Twitter Discourses on Climate Change: Exploring Topics, Presence of Bots

Authors: Thomas Marlow¹, Sean Miller, J. Timmons Roberts²

Abstract

This study explores the role of Twitter bots (automated users) in online discourse on climate change. We examined 6.5 million tweets posted during the days leading up to and the month following President Donald Trump's June 1, 2017 announcement of the United States' withdrawal from the 2015 Paris Climate Agreement. Of a 10% sample of users, we used the machine learning algorithm "Botometer" to identify likely mechanized "bots." Botometer identified 17,509 suspected bot accounts, representing about 9% of users and 17% of all tweets. Query limits on Botometer and the capacity of STM modeling reduced our final sample size to 167,259 tweets. We then used the 'stm' package in the R statistical programming language to implement structural topic models to cluster tweet content into topics. We identified 34 topics. Topics broadly fell into categories related to news of the withdrawal and the responses from various media and government personnel, posts about climate change research, discussions of the denial of climate change, and finally activist topics with campaign goals. Within these topics, bots were often common, representing as much as 38% of tweets. Additionally, among the most common 15 topics, we found evidence that after adjusting for topic size, timing of tweets, and whether a user was suspended, tweets produced by bots were more likely relative to non-bot users to fall into four topics. These topics included tweets sharing news links on the announcement, links related to climate research, links sharing denialist research, and links about White House aids and their views of fossil fuels. These findings suggest a substantial impact of mechanized bots in amplifying denialist messages about climate change, including support for Trump's withdrawal from the Paris Agreement.

Acknowledgements

We gratefully acknowledge the Brown University Data Science Initiative who provided funding to support the purchase of historical Twitter data.

¹ Sociology PhD Candidate, Brown University, thomas_marlow@brown.edu

² Ittleson Professor of Environmental Studies and Sociology, Brown University, j_timmons_roberts@brown.edu

Introduction

There remains persistently high levels of skepticism in the American public about the scientific credibility of climate science (Brenan and Saad 2018). This is despite near unanimous agreement among the scientific community about the origins, trajectories, and consequences of climate change (IPCC 2013). This public skepticism continues to have influence in policy making decisions at the highest levels: since the election of Donald Trump - an avowed climate skeptic - the United States has withdrawn from the Paris Climate Agreement (IPCC 2018), the Environmental Protection Agency (EPA) has repealed the Clean Power Plan which set emissions restrictions for greenhouse gases, removed all discussion of climate change from official agency publications, ended fuel efficiency standards for vehicles, and actively censored public speaking by top EPA scientists on the subject of climate change. During this process, President Trump and his administration have mobilized false and misleading arguments common among climate change denialists. While views on climate change have long been recognized as polarized, there remain questions about how discourses around climate change are initiated, maintained and disseminated.

Polarization on climate change has been amplified by the fossil fuel industry and related industries who have built a set of organizations to generate uncertainty about the existence of the problem and to delay implementation of solutions (Brulle 2019; Brulle and Aronczyk 2019; Dunlap and McCright 2015). Into this already balkanized situation entered a new social media which has reified and increased social polarization: Twitter. One potent mechanism at work in social media are “social bots” (Ferrara et al. 2016), automated accounts on social media platforms. In their simplest incarnations, bots like or repost content generated by a user or groups of users. More sophisticated bots are able to generate unique content, respond to other user posts, and act in groups to coordinate “astroturf” advocacy campaigns. With this basic functionality, bots are effective at propagating a message much faster and more effectively than a non-automated approach.

Our study explores the as yet unstudied role of bots in Twitter climate change discourse. Twitter has proven to be a useful platform for exploring the discourse on climate change (Jang and Hart 2015; Kirilenko and Stepchenkova 2014; Veltri and Atanasova 2017). At the same time, Twitter’s bot problem is both deeply entrenched and well-documented, with some reports indicating that two thirds of links shared on Twitter are being spread by bots (Wojcik et al. 2018). While there has been substantial research on the climate change discourse and denial on social media, none have looked at the prevalence and influence of twitter bots specifically in shaping this discourse. We ask four basic descriptive questions about the state of the climate change discourse on Twitter: What topics are discussed on twitter about climate change? What is the overall prevalence of bots in this discourse? What is the prevalence of bots in each identified topic?

This study examines tweets using some variation of the words “climate change”, “global warming,” and “paris climate agreement” that were posted during the days leading up to and the month following President Donald Trump’s announcement of the U.S. withdrawal from the Paris Climate Agreement. This was a pivotal moment in the history of societal action on climate: by withdrawing, the world’s largest historical emitter of greenhouse gases from a treaty which had taken years to negotiate, President Trump threatened global momentum on the issue. Months of speculation and pleas from major political and business figures to stay in were overtaken by a faction pushing for Trump to keep his campaign promise to get the U.S. out.

We obtained 6.8 million tweets using some variation of the words “climate change”, “global warming,” and “paris climate agreement” that were posted during the days leading up to and the month following President Donald Trump’s 2017 announcement of the United States’ withdrawal from the 2015 Paris Climate Agreement. Of a 10% sample of those tweets, we used the machine learning algorithm “Botometer” to identify likely mechanized “bots.” Botometer identified 17,509 suspected bot accounts, representing about 9% of users and 17% of all tweets.

We then used the ‘stm’ package in the R statistical programming language to implement structural topic models to cluster tweet content into topics, identifying 34. Topics broadly fell into categories related to news of the withdrawal and the responses from various media and government personnel, posts about climate change research, discussions of the denial of climate change, and finally activist topics with campaign goals. Within these topics, bots were often common, representing as much as 38% of tweets. Additionally, among the most common 15 topics, we found evidence that after adjusting for topic size, timing of tweets, and whether a user was suspended, tweets produced by bots were more likely relative to non-bot users to fall into four topics. These topics included tweets sharing news links on the announcement, links related to climate research, links sharing denialist research, and links about White House aids and their views of fossil fuels (labeled Exxon below).

