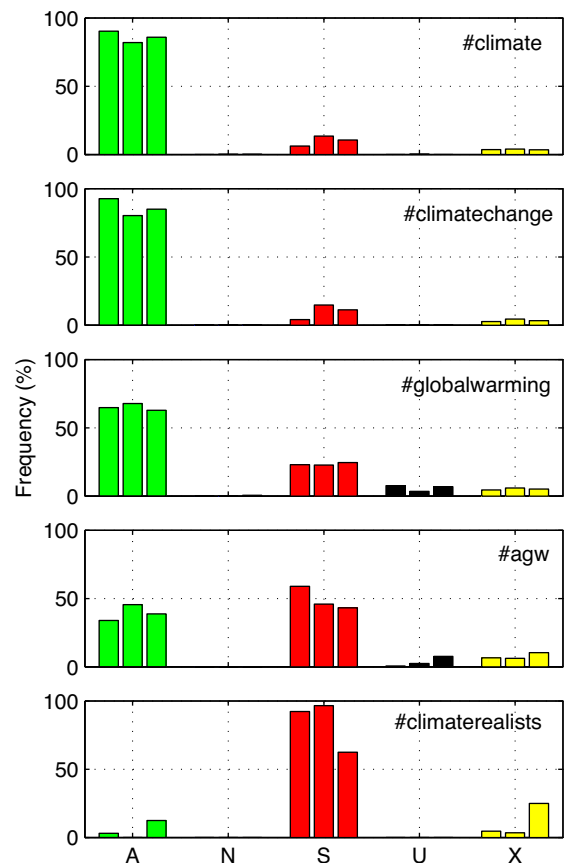
### 3.5. Potential impact of user communication

We measured the expected impact of Twitter communication for a given user based on three simple metrics: number of followers ( $f$ , measured at the time of their latest tweet in our dataset), activity ( $t$ , measured as volume of tweets using any study hashtag during the study period), and reach ( $r = f \times t$ , i.e. the product of followership and tweet volume). We calculated these metrics for each user that was classified as activist ( $n = 1049$ ) or sceptic ( $n = 325$ ) in our assessment of user attitudes.

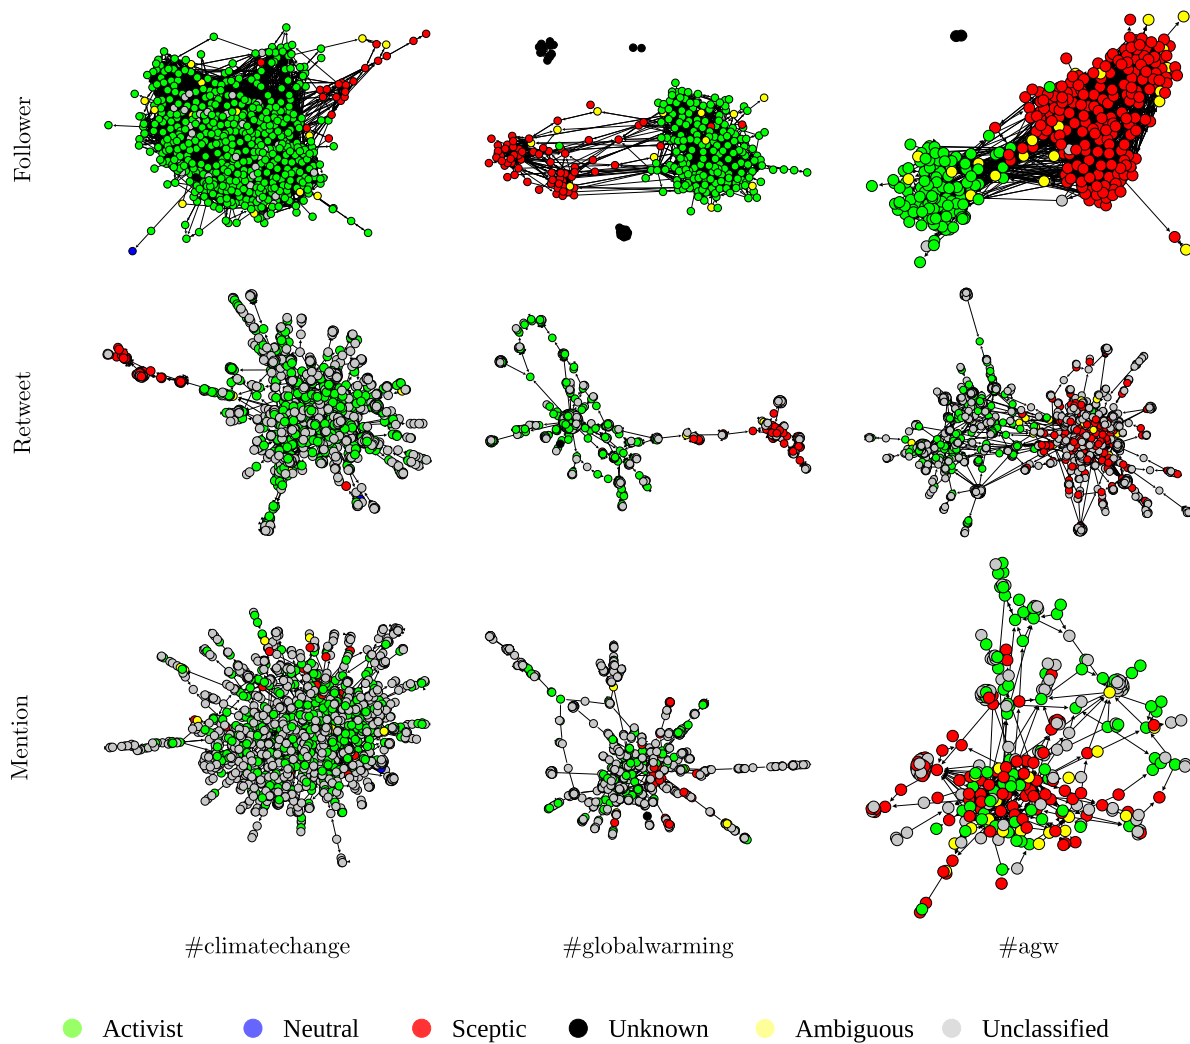
## 4. Results

To understand the balance of opinions expressed about climate change on Twitter, we analysed the distribution of attitudes

amongst users utilising each study hashtag (see Method). Overall, neutral views are effectively absent from the classified sample, while activist and sceptic views are well-represented (Fig. 1 and Table S5). Different hashtags show different attitude distributions, ranging from dominance by activists (*#climate*, *#climatechange*), to dominance by sceptics (*#climateréalists*), to strong representation by both activists and sceptics (*#globalwarming*, *#agw*). Distributions of user attitudes across social interaction networks are strongly heterogeneous (Fig. 2). Follower and retweet networks show striking segregation of users (nodes) based on attitude to climate change (node colours), with large regions dominated by either sceptics or activists. Most interactions (links) appear to occur between like-minded users, with little interaction between users with different views, i.e. these interaction networks appear to show homophily (a tendency for individuals to interact mostly with similar others (McPherson et al., 2001)). Mention networks appear less segregated, with greater levels of interaction between users with different attitudes. The presence of unclassified users (grey nodes) in retweet and mention networks demonstrates that low-activity users may still occupy central positions in these networks if they are frequently retweeted/mentioned by others.



**Fig. 1.** Frequencies of user attitudes towards climate change. Panels show data for different hashtags, colours represent different attitude categories, bars within each cluster represent frequencies for (left to right) follower, retweet and mention networks, respectively. Attitudes of 1545 prolific Twitter users utilising different hashtags to reference climate change were classified by an expert panel as one of: “activist” (A, green; supporting mainstream climate science and/or promoting climate-friendly policies), “neutral” (N, blue; expressing a view on climate change, but not obviously activist or sceptic), “sceptic” (S, red; contrarian view on climate science and/or critical of climate-friendly policies), “unknown” (U, black; no attitude could be distinguished). Only unanimous classifications were accepted; non-unanimous classifications were classified as “ambiguous” (X, yellow). See Table S5 for further detail.



