

What happened? The Spread of Fake News Publisher Content During the 2016 U.S. Presidential Election

Ceren Budak

University of Michigan

Ann Arbor, Michigan

cbudak@umich.edu

ABSTRACT

The spread of content produced by fake news publishers was one of the most discussed characteristics of the 2016 U.S. Presidential Election. Yet, little is known about the prevalence and focus of such content, how its prevalence changed over time, and how this prevalence related to important election dynamics. In this paper, we address these questions using tweets that mention the two presidential candidates sampled at the daily level, the news content mentioned in such tweets, and open-ended responses from nationally representative telephone interviews. The results of our analysis highlight various important lessons for news consumers and journalists. We find that (i.) traditional news producers outperformed fake news producers in aggregate, (ii.) the prevalence of content produced by fake news publishers increased over the course of the campaign—particularly among tweets that mentioned Clinton, and (iii.) changes in such prevalence were closely following changes in net Clinton favorability. Turning to content, we (iv.) identify similarities and differences in agenda setting by fake and traditional news media and show that (v.) information individuals most commonly reported to having read, seen or heard about the candidates was more closely aligned with content produced by fake news outlets than traditional news outlets, in particular for information Republican voters retained about Clinton. We also model fake-ness of retained information as a function of demographics characteristics. Implications for platform owners, news consumers, and journalists are discussed.

CCS CONCEPTS

- Social and professional topics → Political speech;
- Information systems → Social networking sites.

KEYWORDS

news media; fake news; topic modeling; multi-level regression

ACM Reference Format:

Ceren Budak. 2019. What happened? The Spread of Fake News Publisher Content During the 2016 U.S. Presidential Election. In *Proceedings of the 2019 World Wide Web Conference (WWW '19), May 13–17, 2019, San Francisco, CA, USA*. ACM, New York, NY, USA, 12 pages. <https://doi.org/10.1145/3308558.3313721>

This paper is published under the Creative Commons Attribution 4.0 International (CC-BY 4.0) license. Authors reserve their rights to disseminate the work on their personal and corporate Web sites with the appropriate attribution.

WWW '19, May 13–17, 2019, San Francisco, CA, USA

© 2019 IW3C2 (International World Wide Web Conference Committee), published under Creative Commons CC-BY 4.0 License.

ACM ISBN 978-1-4503-6674-8/19/05.
<https://doi.org/10.1145/3308558.3313721>

1 INTRODUCTION

News production and consumption have always played an important role in our democracy—with noted consequences for political participation [23], voting [21], and charitable giving [13]. This role has been repeatedly reshaped due to the changes in the media technology—from the introduction of cheap newsprint [41] to the growth of online news production and consumption [24]. Each change has led to concerns about the quality of news and its consequences for our democracy. Most recently, the concerns have centered around the shift towards news production and consumption through social media platforms [3, 24]—with a specific focus on the increasing influence of fake news. While the spread of fake news content on social media platforms is not new, the 2016 U.S. Presidential election has highlighted its extent, raised fears that this new paradigm of news consumption is misinforming voters and corroding our democracy [22, 53, 64]. These concerns have garnered significant attention in both media and policy circles, with some journalists even going as far as claiming that results of the 2016 election were a consequence of the spread of fake news [22, 53].

There are a number of factors that likely contributed to the spread of fake news publisher content, such as increasing use of social media platforms that lack a centralized control mechanism for news consumption [29], an overall decreasing trust in mainstream media [25, 62, 69], profitability of fake news production [34], and recognition of otherwise fringe ideas and news by high profile politicians [6, 47]. Understanding to what extent the aforementioned reasons contributed to the spread of fake news content will certainly take a long time. Here, instead of the “why”, we focus on the “what”; and examine to what degree individuals shared news from fake news producers as opposed to traditional news outlets during the 2016 U.S. presidential election. Using a rich dataset that includes daily random sample of tweets mentioning the two presidential candidates, the content of external urls shared in those tweets, and nationally representative interviews conducted by Gallup at the daily level, we study the overall prevalence and focus of fake news producers, how the prevalence changed over time, and examine how the production and consumption of such news relate to election campaign dynamics.

By analyzing the overall prevalence of content produced by fake and traditional news outlets and shared on Twitter, we show that:

- (1) An overwhelming majority of news content shared on Twitter was produced by *traditional* news outlets.
- (2) The prevalence of fake news increased over time.
- (3) The increase in prevalence was more drastic in tweets about Hillary Clinton and followed significant campaign moments.
- (4) The changes in fake news prevalence followed changes in net Clinton favorability, highlighting a strong correlation

between candidate favorability and fake news production and consumption.