These findings suggest a substantial impact of mechanized bots in amplifying denialist messages about climate change, including support for Trump’s withdrawal from the Paris Agreement. Before turning to the findings, we first review scholarly findings about polarization in the discourse about climate change, about the role of social media in that polarization, and understanding about the role of mechanized users (bots) in social media. We then describe our data and methods.

Polarization in climate change discourse

Views on climate change and global warming in the United States are distinctly polarized along ideological lines (Brenan and Saad 2018; Carmichael and Brulle 2018; Dunlap, McCright, and Yarosh 2016; Guber 2013; McCright and Dunlap 2011). A 2018 Gallup poll on climate change views found that among Republicans, only 42 percent say they believe climate change is

occurring, as opposed to 86 percent of Democrats (Brenan and Saad 2018). Furthermore, political polarization on views of climate change has worryingly grown over time and displayed evidence of being resistant to change (Carmichael, Brulle, and Huxster 2017; Cook and Lewandowsky 2016; Dunlap et al. 2016; Zhou 2016).

Studies focusing on the sources of this polarization have found that exposure to partisan media is a key driver of this process (Boykoff 2013; Brulle, Carmichael, and Jenkins 2012; Jamieson and Cappella 2008). Using a panel dataset (2001 - 2014) of U.S. national concern about the threat of climate change, Carmichael, Brulle, and Huxster (2017) found that media coverage is one of the most important predictors of Democratic and Republican opinions on climate change. In particular, the more that climate change was covered by news organizations with known ideological positions, the more likely that the partisan group was to respond positively. For example, concern among Democrats about climate change significantly increased as the left-leaning New York Times increased its coverage of climate change. They also found some evidence of what Zhou called the boomerang effect (2016) among Republicans and Democrats. This is the idea that exposure to new information contrary to your established beliefs, actually reinforces your pre-existing notions. For example, the more frequently the left-leaning PBS covered climate change, Republicans had a decline in their belief of the threat of climate change. Similarly, the more often conservative radio-host Rush Limbaugh spoke negatively about belief in climate change, Democrats had increasing concern about the threat of climate change. Carmichael and Brulle found that compared to other explanations of partisan polarization such as differential exposure to science, extreme weather events, and shifts in macro-conditions, the media seemed to explain the largest amount of variance.

Researchers of the media have found that key to this strength is the media's reinforcement of the idea that the science about climate change is uncertain (Dunlap and McCright 2015; McCright and Dunlap 2011). Partially this is the result of how standard print and television media practices interpret and give equal presentation of opposing views. While perhaps unintended, this creates the impression among media consumers of scientific uncertainty (Boykoff and Boykoff 2004, 2007; Brüggemann and Engesser 2017; Oreskes and Conway 2010). However, there is also clear evidence that some parts of the media and political actors are delivering a message of uncertainty regarding climate change as part of an organized effort (Brulle 2014; Dunlap and McCright 2015). In general, research on the denial counter-movement finds that uncertainty regarding climate change is being generated by interested political, social, and intellectual elites. Brulle for example, found that 91 denialist organizations including corporations and their allied trade association, conservative think tanks, and advocacy groups received more than \$7 billion between 2003-2010 (Brulle 2014). Furthermore, in a separate analysis Brulle found that funding towards lobbying on Capitol Hill also represents a significant pathway for money aimed at obfuscating the reality of climate change at the highest levels of decision making (Brulle 2018). Some of this money comes directly from large well known conservative donor groups such as the Koch Affiliated Foundations and ExxonMobil, but worryingly, an increasing amount of this money is being donated anonymously (Brulle 2014). Far from being independent flows of money, tracing the various connections between donors

and denial organizations leads to a dense, and clearly structured network of organizations working towards a common goal.

The flow of money through elite funding networks has real impacts on both the network structure of organized climate denial, but also by qualitatively changing the content produced by denialist organizations (Farrell 2016a). Farrell found that between 1993 and 2013, denialist organizations that received corporate funding, produced texts that were significantly different from other non-funded denial organizations. This funding also fundamentally shaped the network of denialist organizations, with funded organizations gradually coming to occupy central importance over time and have their ideas reflected in mainstream media and politician's speech on climate change (Farrell 2016b, 2016a).

In sum, a handful of self-interested and wealthy elites and corporations have actively worked to fund the manufacturing of partisan uncertainty and skepticism about the science and media reporting about climate change. The success of this work cannot be denied: 69 percent of Republicans believe that the climate change as an issue is exaggerated by the media whereas only 4 percent of Democrats believe this to be true (Brenan and Saad 2018).

Social Media Polarization

Social media has been a relatively recent addition to the media landscape that is getting more attention as both a representation and source of public polarization on a range of issues (see Barberá et al. 2018 for a review of this large body of literature). Studies focusing on climate change in particular have largely found that social media networks and discourse is reflective of the broader social environment – that is to say it is highly polarized (Kirilenko and Stepchenkova 2014; Pearce et al. 2014; Swain 2016, 2017; Williams et al. 2015). Williams et al 2015 for example, analyzed the follower, retweet and mention networks of users coded on a spectrum of climate activist to climate skeptic, Williams and colleagues found that these networks showed strong signs of developing echo chambers - that is, users of a particular ideology were primarily engaging with users who had similar views about climate change.

Beyond just reflecting the media environment at large however, other research has highlighted ways Twitter and other social media platforms may work to promote polarization and ultimately harm democracy (Barberá et al. 2018; Sunstein 2018; Tucker et al. 2018). Research has identified three mechanisms: first, social media exposes people to uncivil conversations around contentious issues which leads to increases in affective polarization (Lelkes 2016; Suhay, Bello-Pardo, and Maurer 2018; Weeks 2015); second, by generating a fragmented news environment that lowers the overall quality and creates spaces for the spread of disinformation (Lazer et al. 2018; Vosoughi, Roy, and Aral 2018) and incorporation of otherwise fringe views (Bail 2012; Farrell 2016a); third, it exposes users to larger number of opposing viewpoints which can activate the type of “boomerang” or “hostile media” response seen in general media (Bail et al. 2018).