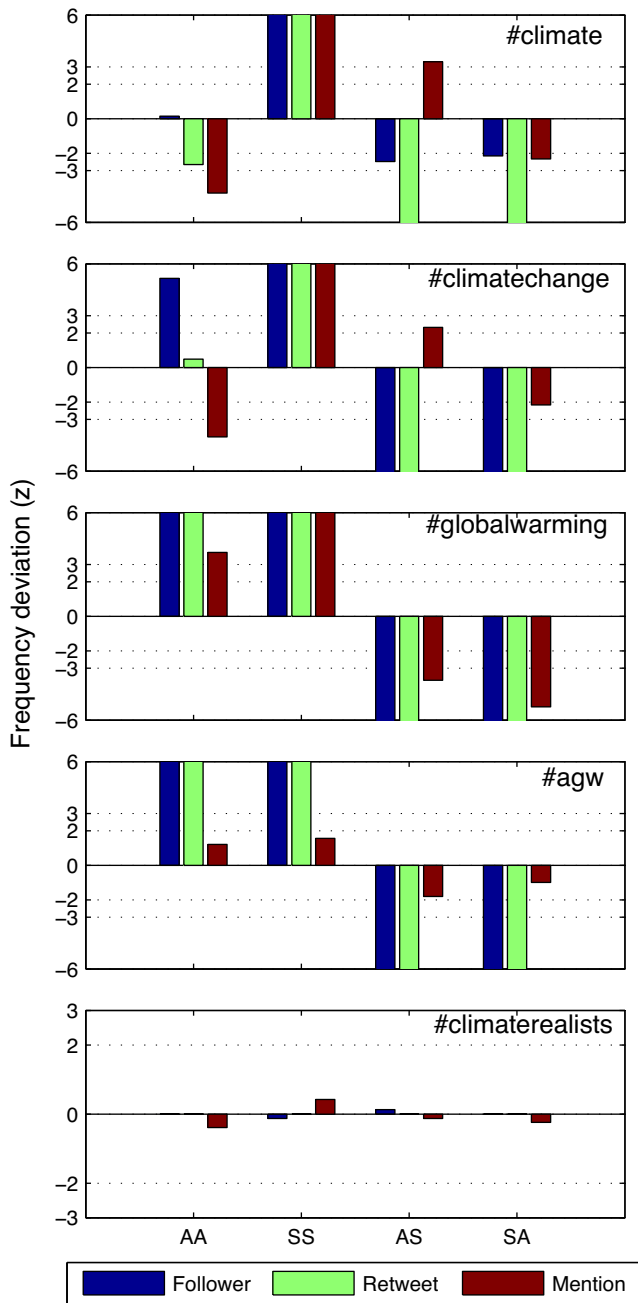
**Fig. 2.** Distribution of attitudes across interaction networks of Twitter users communicating about climate change. Rows show follower, retweet and mention networks, respectively; columns show networks for # *climatechange*, # *globalwarming* and # *agw*, respectively. Each node represents a user and each edge indicates interaction between a pair of users. Nodes are coloured by user attitude classification (see colour legend, unclassified users shown in grey). Network layouts are based solely on network topology and are independent of user attitudes. Networks are filtered for visualisation: follower networks show only users with more than [35, 12, 4] tweets, while retweet and mention networks show only edges with weights greater than [2, 1, 0] retweets and [1, 0, 0] mentions, for [# *climatechange*, # *globalwarming*, # *agw*], respectively.

We formalise and quantify the visually apparent attitude-based homophily in user social networks by examining the frequency of user interactions within and between different attitude categories (groups) (see Methods). Higher-than-expected frequency of within-group interactions and/or lower-than-expected frequency of between-group interactions are considered to be evidence of homophily. We find strong homophily in follower and retweet networks, but a mixed pattern of homophily and heterophily in mention networks (Fig. 3 and Table S6). Mention networks for #*agw* and #*globalwarming* show clear homophily amongst both sceptics and activists. Mention networks for #*climatechange* and #*climatechange* show homophily amongst sceptics and heterophily amongst activists (i.e. higher-than-expected frequency of mentions of sceptics by activists), suggesting that the sceptic minority has disproportionately high visibility in debate using these hashtags.

In order to examine whether users were grouped into like-minded communities, we next used a community-detection algorithm (Blondel et al., 2008) to partition users in each network into local communities of frequent interaction (see Methods). We found partitions giving high modularity scores ((Girvan and

Newman, 2002; Newman, 2006)) for all networks (Tables S2, S3 and S4), indicating strong community structure. We visualised this structure as networks of interactions between communities (Fig. 4). Here each node represents a single community and each edge represents the total volume of interaction between users in the two connected communities. Nodes are coloured according to the mix of attitudes held by community members (see Method). It is visually apparent that most communities in follower and retweet networks have a strongly homogeneous distribution of attitudes towards climate change, dominated by either activist or sceptic views, while very few communities have a mix of both perspectives. However, mention networks are less segregated and show more frequent occurrence of mixed communities containing both activists and sceptics. The community–community interaction networks also give a strong visual suggestion that homophily also occurs at this level of granularity, with a clear tendency for communities to interact most strongly with other communities of similar composition.

To quantify the typical range of attitudes seen by each user, and further explore the nature of user communities around online discussions of climate change, we plotted frequencies of users



**Fig. 3.** Homophily in user interactions based on attitude towards climate change. Homophily was measured as the deviation of the observed frequency of edges connecting sceptic (S) and activist (A) users from the expected frequency derived from the overall frequencies of attitudes in the user population forming each interaction network (see Methods). Deviations are given as z-scores (multiples of standard deviations from expected mean) from bootstrap comparisons of 10,000 null networks. Homophily is indicated by positive z-scores for within-group (AA and SS) interactions and negative z-scores for between-group (AS and SA) interactions. Gridlines at  $z = \pm 2$  and  $z = \pm 3$  indicate estimated statistical significance at levels  $p < 0.05$  and  $p < 0.003$ , respectively (based on extrapolation from observed normal distribution of bootstrap ensemble; exact statistics are given in Table S6). Plotted values are bounded to range  $z \in [-6, 6]$ ; actual values are often considerably more extreme (Table S6).