Next, by analyzing the content of the articles, we provide the following contributions:

- (1) We identify the similarities and differences between content produced by fake and traditional news outlets using semi-supervised machine learning methods.
- (2) We determine that the retained information about Clinton (what likely voters have reported to having *read, seen or heard* about her) was more aligned with fake news agenda compared to traditional news agenda.
- (3) We perform statistical modeling of fake-ness of retained information and show that Republican, younger and male interviewees had a higher tendency to report content more closely aligned with fake news publishers. Race and location had no significant effect.

2 RELATED WORK

Social Media and the News. As more news consumption moves to social media platforms [29]—with social media news consumption even providing predictive signals of consumption through all channels [17]—there has been an increasing interest in studying the trustworthiness of content shared through this channel [57, 73].

Recent increase in the volume of misinformation, conspiracy theories, and disinformation and the likely role social media played in this uptick have sparked growing interest in related research [14, 33, 46, 52, 56, 60, 68]. While earlier studies focus on limiting the spread of misinformation as a theoretical optimization problem [14], more recent studies propose data mining and machine learning techniques to identify and limit spread online [56, 60]. Some studies characterize fake news, misinformation, conspiracy theories, and disinformation online [33, 46, 52, 68, 71], while others focus on the types of agents—such as bots—that disseminate such information [1, 7, 20]. Most relatedly, [71] examines rumors that spread on Twitter between 2006 and 2017 and finds that false news spread faster and deeper than true news. Unlike these studies, we focus specifically on the spread of news during 2016 U.S. presidential election.

News Dynamics During 2016 Election. 2016 election has generated increased interest in the study of news dynamics with a particular focus on fake news (e.g. [3, 7, 10, 31, 35, 42, 44, 47, 58]).

Most relatedly, two studies concurrently published with our work investigate the spread of fake news through Twitter [10, 31]. [31] studies the exposure of registered voters on Twitter to fake news content and shows high concentration—only 1% of individuals accounted for 80% of fake news source exposures. They also study engagement with fake news sources as a function of demographics characteristics. [10] studies networks of information flow on Twitter leading up to the 2016 U.S. elections and show that top influencers spreading center and left leaning news influences Clinton supporters. In contrast, the activity of Trump supporters influences the behavior of top fake spreaders.

Other relevant studies include [47] which studies manipulation and disinformation online, provides a survey of various important case studies, and characterizes potential outcomes of such spread. While the main focus is broader than the 2016 election, the authors

emphasize the relationship between the mis- and disinformation ecosystem and election campaign dynamics. [3] studies the effect of fake news during the 2016 election. The authors show that social media was an important but not dominant source of election news, and that fake news stories about Hillary Clinton received more attention. They further estimate the fake news exposure of an average American. [7], aiming to determine the role bots played during the 2016 election, examines tweets shared during the campaign season and finds that about one-fifth of the entire conversation was produced by bots. [64]—focusing on the most commonly shared urls on Facebook—shows that (i.) hyper-partisan news articles attracted more attention than mainstream news articles on Facebook during the 2016 election, (ii.) hyper-partisan and fake news sites had a strong ideological bias against Hillary Clinton. Also relatedly, [42] studies the peer-production of content related to the two presidential candidates on Wikipedia; [58] provides a new method to estimate media bias; and [35] presents a new dataset for examining misinformation on social media.

Our study complements these investigations by bringing together data sources from social media, news media, and Gallup interviews. Unlike related work, we study the focus and the changes in dynamics of fake news prevalence and relate these to important election dynamics. Furthermore, we present the first study that analyzes the relationship between the favorability of the candidates and the production and consumption of news—with a particular focus on content produced by fake news.

3 DEFINING FAKE AND TRADITIONAL NEWS

Identifying the credibility and quality of news content is a challenging task. A recent publication by leading scholars studying misinformation advocates for classification at the outlet level [45]. The authors advocate for “... focusing on the original sources—the publishers—rather than individual stories, because we view the defining element of fake news to be the intent and processes of the publisher. A focus on publishers also allows us to avoid the morass of trying to evaluate the accuracy of every single news story”. This paper follows this guideline and defines fake and traditional news at the producer (web domain) level.