With regards to climate change, there is some evidence that at least two of these mechanisms could be at work. In one study, Anderson and Huntington (Anderson and Huntington 2017) hand coded a small group of tweets ($n = 4,094$) about a climate related weather disaster in Colorado. They found that while incivility was relatively rare within the full sample - just under 7% of all tweets - of those tweets that were uncivil, they were most likely to occur while discussing politics related to climate change. They also found that both sarcasm and incivility were significantly more likely to occur during comments that contained skeptical views of climate change. This finding strengthened the earlier findings of Williams and colleagues (2015) who used a much larger sample of tweets ($n = 590,608$) about global warming and climate change and found that when partisans on either side of the climate change issue interacted, there was significantly higher amounts of negative sentiment represented.

Content and network analysis of the discourse on climate change on twitter have also found a fragmented and heterogeneous environment of content producers and consumers (Kirilenko and Stepchenkova; Swain 2016; 2017; Williams et al 2015). Kirilenko and Stepchenkova (2014) for example, identified over 46 thousand unique domains linked in tweets about climate change and global warming in 2012. This again, is the type of media environment where misinformation, especially unique misinformation (Vosoughi, Roy and Aral 2018), is able to spread easily. Finally, while no research has looked directly at how people respond to exposure to information contrary to their personal beliefs, there is growing experimental evidence that this type of response is at work in other issues on social media. In a large scale experiment using Twitter users, Bail et al (2018) found that for both self-identified Democrats and Republicans, exposure to a treatment of from the opposite ideology lead to these groups to become more liberal in the case of Democrats, and conservative in the case of Republicans. Therefore on the issue of climate change, where similar effects have been demonstrated in other media types (Carmichael et al 2017), it seems quite likely that this mechanism is also at work.

Bots

Into this social media environment with an elite funded network of organized climate denial organizations, enters a new inexhaustible tool principally designed to rapidly promote, generate, and disseminate messages on social media - the bot. A bot is an algorithm designed to produce content on social media without human input in ways that facilitates interaction with humans (Ferrara et al. 2016). Bots are commonly used by organizations to post content at regular intervals or retweet content related to their specific organization. However, Ferrara et al (2016) summarise the threat of social bots:

“The novel challenge brought by bots is the fact they can give the false impression that some piece of information, regardless of its accuracy, is highly popular and endorsed by many, exerting an influence against which we haven’t yet developed antibodies. Our vulnerability makes it possible for a bot to acquire significant influence” (2016:97-98).

This is made all the more problematic by the increasing sophistication of social bots.

Research has found that humans have difficulty reliably distinguishing between human produced text and bot produced text (Edwards et al. 2014; Everett, Nurse, and Erola 2016; Wang, Angarita, and Renna 2018). Bots are able to produce text that is coherent, and employ strategies like producing text that is counter to the crowd opinion, which boosts their perceived reliability (Everett, Nurse, and Erola 2016).

In sum, the bot is perfectly designed to tirelessly execute propaganda and astroturf campaigns which sow negative sentiment and doubt about an issue. Unfortunately this is not just a warning, social bots have been found to be influential in a wide range of cases including political elections where they have been influential since at least 2010 U.S. midterm elections (Ratkiewicz, et al. 2011). Most recently, social bots were estimated to generate an estimated 3.8 million tweets (19% of the total), about the 2016 U.S. Presidential election (Bessi and Ferrara 2016). To date however, there is no research on the role of bots in producing and disseminating misinformation regarding climate change. Polarization regarding climate change broadly and on social media, as well as a motivated and organized denial machine however, threatens both the presence of social bots, and their having an impact on ongoing polarization. Our study takes the first step in understanding the role of social bots in the discourse of climate change online. We describe the prevalence of bots within distinct domains of discourse around global warming and climate change.

Data

Data for this project comes from two sources. First, the tweets about the climate paris agreement come directly from Twitter via a historical data request using keywords that include combinations and small variations on “climate change”, “global warming”, and “paris climate agreement”. There are two advantages of this approach. There was some concern that these keywords are not representative of the full political spectrum (Barberá and Rivero 2015). That is, we did not have strong prior knowledge about whether both ends of the spectrum are using these same terms in online discussions. To test this, in a pilot analysis of 238,808 from 2018, we identified users from lists of pro-and anti climate change users to generate the most common hashtags, and the keywords used in each group based on the most recent 1000 tweets for each user. For both groups, the most common terms were some variation of “climate change” and “global warming”. Other popular terms included “carbon”, “resilience”, and “weather”. However, data quality including these terms was very poor, and topic models including these terms returned a majority of areas discussing irrelevant areas such as carbon-fiber materials science, cryptocurrency blockchain resilience, and weather reporting for major cities. Filtering the results to only include the initial, simpler query resulted in a representative sample of the Twitter climate discourse.

Our original dataset covered the period of May 1 - June 31 2017. The announcement by President Donald Trump occurred on June 1, 2017 (Figure 1). By Sunday May 28th, the media

began to publish reports that Donald Trump had decided to withdraw the U.S. from the Paris Climate Accord. This is reflected in social media reports on the issue. This too is where we begin sampling data for analysis. During the period of May 1 to June 31, there were over 6.8 million tweets in our data. Over 3.8 of these were retweets. Limiting our study period to May 28th when the majority of reporting began, practically limits the size of the data we have to work with, and theoretically limits our study to understanding the discourse immediately prior to and during the month following Trump's announcement. Even so, this limitation includes 6.5 million tweets, and 1.6 million users.