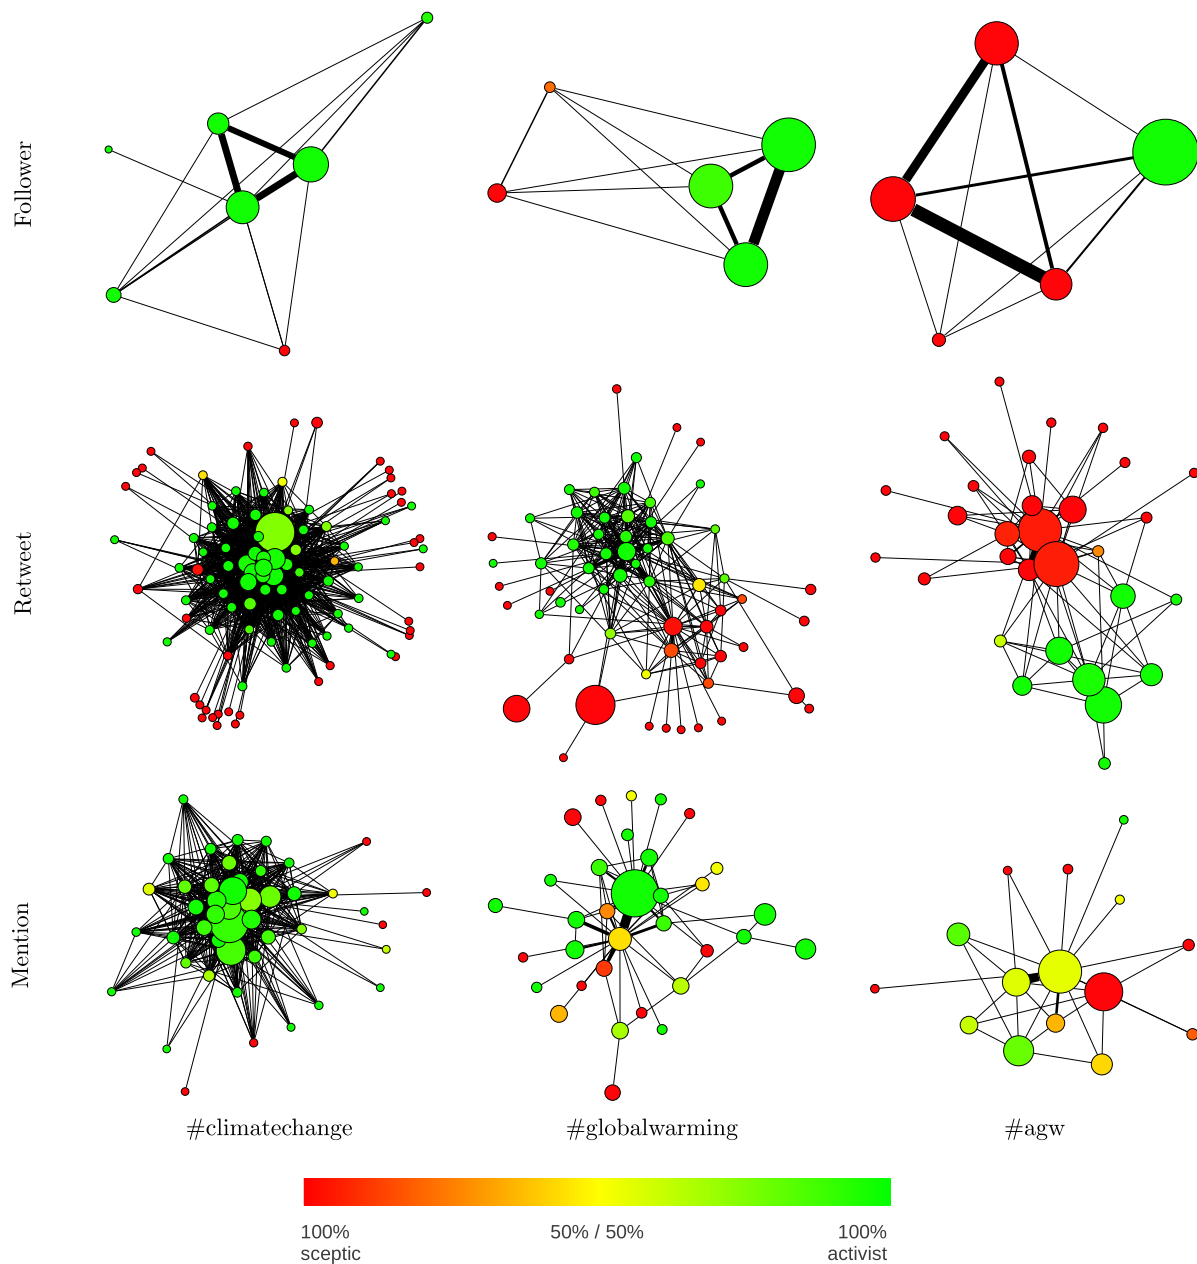
belonging to communities with different compositions (Fig. 5). Community composition was calculated based on the frequency of different attitudes expressed by members (see Method). Almost all users in follower and retweet networks have local communities that are heavily dominated by either activist or sceptic views, with no evidence of mixed-view communities. However, mention

networks show more evidence of mixed-attitude communities, with a substantial number of users embedded in mixed communities for the #globalwarming and #agw hashtags in particular. Adopting simple definitions of an ‘echo chamber’ as a community where at least 90% of members share the same view and an ‘open forum’ as a community where both sceptic and activist views have at least 10% representation, across all hashtags we find that 98%, 94% and 68% of users are members of an echo chamber, and 2%, 3% and 28% are members of an open forum, for follower, retweet and mention networks, respectively. These data confirm that follower and retweet networks are segregated by user attitude, while mention networks show a much greater level of mixed-attitude interactions.

To better understand the nature of interactions between users with differing views, we analysed text content of mentions where both the source and target users had been classified as having either activist or sceptic views, restricting our analysis to mentions with only a single target user to avoid ambiguity (see Methods). The sentiment expressed about the target user by the source user, was classified as ‘positive’, ‘negative’, ‘neutral’, or ‘unknown’ (Fig. 6 and Table S7). We also checked these tweets for the presence of embedded hyperlinks. Overall, most mentions had unknown sentiment, indicating that they did not contain any clearly identifiable sentiment. No tweets were classified with neutral sentiment (i.e. having identifiable sentiment, but neither positive or negative), suggesting a strong element of self-selection bias in this data; users do not tend to express ambivalent or equivocal sentiment when directly referring to another user. Amongst all mention tweets classified as having positive or negative sentiment, 39% also included an embedded hyperlink, while amongst mentions with unknown sentiment, 77% included a link. Also, amongst all mentions including a link, 5% showed positive or negative sentiment, while amongst mentions without a link, 22% showed positive or negative sentiment. Thus we find a strong negative association between the presence of sentiment and the presence of links in mention tweets (Fisher’s exact test, two-tailed  $p < 10^{-32}$ ,  $n = 2601$ ). The 140-character limit on message text may partly explain this finding; including a link (22 characters in compressed form) reduces the available space in which to express sentiment. However, our finding is also consistent with variable usage of mentions, with some being purely informative (with a link but no sentiment) and others more discursive or conversational (with sentiment but no link). The latter interpretation is consistent with our subjective experience of mention usage.

We next explored partisanship between sceptic and activist users by considering sentiment expressed within and between these groups (Table S7). While overall there were fewer between-group mentions ( $n = 558$ ) than within-group mentions ( $n = 2043$ ), the frequency of positive/negative mention sentiment was much higher between attitude groups ( $n = 187 \approx 34\%$ ) than within them ( $n = 69 \approx 3\%$ ). We find strong evidence for positive sentiment in ingroup interactions and negative sentiment in outgroup interactions (where we define activists as the outgroup for sceptics, and vice versa). Sentiment in mentions by activists shows a strongly significant relationship with the attitude of the target user (Fisher’s exact test for association between activist/sceptic attitude of target user and positive/negative sentiment, two-tailed  $p < 0.0001$ ,  $n = 111$ ) with negative sentiment expressed towards sceptics and positive sentiment expressed towards other activists. Similarly, sentiment in mentions by sceptics showed a strongly significant relationship with the attitude of the target user (Fisher’s exact test, two-tailed  $p = 0.0008$ ,  $n = 145$ ) with negative sentiment expressed towards activists and positive sentiment expressed towards other sceptics.

In order to characterise user behaviours in relation to social media communication about climate change, we analysed tweets

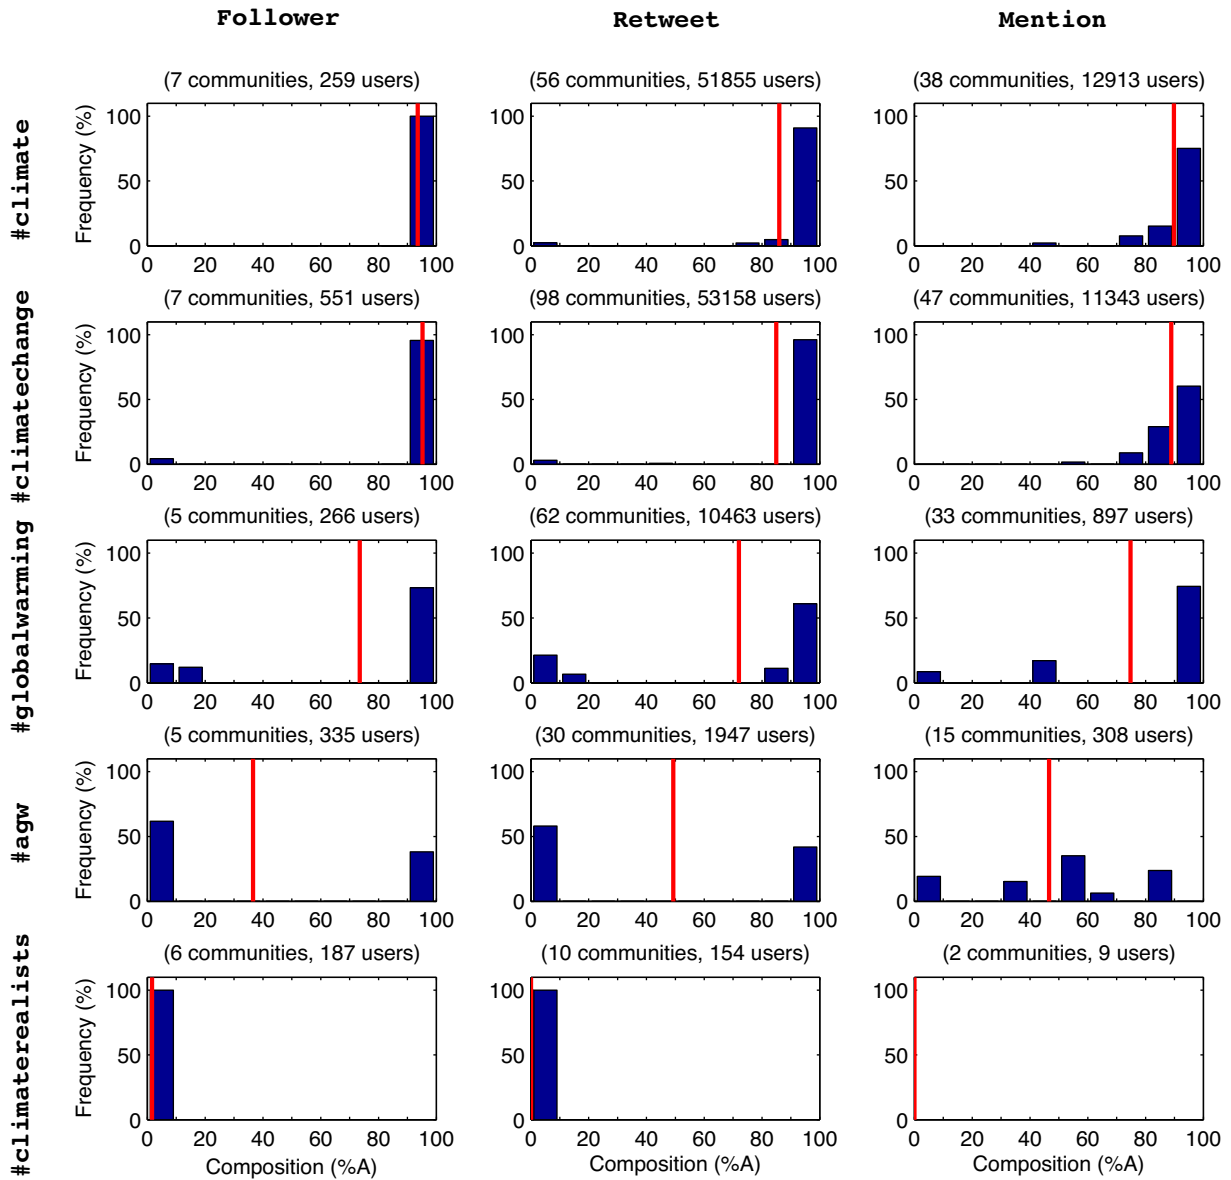


**Fig. 4.** Community-level interaction networks. Here each node represents a user community, while edges represent the total volume of interaction between users in each community (see Methods). Nodes are sized by the number of users in the community and coloured based on the attitudes of their members (scaled from red to green by frequency of activist users, see Methods). Counts of user attitude classifications for each identified community in each network are given in Datasets S1, S2 and S3. Modularity ( $Q$ ) scores for the final community assignment are (top-left to bottom-right): [0.284, 0.443, 0.475; 0.772, 0.883, 0.733; 0.724, 0.825, 0.470].

in which a sceptic/activist user mentioned a single other sceptic/activist user (i.e. where a pairwise interaction between users with identifiable views can be unambiguously defined). We first considered the level of user engagement with opposing viewpoints and whether between-group engagement was reciprocated. Sceptics tended to interact more frequently with their outgroup; overall 57% of mentions by sceptics were directed to activists, while 10% of mentions by activists were directed to sceptics. However, due to the relatively small number of sceptics overall (Fig. 2 and Table S5), the null expectation would be for 72% of mentions by sceptics and 20% of mentions by activists to be targeted to their outgroups. Thus the observed imbalance still represents a homophilic pattern of interaction (Fig. 3 and Table S6). With regard to reciprocity of cross-group engagement, considering

both sceptic and activist users together, there was a positive relationship (Pearson's  $r = 0.65$ ,  $p < 0.0001$ ,  $n = 510$ ) between the frequencies of inward and outward mention interactions involving the outgroup. This relationship was stronger for sceptics (Pearson's  $r = 0.82$ ,  $p < 0.0001$ ,  $n = 87$ ) and weaker, but still present, for activists (Pearson's  $r = 0.49$ ,  $p < 0.0001$ ,  $n = 423$ ). Thus it appears that efforts to engage users with differing views are typically reciprocated.

We next analysed the level of user negativity towards their outgroup and whether negativity was reciprocated. Overall both sceptics and activists expressed similar levels of negative sentiment towards the outgroup. We defined outward negativity for each user as the proportion of mention tweets in which they expressed negative sentiment about another user, and inward



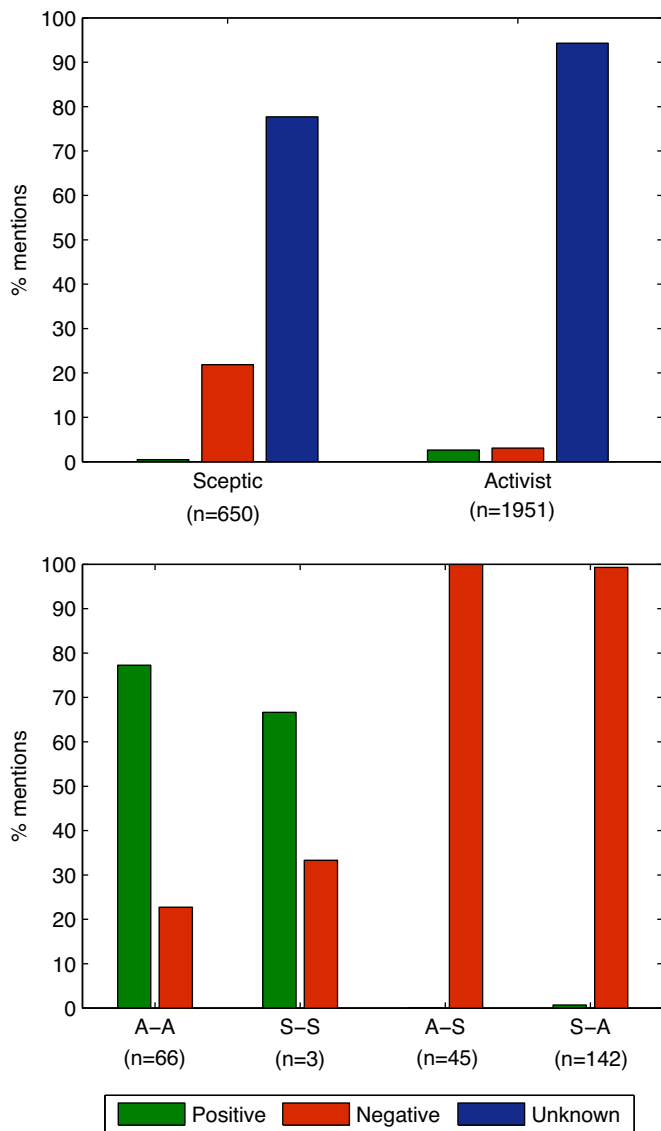
**Fig. 5.** Composition of communities in user social networks. Plots show frequencies of users belonging to communities with different distributions of attitudes amongst their members. Community composition is given here as the frequency of members with activist attitudes (see Method). Frequencies are binned into 10% ranges. Red lines show composition of overall user population for each network, indicating the expected group composition if user types were distributed homogeneously. Only communities with  $\geq 10$  unanimously classified users were included in this analysis; this condition gives coverage of:  $\frac{249}{259}$ ,  $\frac{39,542}{51,855}$ ,  $\frac{11,622}{12,913}$ ,  $\frac{549}{551}$ ,  $\frac{28,269}{53,158}$ ,  $\frac{10,244}{11,343}$ ,  $\frac{266}{266}$ ,  $\frac{4959}{10,463}$ ,  $\frac{351}{897}$ ,  $\frac{327}{335}$ ,  $\frac{1383}{1947}$ ,  $\frac{281}{308}$ ,  $\frac{133}{187}$ ,  $\frac{115}{154}$  and  $\frac{0}{9}$  users in each network (reading left-to-right, top-to-bottom in the plot above). See Datasets S1, S2 and S3 for detailed data on community sizes and composition. Insufficient users were classified to perform this analysis for the #climateréalists mention network.