The second source of data, is a predicted probability that a given user is a bot. This score comes from the Botometer API developed at the Indiana University Observatory on Social Media (OSoMe), a joint project with the Network Science Institute and Center for Complex Networks and Systems Research (Varol et al. 2017). Their freely available API allows users to submit twitter screen names and get a json object of scores in return. We write these responses to a SQL database for future use. Rate limits imposed by Botometer and Twitter are however a major limitation. Instead of classifying the entire dataset including 1.6 million users, we took an initial 10% random sample, and an additional users classified within our time-frame for a total of 184,767 users responsible for 885,164 tweets. STM model convergence was slow using even a sample of this size, so we took two additional steps. First, we took a random 20% sample of the remaining tweets and pruned the vocabulary words present in at least 10 tweets About 80,000 words are removed during this process which dropped an additional 1,700 documents. The final sample size is 167,259 tweets, none of which are retweets (See Table 1).

Figure 1: Climate Change Tweets, May-June 2017

Methods

Detecting Bots

We used Botometer (<https://botometer.iuni.iu.edu>) to estimate the probability a user is a bot. Botometer is a collaborative effort between two departments at Indiana University: the Network Science Institute and the Center for Complex Networks and Systems Research. Given a Twitter handle, Botometer extracts up to 200 of their most recent tweets and up to 100 of their most recent mentions along with metadata associated with the user and their tweets. This metadata falls into six categories including: network features, user features, friend features, temporal features, content, and sentiment. Each category is evaluated separately by Botometer for its similarity to bot like behavior. Network evaluation is based on the distribution of likes, retweets, and hashtag usage. User feature evaluation is tied to the user's metadata such as when the account was created, where the account was created, and whether or not the account is verified. Friends features are evaluated based on the distribution of their friends list and the distribution of those friends' tweets. Temporal features evaluates a user based on the timing between retweets and responses. Content features are English-exclusive features based on natural language processing methods to determine bot likelihood by looking at the length of the tweets themselves and tagging parts of speech within the tweet and comparing their distributions to those generated by bots. Sentiment features are also English-exclusive and determine the general emotion behind a user's tweets, building upon established literature by comparing arousal, valence, and dominance scores to prior distributions. These features are

applied to a boosted random forest classifier. Botometer computes a score for each category and then weights them to calculate a final score between 0 and 1.

Botometer can achieve an AUROC (area under the receiver operating characteristic curve) of as high as 0.93 (Varol et al. 2017). AUROC is a metric used to evaluate the performance of binary classifiers. In this context, this value means that 93% of the time, a bot account will receive a higher Botometer score than a non-bot account. One shortcoming of Botometer and all algorithms like it however is that they frequently misclassify organizational or institutional accounts as bots. This can be attributed to the fact that these accounts are often partially or fully automated and thus exhibit similar behavior to bot accounts in user, temporal, and content features. To limit this source of bias, we first chose an appropriate cutoff (between 0 and 1) between bot and non-bot users that balances false-positives and false-negatives. Following previous research that used human coding checks to estimate optimal cutoff points (Wojcik et al. 2018), we used a cutoff of .43. We also inspected influential bot accounts to provide an informal validity check. Both of these steps indicate that while some organizations are classified as bots, these are not common enough to significantly impact our estimates of bot prevalence or influence in the climate change discourse. With this tool we classified 184,767 users and identified 17,509 users as suspected bots (Table 1).

Table 1: Tweets and Users Classified Using Botometer

	Count	Percentage
Sample Tweets	885,164	-
Users	184,767	-
Suspected Bots	17,509	9.47%
Suspected Bot Tweets	157,425	17.78%
Suspended Users	3,863	2.10%
Suspended User Tweets	21,474	2.43%

Structural Topic Modeling

To cluster the climate change discourse on Twitter, we used a Structural Topic Modeling (STM) approach implemented with the R package “stm” (Roberts, Stewart, and Tingley 2019). STM is a recently-developed method of clustering text objects that builds on the more common latent Dirichlet allocation (LDA) model by assessing not only textual similarity (content) but also metadata features (prevalence) like date of publication or source. Farrell for example used

continuous metadata about the date and time of publication to see how the text produced by climate change denial organizations evolved over time (Farrell 2016b).

The algorithm used by STM is an expectation-maximization process. STM assumes documents are generated from topics in the following way. First, the distribution of topics is generated by a logistic normal distribution (as opposed to the Dirichlet distribution used by LDA. This usage of the lognormal distribution at this step is what separates STM from LDA. Unlike the lognormal distribution, the Dirichlet distribution is parametrized by a vector of word priors, as opposed to a mean vector and variance matrix. This parameterization allows for covariance to be present between topic proportions, unlike LDA, which assumes that these proportions are independent. It then uses the empirical word distribution and the document-specific metadata features to construct a logistic regression model for the likelihood of each word in each topic. Finally, it generates words, first conditioned on topic proportions, then conditioned on the document-specific word distribution for each topic. Each step of the EM process tunes these parameters until the values that maximize the likelihood of the observed topics being generated are found.

We omit retweets during the model's training, as these acted as duplicate documents that would center the topics around heavily-retweeted texts. We do, however, include document-level metadata to model topic formation. In particular, we model topics as a function of the time at which a tweet is published, their user id, an indicator of whether a user is a bot or not, and an indicator of whether the user was suspended some time after the study period. Thus, topics are dependent on these factors. This is advantageous because it allows us to better distinguish topics that over time and contain different proportions of bots. Additionally, the decision was made to only include tokens that appeared in 25 or more tweets. This was done so that words that appeared in only a few tweets would not be unduly associated with the topic(s) in which they appear.

Choosing the number of topics

An important source of researcher input is the decision on the appropriate number of topics (k). We use an approach to choosing the number of topics that blends computational and interpretive approaches. STM offers a number of useful metrics when selecting the number of topics. The three most important of these are *semantic coherence*, *exclusivity*, and *variational lower bound*. Semantic coherence is a measure of the “dependence” of words in a topic. It scores topics based on the probability that a topics most common words appear together, with common word co-occurrence scoring highly. The literature has empirically shown that this measure is a suitable proxy for human judgments on topic quality (Mimno et al. 2011). Exclusivity scores topics highly when words appear exclusively within a single topic. Finally, the held-out likelihood is the likelihood of a prediction using 5,000 tweets held-out and residuals are the error produced by this prediction.