negativity as the proportion of mentions in which another user expressed negative sentiment about them. Amongst those users who engaged with (i.e. directed at least one mention towards) their outgroup, mean values for outward negativity towards the outgroup were similar for sceptics (22.1%) and activists (22.4%), with no significant difference in the distributions (Mann–Whitney  $U = 1300.5$ ,  $n_A = 36$ ,  $n_S = 37$ ,  $p = 0.7194$ ). The higher overall frequency of negative sentiment in mentions by sceptics (Fig. 6) arises from their greater frequency of interaction with their outgroup. With regard to reciprocity of negativity, considering both sceptics and activists together, there was a positive relationship (Pearson’s  $r = 0.36$ ,  $p = 0.0092$ ,  $n = 51$ ) between inward and outward negativity in mentions involving the outgroup. This relationship was stronger for sceptics (Pearson’s  $r = 0.59$ ,  $p = 0.0047$ ,  $n = 21$ ) and present, but not significant, for activists

(Pearson’s  $r = 0.26$ ,  $p = 0.1679$ ,  $n = 30$ ). This suggests a pattern of reciprocity in the expression of negative sentiment; those who are negative towards their outgroup are themselves the target of negative sentiment from their outgroup.

To determine whether there was any association between expressed negativity and user engagement with the outgroup, we tested for correlation between outward negativity towards the outgroup and overall frequency of interaction with the outgroup (summing both inward and outward mentions). Amongst those users with at least one mention directed towards the outgroup, considering both sceptics and activists together, we found no linear relationship (Pearson’s  $r = -0.02$ ,  $p = 0.8577$ ,  $n = 73$ ) due to a large proportion of users who never expressed negative sentiment. Further restricting the sample to only those users who authored at least one negative tweet targeting an outgroup user, we found a





**Fig. 6.** Sentiments in user–user mentions were classified as one of: “positive” (agreement, approval, praise), “negative” (disagreement, disapproval, criticism), “neutral” (sentiment, but neither positive nor negative), “unknown” (no sentiment could be distinguished). See Method for details of classification. No tweets were classified as having neutral sentiment. The upper panel shows sentiment in all tweets made by sceptics and activists, irrespective of whether their target was activist or sceptic. The lower panel shows positive/negative sentiment in tweets within and between the sceptic (S) and activist (A) user populations. Labels have ‘source–target’ format, e.g. A–S indicates a tweet written by an activist user in which a sceptic user was mentioned.

clear negative relationship between interaction frequency and outward negativity (Spearman’s  $r = -0.54$ ,  $p = 0.0003$ ,  $n = 40$ ), present for both activists (Spearman’s  $r = -0.72$ ,  $p = 0.0007$ ,  $n = 18$ ) and sceptics (Spearman’s  $r = -0.48$ ,  $p = 0.0244$ ,  $n = 22$ ). This may suggest either that users who take a negative approach to cross-group engagement are less likely to sustain repeated interactions, or alternatively, that users who frequently interact with their outgroup become less negative towards them over time.

To explore the effect of social context on user behaviour, we first examined the relationship between user attitude and the balance of attitudes to which that user was frequently exposed on Twitter. For each community identified in the mention networks for #climatechange, #globalwarming and #agw (Fig. 4), we measured heterogeneity in attitudes amongst its members (see

Methods). Across all communities in which  $\geq 10$  users were classified for attitudes to climate change, there was a negative relationship (Pearson’s  $r = -0.45$ ,  $p = 0.0007$ ,  $n = 52$ ) between heterogeneity and the proportion of classified members holding a polarised attitude (i.e. sceptic or activist). Consistent with this finding, there was a positive relationship (Pearson’s  $r = 0.48$ ,  $p = 0.0003$ ,  $n = 52$ ) between community heterogeneity and the proportion of members with attitude classified as ambiguous, i.e. members for which our expert panel did not reach a unanimous decision. The panel often disagreed about a user’s attitude when their messages were less extreme; while we exclude ambiguous classifications from most of our analysis, in this context they are suggestive of a more moderate or equivocal attitude. Thus overall we find lower levels of polarisation in mixed-attitude communities that have a balance of sceptic and activist views. Next we analysed the relationship between social context and the expression of negative sentiment towards the outgroup, comparing community heterogeneity and outward negativity towards the outgroup amongst community members. While polarised views were less frequent in heterogeneous communities, those members who did hold a polarised view were more likely to express negative sentiment towards their outgroup. For communities with at least one member who directed at least one mention to the outgroup, there was a positive relationship (Pearson’s  $r = 0.43$ ,  $p = 0.0036$ ,  $n = 43$ ) between community heterogeneity and the proportion of all outgroup mentions by its members that contained negative sentiment. That is, members of heterogeneous communities expressed higher levels of negative sentiment towards their outgroups than did members of homogeneous communities, after normalisation for frequency of interaction with the outgroup. Consistent with this community-level pattern, individual users for whom  $\geq 10$  mentions of the outgroup were classified for sentiment showed a strong positive relationship (Pearson’s  $r = 0.65$ ,  $p = 0.0091$ ,  $n = 15$ ) between the heterogeneity of their community and their negativity towards the outgroup. This relationship was present for both activists (Pearson’s  $r = 0.68$ ,  $p = 0.1365$ ,  $n = 6$ ) and sceptics (Pearson’s  $r = 0.73$ ,  $p = 0.0263$ ,  $n = 9$ ), though low sample size for the former group reduced its significance. Overall we find that exposure to a balance of views is associated with reduced levels of polarisation, but increased levels of partisan negativity from those users who hold a polarised view.

The last component to our analysis was to examine the potential impact of communications by activist and sceptic users, based on numbers of followers, activity (volume of tweets), and overall reach (activity  $\times$  followership) (see Method). All three measures showed skewed (approximately log-normal) distributions. With regard to followership, activists typically had more followers than sceptics, but the difference was not significant (activists:  $min = 2$ ,  $lq = 272.25$ ,  $median = 807$ ,  $uq = 2878.25$ ,  $max = 31913887$ ; sceptics:  $min = 3$ ,  $lq = 224.25$ ,  $median = 753$ ,  $uq = 2014.5$ ,  $max = 67451$ ; Mann-Whitney  $U$ -test,  $p = 0.07$ ). However, the activist group was distinguished by a small number of users with very high followership (e.g. 29 users with  $>100000$  followers) compared to users in the sceptic group. Activity levels were significantly higher for activists, who typically released more tweets than sceptics (activists:  $min = 1$ ,  $lq = 23.75$ ,  $median = 55$ ,  $uq = 107$ ,  $max = 3993$ ; sceptics:  $min = 1$ ,  $lq = 9$ ,  $median = 16$ ,  $uq = 35$ ,  $max = 2397$ ; Mann-Whitney  $U$ -test,  $p < 0.0001$ ). Typical reach was also significantly higher for activists than for sceptics (activists:  $min = 14$ ,  $lq = 10550.5$ ,  $median = 52128$ ,  $uq = 215167.5$ ,  $max = 159569435$ ; sceptics:  $min = 68$ ,  $lq = 3082.5$ ,  $median = 12384$ ,  $uq = 65285.75$ ,  $max = 19283865$ ; Mann-Whitney  $U$ -test,  $p < 0.0001$ ). Thus while activists and sceptics generally have similar numbers of followers, activist users send a greater number of messages and thereby have greater potential reach with their communication on Twitter.

## 5. Discussion

In this study we used the social media platform Twitter to analyse online many-to-many communication about climate change. We found a high degree of polarisation in attitudes, consistent with self-selection bias; those users who were most active in online discussions of climate change tended to have strong attitudes (either activist or sceptic) and neutral views were largely absent. Our results show that social media discussion of climate change is characterised by strong attitude-based homophily and widespread segregation of users into like-minded communities. Most users in our study interacted only with like-minded others and exposure to alternative views was relatively rare. However, there were also a minority of users who frequently interacted with others with differing views and we identified coherent mixed-attitude communities where such interactions occurred.

We found several regularities in user behaviour that help illustrate the nature of the Twitter debate about climate change. Users varied in the level and tone of their engagement with others holding different views, but efforts to engage with the outgroup, and the sentiment associated with these interactions, were typically reciprocated. Users exposed to diverse views in mixed-attitude communities were less likely to hold a polarised view; however, those members that did hold polarised views were more likely to express negative sentiment towards others with differing views. For the limited set of behavioural metrics we studied (primarily homophily, engagement with the outgroup, and partisan sentiment), we found no substantive differences between activist and sceptic users. Observed high negativity by sceptic users is a statistical feature of their minority status; since sceptics on average interacted more with their outgroup, they tended to express more negative sentiment, but they were not more negative than activists on a per-interaction basis. However, while sceptics and activists typically had similar numbers of followers, activists tended to be more active communicators and hence had greater potential reach with their messaging.

In this study we used a selection of five climate-related hashtags to identify tweets about climate change. Since our research aim was to study interactions between social network structure and polarisation of attitudes, our data collection strategy aimed to capture both the bulk of Twitter communication about climate change and also a diversity of user views. Three of the hashtags used (*#climate*, *#climatechange*, *#globalwarming*) are generic and widely used, representing 97.7% of the tweets collected in this study. The other two hashtags (*#agw*, *#climaterealist*s) were identified in preliminary analysis as having relatively high usage by users expressing sceptic views about climate science and climate politics. These were deliberately included to increase representation of sceptic users in our dataset – sceptic views were relatively rare in tweets using the more generic hashtags. An important finding of our study is that different hashtags are associated with different distributions of user attitudes (Fig. 1). This specificity of attitude distributions across hashtags motivates our decision to report results on a 'per hashtag' basis and avoid merging or agglomerating data across tags, as has been done elsewhere. It also raises the issue of generality – are our results based on a representative sample of Twitter activity? In a different study (unpublished data) that collected tweets using a larger set of 27 climate-related hashtags, we found that *#climate*, *#climatechange* and *#globalwarming* were by far the most heavily used hashtags in that dataset, together accounting for 80% of tweets over a 9-month period spanning September 2013–May 2014. In contrast, *#agw* accounted for 1.4% of tweets while *#climaterealist*s had effectively fallen out of use. While we do not account for tweets without hashtags, this finding confirms the

importance of *#climate*, *#climatechange* and *#globalwarming* as the most representative hashtags (by volume of use) for climate change discussion on Twitter. Furthermore, while generalisation of hashtag-level results may not always be appropriate (e.g. the balance of expressed attitudes towards climate change), the social network-related phenomena reported here were seen across all hashtags and are likely to be general; this includes polarisation of attitudes, homophily, presence of echo chambers (and open forums), and individual-level trends in behaviour (e.g. reciprocation of negativity and engagement with outgroups).

Social networks based on different forms of user interaction have different properties. We consistently observed strong attitude-based homophily in follower and retweet networks, but much less consistent and weaker homophily in mention networks. This may reflect different user goals when engaging in different kinds of interaction, as well as overall motivations for using social media. Online social networks are recognised sites of both the construction of social identities (Zhao et al., 2008) and their linguistic performance (Tamburrini et al., 2015). Moreover, the public nature of communication on Twitter is likely to encourage users towards behaviours consistent with the image they wish to express. Following another user is a public decision to associate with and receive content from that user. Retweeting often implies (public) endorsement of either the individual tweet or its original author. Thus users are likely to follow/retweet others with views consistent with their own. Meanwhile, mentioning another user in a tweet might have several purposes. Most mention tweets do not express identifiable sentiment and include a link to an external web resource, suggesting a purely informative purpose. However, mentions can also form part of a discussion or conversation, or offer (possibly critical) comment on the target user's activities or expressed attitudes. Thus mentioning another user with a conflicting view can still be a coherent expression of social identity, so long as it is framed appropriately. Social identity theory (Tajfel and Turner, 1979) asserts that group identity is often defined by contrast with an outgroup, and suggests that willingness to negatively engage with outgroup members is a way of affirming membership of the ingroup. Indeed, here mentions of users with alternative views were often accompanied by negative sentiment; the lack of universal homophily in mention interactions does not appear to imply a lack of partisan feeling. Thus the differing levels of homophily we observe in different kinds of interaction network are consistent with a view of social media activity as an expression of social identity.

For the contentious topic of climate change, we find that most individuals engaged in online discussions are embedded within communities of like-minded users; such self-reinforcing "echo chambers" can prevent engagement with alternative viewpoints and promote extreme views (Sunstein, 2007). Partisan online communities may also act as selective filters that impede transmission of unfavoured ideas across the broader social network. However, here we also identified mixed-attitude communities in which users were frequently exposed to a diversity of viewpoints. We characterise such communities as "open forums", in which cross-constituency discussions and exchange of ideas can take place, and speculate that the reduced likelihood of polarised views that we observed for these communities is indicative of a moderating effect of such interactions. Cross-group interactions at least indicate an open channel for information flow and potential influence; although such interactions were often acrimonious, it is hard for a user to be influenced by another user with whom they have no interaction at all.

While we have not measured change in attitudes over time in this study, and thus cannot directly quantify influence, we observed significant relationships between the diversity of attitudes expressed in a user's online community and their

expressed views and communicative behaviour. Exposure to a diversity of views was associated with a lower likelihood of holding a polarised view; this is consistent with experimental studies showing that individual attitudes tend to be weaker in attitudinally diverse groups than in attitudinally congruent groups (Visser and Mirabile, 2004; Levitan and Visser, 2009). However, within the same mixed-view communities, users who did express a polarised view were more negative in their interactions with opposing viewpoints. We speculate that these users, who hold a strong view about climate change but inhabit a community where conflicting views are often expressed, have a strong perception that there exists a genuine and contested debate about climate change. Thus these users make greater efforts to influence the debate towards their own view and appear more confrontational or aggressive towards users who express conflicting views. Willingness to confront (Czopp, 2013) or express disapproval for (Swim and Bloodhart, 2013) “anti-environmental” behaviour in a social context has been shown to influence the future environment-related behaviours of both the target of the confrontation and passive bystanders, perhaps explaining its use in the present context in attempts to influence others. In contrast, we suggest that users in communities dominated by views from one side of the debate may have less perception that an active debate exists amongst their interaction group. These users may be less motivated to engage in a negative or confrontational manner, since they have little to gain and there may be a social cost to confrontation (Nolan, 2013). Users from a minority group within a community may also avoid negativity in order to maintain interaction with an outgroup majority, consistent with our observation that the users with the highest frequency of interaction with the outgroup tended to express a low level of negative sentiment.

In general, the structure, content and dynamics of social media interactions and online discourses remain poorly understood. Social media content varies in type, including text, images, video and conversational media, requiring different forms of analysis. Social media debates are typically decentralised, fragmented and diffuse, with very large numbers of participants each making a relatively small contribution. The resulting difficulty of locating and demarcating a coherent and representative data corpus for a particular topic poses a significant challenge to established methods of textual analysis, such as discourse or conversation analyses (Edwards and Potter, 1992). Here we have aggregated topical samples of short text messages from Twitter, which individually have limited information content, but can collectively offer rich datasets. We have shown that social network analysis and automated community-detection, in conjunction with qualitative analysis of textual content, can be used to identify different kinds of communities and begin to understand the processes of online social influence. While mixed-attitude communities are rare in our dataset, relative to the larger number of communities dominated by a single viewpoint, they offer an intriguing possibility for examining several previously unexplored issues regarding online social interactions around contentious topics. For example, longitudinal study of mixed-attitude communities may reveal key processes of online community formation, social influence, and conflict resolution; do mixed-attitude communities reach a consensus over time, or do they fragment along ideological faultlines? Due to the pervasive use of social media to communicate about almost all aspects of human activity, the ability to examine these generic processes might be fruitfully applied to a wide variety of domains. Study of social media can also offer an improved understanding of public opinion (here about climate change) by considering how individual attitudes are distributed across social networks, complementing existing data on bulk frequencies of different attitudes from (e.g.) segmentation-based survey methods (Leiserowitz et al., 2008).