There is a trade-off between semantic coherence and exclusivity. As the number of topics increases, semantic coherence initially increases followed by a general decrease while exclusivity generally increases. This makes intuitive sense; with a small k , topics will be formed around groups of co-occurring words that may not be exclusive to their topic, while with a large k , topics will be formed around individual words without any other correlated ones. Variational lower bound behaves similarly to exclusivity, increasing with the number of topics with diminishing returns. The optimal strategy therefore is to evaluate these metrics for each possible k , starting at $k=2$ and continuing until the increase in exclusivity begins to be dwarfed by the decrease in semantic coherence. At this point then, a human examiner can enter the process and select the k in this range according to the examiner's expert opinion and the needs of the research problem.

Applying this approach to our corpus, an optimal number of topics was determined to be 34 (See Figure 1). As can be seen by the plots, semantic coherence declines continuously while exclusivity increases rapidly until $K=40$ at which point, increases become much smaller although they never fully maximize. Plotting the scatter plot of topic scores according to both exclusivity and semantic coherence (not shown) indicates that the optimal balance between the two is between 34 and 40, inclusive. This furthermore matches the results of the predicted model using held out data. While the likelihood never maximized, the residuals for the model were smallest between 34 and 36, inclusive. At this point, the topic quality for each model was examined, and the decision was made to choose the lowest K in the range. In this case, that K ended up to be 34. STM considers each document to be comprised of a mixture of multiple topics; however, for the purpose of analysis, each tweet was assigned the topic value for which it was most prevalent.

Figure 2: Model Diagnostics by Number of Topics

Results

Prevalence of Bots in the Climate Discourse:

For the period May 28- June30, 2017, we classified 184,767 unique users using Botometer and identified 17,509 users as suspected bots (Table 1).

Figure 3: Frequency of Tweets about Climate Change with Prevalence as Frequency and Proportion, May 28-June 30, 2017.

On a typical day, bots produce on an average day 25 percent of all tweets about climate change (Figure 2). However during the spike in activity around President Trump's announcement of withdrawal from the Paris Agreement, this proportion dropped by about half, to about 12-13 percent. While tweets produced by suspected bots did increase substantially during the few days around Trump's announcement (from hundreds a day to over 25,000 a day), they did not do so proportionally relative to the increase in tweets overall (this includes retweets). This suggests that suspected bots (as detected by Botometer), are relatively static in their behavioural patterns, and unable to respond to events which produce large numbers of activity in a short amount of time.

Figure 4: Expected topic proportions with each topic labeled with a keyword, and three words with the highest FREX score. FREX is a measure that combines the probability a word is in a topic, with the word's exclusivity to the topic.

Structural topic models identified 34 topics which were discussed in the 167,259 tweets we classified (see Methods section for details). STM treats each tweet as a “bag of words,” and classifies these into “topics” which tend to include the same sets of words. We reviewed the fifty tweets with the highest probabilities of being in each topic.

The most frequent tweets in the days around President Trump’s withdrawal from the Paris Agreement fell into topics which could be characterized as both supportive and outraged at his actions. Many were simply reporting his action, and a number discussed just one aspect of the decision. Some were simply reporting climate science news, as was seen in the most common topic (21). Many had identical text, especially in cases where users appeared to be tweeting out a new story directly from the news outlet’s website. This is the case with the second most frequent topic as well, but this one from the side of those approving of the president’s action. The third most frequent topic appeared to be from an action alert of a climate NGO, calling on readers to resist Trump’s action. The repetitive nature of many tweets reflects how Twitter is largely one directional, with interactions and discussion fairly unusual.

Figure 5: Predicted Probability of Topic Membership by User Type.

In which Twitter topics are bots more likely to be an important presence? Figure 5A displays the point estimate and confidence intervals of the predicted probability a tweet will be assigned to the top 15 most probable topics, given that that a user is labeled a suspected bot or not and controlling for other confounders. The vertical line displays the average effect size. The topics

are ordered by their expected proportions (See Figure 4) so that, on average, any tweet is more likely to be identified in the most common topics (e.g. 21 or 3).

Given that they are only about 18 percent of all tweets during our sample period, those from Twitter bots are less frequent than those from human users in a majority (8) of the top 15 topics (Figure 5B). Put more precisely, being labeled a bot lowers the probability of being in those topics. Three more topics had nearly identical probability of tweets being sent by humans or bots (14 denial news, 30 warning, 32 thanks). The overall picture, then, is that bots are focused on driving the frequency of some topics much more than others. In only four of the top 15 topics does being labeled a bot make it more likely a tweet is in a topic.

But in some cases, bots are highly common. Of the four topics where bots are most focused, two form the largest topics overall, and cover the subjects of announcement news and climate research. The gap between users is also the largest in these topics. Being labeled a bot increases the probability of being in topic 21 by over 3% and increases the probability of being in topic 3 by over 6%. For topic 20 (“resistance”), the opposite is true. Bot users were less likely to produce tweets in this topic. This finding suggests that the “resistance” had large numbers of real users tweeting relative to bots.

It is informative to look at the topics where bots are more likely to be classified than human users. There are two other topics in the top 15 where bot users are more common. Topic 8, labeled “denial research,” includes tweets promoting research which is interpreted as contradicting the scientific consensus on human-caused climate change. As indicated by the top FREX words and a selection of tweets, these often revolve around the research purporting to find global cooling or stable temperatures. A user labeled as a bot has on average a .5% greater probability of being in topic 8. Topic 2, is characterized by large numbers of tweets directed at French President Emmanuel Macron. On average, a user labeled as a bot has a .5% greater probability of being in topic 2.