Although we did not measure the evolution of opinions over time, our results have some relevance for understanding opinion leadership on Twitter with regard to climate change. The “two-step flow” model of communication (Katz and Lazarsfeld, 1955; Katz, 1957) states that media influence is exerted through the intermediary actions of “opinion leaders”, individuals who are highly engaged with media content around an issue and who act to interpret and disseminate new information to others. The two-step flow model has been examined in the context of climate change campaigns by Nisbet and Kotcher (2009), while various studies (e.g. (Cha et al., 2010, 2012; Wu et al., 2011)) have given empirical support for two-step flow as a general model for media dissemination on Twitter. Node centrality is often used as a proxy measure to identify opinion leaders in social networks (Valente, 2012). One of the simplest centrality measures is degree centrality, which in the context of Twitter corresponds to the number of followers. Our results show that activists and sceptics typically had similar numbers of followers, suggesting similar status as opinion leaders. However, activists were more abundant than sceptics, and typically more active (by volume of tweets), giving activist views greater overall exposure. The activist group also contained a small number of users with very high follower numbers. All else being equal, these results suggest that opinion leadership on Twitter is in general likely to promote (activist) views supportive of the scientific consensus and actions to mitigate or adapt to climate change. However, there are several important caveats concerning the validity of equating followership with influence – for example, followership does not measure how many tweets are actually seen, or the impact of those tweets on user attitudes. The dynamics of influence are further complicated by possible effects of echo chambers and partisan filtering, as identified above.

The majority of Twitter users directly engaged with climate change (e.g. the 179,180 users in our dataset who used climate-related hashtags) were connected by their interactions into large social networks, creating the potential for influence and opinion leadership to be enacted. However, examination of follower numbers suggests that many (or most) followers of these engaged users were not themselves participants in discussions of climate change. Thus to some extent all of the engaged users in our dataset could be argued to act as opinion leaders for the topic of climate change on Twitter, in that they disseminate information from an engaged user population (who are themselves the producers of social media content relating to climate change) to a larger non-engaged user population. The most-followed users have an especially important position in this regard; several users we examined had follower populations larger than the entire climate-engaged user population represented by our dataset, giving them exceptional potential reach. Elaboration of the multi-step flow of online communication about climate change, and the ways in which it interacts with network structure and other media to shape the evolution of opinions amongst social media users and others, is an important topic for future research.

Discussion of climate change on social media is likely to affect the wider ‘offline’ climate debate. While the impact of social media is likely to be greatest in countries with high levels of Internet use, impacts may still be felt in countries with lower Internet usage, where those who do access the Internet are likely to be opinion leaders. For example, a recent large survey of climate change awareness in India showed that individuals leading local opinion about climate change were likely to have Internet access and make frequent use of online social networks (Gambhir and Kumar, 2013). Social media are increasingly a key locus for information exchange, opinion formation, political debate and political activism (Sunstein, 2007; Farrell, 2012). Online media are already an important and credible source of information about climate change (Schafer, 2012), while social media are known to influence mainstream

media on scientific topics (Bonetta, 2007; Schafer, 2012). In comparison with broadcast media such as print or television, the social and participatory aspects of social media may give it a more central role in shaping individual attitudes and behaviours. Peer influence is argued to be an important factor in assessment of the risks of climate change (Kahan et al., 2012), to the extent that some authors have proposed its use in climate change campaigns (Nisbet and Kotcher, 2009), while social network influence is known to affect individual attitudes and behaviours in general (Katz and Lazarsfeld, 1955; Katz, 1957; Salganik et al., 2006; Bond et al., 2012; Centola, 2010; Wu et al., 2011; Muchnik et al., 2013; Kramer et al., 2014). Social media may thus be expected to have a significant impact on the evolution of public opinion about climate change in the general population, by direct effects on attitudes amongst the engaged user population, secondary effects on the non-engaged user population, cross-media interactions (O'Neill et al., 2015), and social percolation of opinions to non-users. The nature of this impact is hard to predict. The presence of echo chambers may have a stabilising effect by reinforcing existing views, while open forums might promote greater change. Social media discussions of climate change are underpinned by social networks, which both facilitate and constrain the discussions that can take place. As social media usage increases, the complex structural features of these networks are likely to have an increasingly important role in shaping public engagement with this important global challenge.

## Acknowledgments

This work was supported via the University of Exeter's Bridging the Gaps project (funded by the UK EPSRC's Bridging the Gaps initiative - grant reference EP/I001433/1). We thank three anonymous referees for their comments, which have greatly improved the manuscript.

## Appendix A. Supplementary data

Supplementary data associated with this article can be found, in the online version, at <http://dx.doi.org/10.1016/j.gloenvcha.2015.03.006>.