Figure 5B shows the raw counts of tweets within a topic, categorized by tweets produced by suspected bots or other users. Following Figures 4 and 5A, the plot is ordered by the expected proportions of each topic within the corpus. However, there is some discrepancy between the expected proportions and counts in 5B. Topic 8 which includes discourse on denial research for example, has low expected proportions but has noticeably higher numbers of tweets than its neighbors in Figure 5B. This discrepancy comes from the way topics are assigned. Each document (tweet) is a mixture of words from different topics. To assign a topic, we give a document the value which corresponds to the largest portion of its words. So if 75% of the words come from topic 21 and 25% from topic 3, the document is assigned to topic 21. While topic 8 is present in fewer documents than many above it, it is the largest proportion in those documents in which it appears.

There are a few notable things about Figure 5B. First, tweets from bot accounts are never more common than other user tweets. This is not surprising, given that bots account for an estimated

17% minority share of all tweets during the study period. This finding suggests that at least for the online discourse around a major media event like Trump's announcement of withdrawal from the Paris Climate Agreement, regular users are active enough to overwhelm bot activity. We should be careful not to generalize this finding beyond the context of a major news event. As evidenced in Figures 3A-B, non-bot tweets spiked during the days around the announcement causing the proportion of tweets from bots to dip. The rest of the time, bots are producing on average one quarter of all tweets on climate change and global warming. Furthermore, we know from other research, bots can be deployed in more targeted ways to overwhelm a discourse (Wang et al. 2018)

Second, bots are not active evenly across topics. The three topics with the highest numbers of bot tweets were on climate research (topic 21) with 12,516 tweets generated by bots, news links (topic 3) with 9,433 tweets generated by bots, and denial research promotion (topic 8) with 1,908 tweets generated by bots. Top topics by proportion of tweets from bots were 25 (Fake Science, 38%), 6 (How Coffee Cultivation will be Affected by Climate Change, 30%), 2 (Exxon, 28%), 7 (No Consistent Theme, 28%), 3 (News, 26%), and 9 (International Response, 25%). Interestingly, the third most common topic overall, topic 20, had only 5% of its tweets produced by users suspected of being a bot. This matches the finding in Figure 5A that being labeled a bot lowers the probability that a tweet will fit in a topic characterized as focusing on the resistance. We labeled topic 20 as "resistance" because the top tweets were online activism aimed at building solidarity around the Paris Climate Agreement and urging government leaders to support climate change legislation. Hashtags like #actonclimate and #resist were common and their stemmed versions were among the top words measured by FREX (See Figure 3). The low percentage of bots within this topic, suggests that this activism was a genuinely popular movement of individuals. Despite bots' demonstrated skill at promoting activist messages, they were not employed widely.

Overall, these findings suggest that bots are not just prevalent, but disproportionately so in topics that were supportive of Trump's announcement or skeptical of climate science and action.

Discussion and Conclusion

Bots are having a significant impact upon discourse about climate change on Twitter, and that impact is not random. This research has documented a massive surge in the Twitter discourse on climate change in the immediate aftermath of President Trump's announcement of the U.S.'s planned withdrawal from the Paris Agreement. In the two months of May and June, 2017, over 6.8 million tweets discussed the topic, with about 3.8 million of those being retweets. On a typical day, about 25 percent of all tweets on climate change come from users that exhibit patterns of use that are characteristic of mechanized bots. That proportion dropped in the immediate aftermath of Trump's announcement, but the absolute number of bot tweets soared. Suspected bot tweeting was focused on a relatively small number of topics (4 of the top 15 most

frequent topics), and those topics were more likely to be appreciative of Trump's action, or skeptical or denying of the wide scientific consensus on climate change.

To return to our main theoretical points and contribution to the social science literature on climate denial, polarization, and the role of social media and bots, this study informs several issues. First, We found a distinct divide in the way people are discussing climate change--this study confirms the polarization that exists in American discourse on the issue. Top topics condemned or praised the President for taking this drastic step; even climate news was divided in how it was reported. An important finding is that suspected mechanized Twitter users (bots) were more likely to be present amplifying denialist discourses than others.

These findings are in line with and expand upon those from the small but emerging literature on climate change discourses in media and social media. Wetts (2019) and Farrell. Social media is highly polarized, not just reflecting larger media environment, but driving more polarization in very specific ways through specific mechanisms (Barberá et al. 2018), (these are consistent with our findings, uncivil debate in discourse with contention). Social media exposes people to uncivil conversations on contentious issues (Lelkes 2016; Suhay et al. 2018). Social media generates a more fragmented news environment, more sources with lower quality and channels for bad or fake news (Lazer et al. 2018). Social media exposes users to opposing viewpoints which can activate the boomerang effect of entrenching/reinforcing extreme positions (Bail et al. 2018).

Bots, are tools designed to generate, disseminate, and promote positions on social media. Because they can impersonate trusted sources effectively, bots are sometimes more effective than human producers of social media content at driving public opinion (Edwards et al. 2014; Everett et al. 2016; Wang et al. 2018). This study provides a first look at bots in these discourses, finding that some discourses have much greater presence of bots, and those seem to be associated with letter-writing campaigns. Influence on policy-makers is future research.

There are several caveats and limitations of this research that bear remembering as we conclude. The data collected used very general search terms ("climate change," "global warming," "Paris," "climate," and "agreement"); different search terms might reveal a new set of topics, and there is no doubt that millions more tweets never used these terms. The tweets collected were also from a very short and unusual period of time; it would be revealing to examine historical patterns over longer periods. Using mechanical sorting of tweets into topics also raises a number of methodological concerns; there would be value in comparing the findings to hand-coded, deductive categories.

But the findings are striking: discourses on climate change on the highly influential social media platform are being manipulated by mechanized users. Their influence is selectively focused on more skeptical and denialist discourses. The implications for social movements, democracy, and the ability of our society to confront the existential crisis of climate change is at stake.

References Cited

- Anderson, Ashley A. and Heidi E. Huntington. 2017. "Social Media, Science, and Attack Discourse: How Twitter Discussions of Climate Change Use Sarcasm and Incivility." *Science Communication* 39(5):598–620.
- Bail, Christopher A. 2012. "The Fringe Effect: Civil Society Organizations and the Evolution of Media Discourse about Islam since the September 11th Attacks." *American Sociological Review* 77(6):855–79.
- Bail, Christopher A., Lisa Argyle, Taylor Brown, John Bumpus, Haohan Chen, M. B. Fallin Hunzaker, Jaemin Lee, Marcus Mann, Friedolin Merhout, and Alexander Volfovsky. 2018. "Exposure to Opposing Views Can Increase Political Polarization: Evidence from a Large-Scale Field Experiment on Social Media." *SocArXiv*.
- Barberá, Pablo, Joshua A. Tucker, Andrew Guess, Cristian Vaccari, Alexandra Siegel, Sergey Sanovich, Denis Stukal, and Brendan Nyhan. 2018. *Social Media, Political Polarization, and Political Disinformation: A Review of the Scientific Literature*. California, USA: William + Flora Hewlett Foundation.
- Boykoff, Maxwell T. 2013. "Public Enemy No. 1?: Understanding Media Representations of Outlier Views on Climate Change." *American Behavioral Scientist* 57(6):796–817.
- Boykoff, Maxwell T. and Jules M. Boykoff. 2004. "Balance as Bias: Global Warming and the US Prestige Press." *Global Environmental Change* 14(2):125–36.
- Boykoff, Maxwell T. and Jules M. Boykoff. 2007. "Climate Change and Journalistic Norms: A Case-Study of US Mass-Media Coverage." *Geoforum* 38(6):1190–1204.
- Brenan, Megan and Lydia Saad. 2018. "Global Warming Concern Steady Despite Some Partisan Shifts." *Gallup.Com*. Retrieved August 31, 2018 (<https://news.gallup.com/poll/231530/global-warming-concern-steady-despite-partisan-shifts.aspx>).
- Brüggemann, Michael and Sven Engesser. 2017. "Beyond False Balance: How Interpretive Journalism Shapes Media Coverage of Climate Change." *Global Environmental Change* 42:58–67.
- Brulle, Robert J. 2014. "Institutionalizing Delay: Foundation Funding and the Creation of U.S. Climate Change Counter-Movement Organizations." *Climatic Change* 122(4):681–94.
- Brulle, Robert J. 2018. "The Climate Lobby: A Sectoral Analysis of Lobbying Spending on Climate Change in the USA, 2000 to 2016." *Climatic Change* 149(3):289–303.
- Brulle, Robert J. 2019. "Networks of Opposition: A Structural Analysis of U.S. Climate Change Countermovement Coalitions 1989–2015." *Sociological Inquiry* 50.12333.
- Brulle, Robert J. and Melissa Aronczyk. 2019. "Organised Opposition to Climate Change Action in the United States." *Routledge Handbook of Global Sustainability Governance* 145.
- Brulle, Robert J., Jason Carmichael, and J. Craig Jenkins. 2012. "Shifting Public Opinion on Climate Change: An Empirical Assessment of Factors Influencing Concern over Climate Change in the U.S., 2002–2010." *Climatic Change* 114(2):169–88.
- Carmichael, Jason T. and Robert J. Brulle. 2018. "Media Use and Climate Change Concern." *International Journal of Media & Cultural Politics* 14(2):243–53.
- Carmichael, Jason T., Robert J. Brulle, and Joanna K. Huxster. 2017. "The Great Divide: Understanding the Role of Media and Other Drivers of the Partisan Divide in Public Concern over Climate Change in the USA, 2001–2014." *Climatic Change* 141(4):599–612.
- Cook, John and Stephan Lewandowsky. 2016. "Rational Irrationality: Modeling Climate Change

- Belief Polarization Using Bayesian Networks.” *Topics in Cognitive Science* 8(1):160–79.
- Dunlap, Riley E. and Aaron M. McCright. 2015. “Challenging Climate Change: The Denial Counter-movement.” Pp. 300–332 in *Climate Change and Society*. New York: Oxford University Press.
- Dunlap, Riley E., Aaron M. McCright, and Jerrod H. Yarosh. 2016. “The Political Divide on Climate Change: Partisan Polarization Widens in the U.S.” *Environment: Science and Policy for Sustainable Development* 58(5):4–23.
- Edwards, Chad, Autumn Edwards, Patric R. Spence, and Ashleigh K. Shelton. 2014. “Is That a Bot Running the Social Media Feed? Testing the Differences in Perceptions of Communication Quality for a Human Agent and a Bot Agent on Twitter.” *Computers in Human Behavior* 33:372–76.
- Everett, Richard M., Jason R. C. Nurse, and Arnau Erola. 2016. “The Anatomy of Online Deception: What Makes Automated Text Convincing?” Pp. 1115–20 in *Proceedings of the 31st Annual ACM Symposium on Applied Computing - SAC '16*. Pisa, Italy: ACM Press.
- Farrell, Justin. 2016a. “Corporate Funding and Ideological Polarization about Climate Change.” *Proceedings of the National Academy of Sciences* 113(1):92–97.
- Farrell, Justin. 2016b. “Network Structure and Influence of the Climate Change Counter-Movement.” *Nature Climate Change* 6(4):370–74.
- Ferrara, Emilio, Onur Varol, Clayton Davis, Filippo Menczer, and Alessandro Flammini. 2016. “The Rise of Social Bots.” *Commun. ACM* 59(7):96–104.
- Guber, Deborah Lynn. 2013. “A Cooling Climate for Change? Party Polarization and the Politics of Global Warming.” *American Behavioral Scientist* 57(1):93–115.
- IPCC. 2018. *Global Warming of 1.5°C. An IPCC Special Report on the Impacts of Global Warming of 1.5°C above Pre-Industrial Levels and Related Global Greenhouse Gas Emission Pathways, in the Context of Strengthening the Global Response to the Threat of Climate Change, Sustainable Development, and Efforts to Eradicate Poverty*.
- Jamieson, Kathleen Hall and Joseph N. Cappella. 2008. *Echo Chamber: Rush Limbaugh and the Conservative Media Establishment*. First Edition edition. Oxford ; New York: Oxford University Press.
- Jang, S. Mo and P. Sol Hart. 2015. “Polarized Frames on ‘Climate Change’ and ‘Global Warming’ across Countries and States: Evidence from Twitter Big Data.” *Global Environmental Change* 32:11–17.
- Kirilenko, Andrei P. and Svetlana O. Stepchenkova. 2014. “Public Microblogging on Climate Change: One Year of Twitter Worldwide.” *Global Environmental Change* 26:171–82.
- Lazer, David M. J., Matthew A. Baum, Yochai Benkler, Adam J. Berinsky, Kelly M. Greenhill, Filippo Menczer, Miriam J. Metzger, Brendan Nyhan, Gordon Pennycook, David Rothschild, Michael Schudson, Steven A. Sloman, Cass R. Sunstein, Emily A. Thorson, Duncan J. Watts, and Jonathan L. Zittrain. 2018. “The Science of Fake News.” *Science* 359(6380):1094–96.
- Lelkes, Yphtach. 2016. “Mass Polarization: Manifestations and Measurements.” *Public Opinion Quarterly* 80(S1):392–410.
- McCright, Aaron M. and Riley E. Dunlap. 2011. “The Politicization of Climate Change and Polarization in the American Public’s Views of Global Warming, 2001–2010.” *The Sociological Quarterly* 52(2):155–94.
- Mimno, David, Hanna M. Wallach, Edmund Talley, Miriam Leenders, and Andrew McCallum. 2011. “Optimizing Semantic Coherence in Topic Models.” Pp. 262–272 in *Proceedings of the Conference on Empirical Methods in Natural Language Processing, EMNLP '11*.