## 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. ACM, pp. 36–43.
- Asur, S., Huberman, B.A., 2010. Predicting the future with social media. In: Proc. Web Intelligence and Intelligent Agent Technology (WI-IAT), 2010. IEEE, pp. 492–499.
- Auer, M.R., Zhang, Y., Lee, P., 2014. The potential of microblogs for the study of public perceptions of climate change. Wiley Interdiscip. Rev.: Clim. Change.
- Bastian, M., Heymann, S., Jacomy, M., 2009. Gephi: an open source software for exploring and manipulating networks. In: Proc. International Conference on Weblogs and Social Media (ICWSM 2009).
- Blondel, V.D., Guillaume, J.L., Lambiotte, R., Lefebvre, E., 2008. Fast unfolding of communities in large networks. J. Stat. Mech.: Theory Exp. 2008, P10008.
- Bond, R.M., Fariss, C.J., Jones, J.J., Kramer, A.D.L., Marlow, C., Settle, J.E., Fowler, J.H., 2012. A 61-million-person experiment in social influence and political mobilization. Nature 489, 295–298.
- Bonetta, L., 2007. Scientists enter the blogosphere. Cell 129, 443–445.
- Borge-Holthoefer, J., Rivero, A., Garcia, I., Cauhe, E., Ferrer, A., Ferrer, D., Francos, D., Iniguez, D., Pilar Perez, M., Ruiz, G., Sanz, F., Serrano, F., Vinas, C., Tarancon, A., Moreno, Y., 2011. Structural and dynamical patterns on online social networks: the Spanish May 15th movement as a case study. PLoS ONE 6, e23883.
- Capstick, S.B., Pidgeon, N.F., 2013. What is climate change scepticism? examination of the concept using a mixed methods study of the UK public. Glob. Environ. Change, Early edition.
- Carvalho, A., 2010. Media(ted) discourses and climate change: a focus on political subjectivity and (dis)engagement. Wiley Interdiscip. Rev.: Clim. Change 1, 172–179.
- Centola, D., 2010. The spread of behavior in an online social network experiment. Science 329, 1194–1197.
- Cha, M., Benevenuto, F., Haddadi, H., Gummadi, K., 2012. The world of connections and information flow in Twitter. Syst. Man Cybern. A: Syst. Hum. 42, 991–998.
- Cha, M., Haddadi, H., Benevenuto, F., Gummadi, P.K., 2010. Measuring user influence in Twitter: the million follower fallacy. In: Proceedings of the Fourth International AAAI Conference on Weblogs and Social Media (ICWSM-10). pp. 10–17.
- Christakis, N.A., Fowler, J.H., 2007. The spread of obesity in a large social network over 32 years. N. Engl. J. Med. 357, 370–379.
- Christakis, N.A., Fowler, J.H., 2008. The collective dynamics of smoking in a large social network. N. Engl. J. Med. 358, 2249–2258.
- Conover, M., Ratkiewicz, J., Francisco, M., Gonçalves, B., Menczer, F., Flammini, A., 2011. Political polarization on Twitter. In: Proc. International Conference on Weblogs and Social Media (ICWSM 2011).
- Conover, M.D., Gonçalves, B., Flammini, A., Menczer, F., 2012. Partisan asymmetries in online political activity. EPJ Data Sci. 1, 1–19.
- Czopp, A.M., 2013. The passive activist: negative consequences of failing to confront anti-environmental statements. Ecopsychology 5, 17–23.
- Edwards, D., Potter, J., 1992. Discursive Psychology. Sage, London.
- Farrell, H., 2012. The consequences of the internet for politics. Annu. Rev. Polit. Sci. 15, 35–52.
- Fowler, J.H., Christakis, N.A., 2008. Dynamic spread of happiness in a large social network: longitudinal analysis over 20 years in the Framingham heart study. BMJ: Br. Med. J. 337.
- Gallois, C., Ogay, T., Giles, H., 2005. Communication accommodation theory: a look back and a look ahead. In: Gudykunst, W. (Ed.), Theorizing About Intercultural Communication. Sage, Thousand Oaks, CA, USA, pp. 121–148.
- Gambhir, V., Kumar, P., 2013. How the people of India live with climate change and what communication can do: Technical Report. BBC Media Action: Climate Asia Project.
- Girvan, M., Newman, M.E., 2002. Community structure in social and biological networks. Proc. Natl. Acad. Sci. 99, 7821–7826.
- Hagberg, A.A., Schult, D.A., Swart, P.J., 2008. Exploring network structure, dynamics, and function using NetworkX. In: Proc. 7th Python in Science Conference (SciPy2008) pp. 11–15.
- Hulme, M., 2009. Why We Disagree About Climate Change: Understanding Controversy, Inaction and Opportunity. Cambridge University Press, Cambridge, UK.
- IPCC, 2013. Summary for Policymakers. In: Stocker, T., Qin, D., Plattner, G.K., Tignor, M., Allen, S.K., Boschung, J., Nauels, A., Xia, Y., Bex, V., Midgley, P. (Eds.), Climate Change 2013: The Physical Science Basis. Contribution of Working Group I to the Fifth Assessment Report of the Intergovernmental Panel on Climate Change. Cambridge University Press, Cambridge, United Kingdom and New York, NY, USA.
- Jackson, M.O., 2010. Social and Economic Networks. Princeton University Press.
- Jacomy, M., Heymann, S., Venturini, T., Bastian, M., 2012. ForceAtlas2: A Continuous Graph Layout Algorithm for Handy Network Visualization. Technical Report. Sciences Po – Medialab. Paris.
- Kahan, D.M., Peters, E., Wittlin, M., Slovic, P., Ouellette, L.L., Braman, D., Mandel, G., 2012. The polarizing impact of science literacy and numeracy on perceived climate change risks. Nat. Clim. Change 2, 732–735.
- Katz, E., 1957. The two-step flow of communication: an up-to-date report on an hypothesis. Pub. Opin. Q. 21, 61–78.
- Katz, E., Lazarsfeld, P.F., 1955. Personal Influence: The Part Played by People in the Flow of Mass Communication. Free Press, Glencoe, IL, USA.
- Kirilenko, A.P., Stepchenkova, S.O., 2014. Public microblogging on climate change: one year of Twitter worldwide. Glob. Environ. Change, Online early edition.
- Kramer, A.D.L., Guillory, J.E., Hancock, J.T., 2014. Experimental evidence of massive-scale emotional contagion through social networks. Proc. Natl. Acad. Sci., Online early edition.
- Leiserowitz, A., Maibach, E., Roser-Renouf, C., 2008. Global Warming's 'Six Americas': An Audience Segmentation. Technical Report Yale University and George Mason University, New Haven, CT Yale Project on Climate Change.
- Leiserowitz, A., Maibach, E., Roser-Renouf, C., Feinberg, G., Rosenthal, S., Marlon, J., 2014. Climate Change in the American Mind: Americans' Global Warming Beliefs and Attitudes in November, 2013. Technical Report Yale University and George Mason University, New Haven, CT Yale Project on Climate Change Communication.
- Levitan, L.C., Visser, P.S., 2009. Social network composition and attitude strength: exploring the dynamics within newly formed social networks. J. Exp. Soc. Psychol. 45, 1057–1067.
- McPherson, M., Smith-Lovin, L., Cook, J.M., 2001. Birds of a feather: homophily in social networks. Annu. Rev. Sociol. 27, 415–444.
- Moser, S.C., 2010. Communicating climate change: history, challenges, process and future directions. Wiley Interdiscip. Rev.: Clim. Change 1, 31–53.
- Muchnik, L., Aral, S., Taylor, S.J., 2013. Social influence bias: a randomized experiment. Science 341, 647–651.
- Newman, M.E., 2006. Modularity and community structure in networks. Proc. Natl. Acad. Sci. 103, 8577–8582.
- Nisbet, M.C., Kotcher, J.E., 2009. A two-step flow of influence? Opinion-leader campaigns on climate change. Sci. Commun. 30, 328–354.
- Nolan, J.M., 2013. Creating a culture of conservation: willingness to confront environmental transgressors. Ecopsychology 5, 3–8.
- O'Connor, B., Balasubramanian, R., Routledge, B.R., Smith, N.A., 2010. From tweets to polls: linking text sentiment to public opinion time series. In: Proc. ICWSM 2010, pp. 122–129.
- O'Neill, S., Williams, H.T., Kurz, T., Wiersma, B., Boykoff, M., 2015. Dominant frames in legacy and social media coverage of the IPCC fifth assessment report. Nat.

- Clim. Change. <http://www.nature.com/nclimate/journal/v5/n4/full/nclimate2535.html>.
- O'Neill, S.J., Boykoff, M., 2010. Climate denier, skeptic, or contrarian? *Proc. Natl. Acad. Sci.* 107, E151.
- Pearce, W., Holmberg, K., Hellsten, I., Nerlich, B., 2014. Climate change on Twitter: topics, communities and conversations about the 2013 IPCC Working Group 1 report. *PLoS ONE* 9, e94785.
- Postmes, T., Spears, R., Lea, M., 2000. The formation of group norms in computer-mediated communication. *Hum. Commun. Res.* 26, 341–371.
- Salganik, M.J., Dodds, P.S., Watts, D.J., 2006. Experimental study of inequality and unpredictability in an artificial cultural market. *Science* 311, 854–856.
- Schafer, M.S., 2012. Online communication on climate change and climate politics: a literature review. *Wiley Interdiscip. Rev.: Clim. Change* 3, 527–543.
- Shalizi, C.R., Thomas, A.C., 2011. Homophily and contagion are generically confounded in observational social network studies. *Sociol. Methods Res.* 40, 211–239.
- Sharman, A., 2014. Mapping the climate sceptical blogosphere. *Glob. Environ. Change* 26, 159–170.
- Sørensen, T., 1957. A method of establishing groups of equal amplitude in plant sociology based on similarity of species and its application to analyses of the vegetation on Danish commons. *Kongelige Danske Videnskabernes Selskab* 5, 1–34.
- Sunstein, C.R., 2007. *Republic.com 2.0*. Princeton University Press Princeton, NJ, USA.
- Swim, J.K., Bloodhart, B., 2013. Admonishment and praise: interpersonal mechanisms for promoting pro-environmental behavior. *Ecopsychology* 5, 24–35.
- Tajfel, H., Turner, J.C., 1979. An integrative theory of intergroup conflict. In: Austin, W.G., Worchel, S. (Eds.), *The Social Psychology of Intergroup Relations*. Brooks/Cole, Monterey, CA, pp. 33–47.
- Tamburrini, N., Cinnirella, M., Jansen, V.A., Bryden, J., 2015. Twitter users change word usage according to conversation-partner social identity. *Soc. Netw.* 40, 84–89.
- Tumasjan, A., Sprenger, T.O., Sandner, P.G., Welpe, I.M., 2010. Predicting elections with Twitter: what 140 characters reveal about political sentiment. In: *Proc. ICWSM 2010*, pp. 178–185.
- Twitter, 2013. *Twitter REST API v1.1 Documentation* (accessed 30.11.13).
- United States Securities Exchange Commission, 2013. *Form S-1 Registration Statement*. Twitter, Inc. (accessed 11.08.14).
- Valente, T.W., 2012. Network interventions. *Science* 337, 49–53.
- Visser, P.S., Mirabile, R.R., 2004. Attitudes in the social context: the impact of social network composition on individual-level attitude strength. *J. Pers. Soc. Psychol.* 87, 779–795.
- Wu, S., Hofman, J.M., Mason, W.A., Watts, D.J., 2011. Who says what to whom on Twitter. In: *Proceedings of the 20th International Conference on World Wide Web*. pp. 705–714.
- Zhao, S., Grasmuck, S., Martin, J., 2008. Identity construction on Facebook: digital empowerment in anchored relationships. *Comput. Hum. Behav.* 24, 1816–1836.