- Stroudsburg, PA, USA: Association for Computational Linguistics.
- Oreskes, Naomi and Erik M. Conway. 2010. "Defeating the Merchants of Doubt." *Nature*. Retrieved September 3, 2018 (<https://www.nature.com/articles/465686a>).
- Pearce, Warren, Kim Holmberg, Iina Hellsten, and Brigitte Nerlich. 2014. "Climate Change on Twitter: Topics, Communities and Conversations about the 2013 IPCC Working Group 1 Report." *PLOS ONE* 9(4):e94785.
- Roberts, Margaret E., Brandon M. Stewart, and Dustin Tingley. 2019. "stm: An R Package for Structural Topic Models." *Journal of Statistical Software* 91(1):1–40.
- Suhay, Elizabeth, Emily Bello-Pardo, and Brianna Maurer. 2018. "The Polarizing Effects of Online Partisan Criticism: Evidence from Two Experiments." *The International Journal of Press/Politics* 23(1):95–115.
- Sunstein, Cass R. 2018. *#Republic: Divided Democracy in the Age of Social Media*. Updated edition. Princeton, N.J. Oxford: Princeton University Press.
- Swain, John. 2016. "Mapped: The Climate Change Conversation on Twitter." *Carbon Brief*. Retrieved September 5, 2018 (<https://www.carbonbrief.org/mapped-the-climate-change-conversation-on-twitter>).
- Swain, John. 2017. "Mapped: The Climate Change Conversation on Twitter in 2016." *Carbon Brief*. Retrieved September 5, 2018 (<https://www.carbonbrief.org/mapped-the-climate-change-conversation-on-twitter-in-2016>).
- Tucker, Joshua, Andrew Guess, Pablo Barbera, Cristian Vaccari, Alexandra Siegel, Sergey Sanovich, Denis Stukal, and Brendan Nyhan. 2018. "Social Media, Political Polarization, and Political Disinformation: A Review of the Scientific Literature." *SSRN Electronic Journal*.
- Varol, Onur, Emilio Ferrara, Clayton A. Davis, Filippo Menczer, and Alessandro Flammini. 2017. "Online Human-Bot Interactions: Detection, Estimation, and Characterization." *ArXiv:1703.03107 [Cs]*.
- Veltri, Giuseppe A. and Dimitrinka Atanasova. 2017. "Climate Change on Twitter: Content, Media Ecology and Information Sharing Behaviour." *Public Understanding of Science* 26(6):721–37.
- Vosoughi, Soroush, Deb Roy, and Sinan Aral. 2018. "The Spread of True and False News Online." *Science* 359(6380):1146–51.
- Wang, Patrick, Rafael Angarita, and Ilaria Renna. 2018. "Is This the Era of Misinformation yet: Combining Social Bots and Fake News to Deceive the Masses." Pp. 1557–61 in *Companion of the The Web Conference 2018 on The Web Conference 2018 - WWW '18*. Lyon, France: ACM Press.
- Weeks, Brian E. 2015. "Emotions, Partisanship, and Misperceptions: How Anger and Anxiety Moderate the Effect of Partisan Bias on Susceptibility to Political Misinformation." *Journal of Communication* 65(4):699–719.
- Williams, Hywel T. P., James R. McMurray, Tim Kurz, and F. Hugo Lambert. 2015. "Network Analysis Reveals Open Forums and Echo Chambers in Social Media Discussions of Climate Change." *Global Environmental Change* 32:126–38.
- Wojcik, Stefan, Solomon Messing, Aaron Smith, Lee Rainie, and Paul Hitlin. 2018. "Bots in the Twittersphere." *Pew Research Center: Internet, Science & Tech*. Retrieved July 16, 2018 (<http://www.pewinternet.org/2018/04/09/bots-in-the-twittersphere/>).
- Zhou, Jack. 2016. "Boomerangs versus Javelins: How Polarization Constrains Communication on Climate Change." *Environmental Politics* 25(5):788–811.