Fake News Publishers. In this study, we rely on the fake news classification provided by [4]. The set of fake news producers published by [4] are gathered from 5 distinct lists: (i) a research project by Grinberg et al. [30] (490 publishers); (ii) Politifact’s article titled “Politifact’s guide to fake news websites and what they peddle” [27] (325 publishers); (iii) the list provided in the three articles published by BuzzFeed on fake news [63, 65, 66] (223 sites); (iv) a research project by Guess et al. [32] (92 publishers); and (v) FactCheck’s article titled “Websites that post fake and satirical stories” [59]. [27, 59, 63, 65, 66] are published by organizations that apply journalistic standards in evaluating the validity of news articles. [30, 32] use a combination of these resources as well as lists published by [12, 75] with some further verification from Snopes.com.

There are a number of fake news producer lists available online (e.g. [27, 30, 32, 59, 63, 65, 66] as described briefly above). We rely on the classification provided by [4] for two main reasons. First, this list provides the widest coverage. Second, the authors provide the most detailed robustness checks of their list compared to other

efforts (e.g. [75]). A key concern here is the potential selection bias in the list of fake news sites gathered. The authors perform two robustness checks. First, they focus on sites that are identified as fake news sites by at least two or three lists instead of one. Second, they consider lists of sites assembled from any four out of the five original lists. They show that the general trends of fake news coverage is consistent across these definitions.

Traditional news. We use Alexa [2] to identify traditional news domains. Alexa provides commercial web traffic analytics using a large panel of Internet users and their web-browsing patterns. This site also provides a rich ontology of the web¹. Using the Alexa ontology, we identify the set of websites listed under the news category. There are in total 5,497 domains listed under this category including national, regional as well as international news sources. We remove the set of web domains listed by [4] and classify the rest of the sites as traditional news sources. Note that the news categorization provided by Alexa has also been used in past work to identify news outlets (e.g. [3, 58]).

We note that future work is needed to measure and improve the accuracy and coverage of the aforementioned lists of traditional and fake news outlets. Furthermore, such lists need to be maintained in a dynamic fashion to account for changes in the news media space.

4 DATA

Social and News Media. The social media dataset used in this study is described in detail in [54]. Data are collected through Sysomos MAP². This service allows gathering of tweets sampled uniformly at random on a daily level. The dataset covers the time period between May 23, 2014 and January 1, 2017. For each day within that time period, we use (i.) 5,000 tweets that were uniformly randomly sampled from all tweets that included the keyword *Trump*, and (ii.) 5,000 tweets that are similarly sampled from all tweets that mention *Clinton* [54]. The resulting dataset includes approximately 4.8 million tweets about Donald Trump and, similarly, approximately 4.8 million tweets about Hillary Clinton. Given that our goal is to analyze news content, we then remove tweets that do not include an external url. The dataset after filtering tweets without urls includes: (i.) 3,072,251 shares of 1,529,251 unique shortened urls by 600,915 unique Twitter users, and (ii.) 4,689,087 shares of 2,025,203 unique shortened urls by 602,613 unique Twitter users for Trump and Clinton.

Given the set of tweets that include a shortened url (of the form: <http://t.co/xxxxxxxxxx>), we next use automated scripts to identify the url redirect and extract the content of the webpage. We remove urls internal to Twitter since such urls correspond to quoted retweets as opposed to external news content. The resulting dataset includes approximately 1.7 million unique urls. For each share (i.e. tweet including a valid url), we record the following data: i.) the shortened url, ii.) the original url, iii.) domain name (e.g. wsj.com), iv.) title of the document, v.) text body of the document, vi.) the date of the tweet, vii.) Twitter id of the user sharing the url, and (viii.) a binary categorization that indicates whether this tweet is in our Clinton or Trump sample. When we limit our data to domains that are included in the fake and traditional news domains lists (detailed

¹Since our data collection, Alexa has stopped sharing the ontology.

²<http://sysomos.com/products/overview/sysomos-map>

below), there are approximately 270K unique articles that were shared in approximately 564K tweets about Trump (in tweets that mention Trump), and approximately 350K unique articles about Clinton that were shared in 1.1M Clinton tweets.

Gallup Data. In order to characterize the election dynamics as perceived by the U.S. population, we also use responses to two open-ended questions fielded as part of the Gallup Daily survey. Respondents were asked, “What have you recently read, seen, or heard about [CANDIDATE]?” for both Donald Trump and Hillary Clinton (asked in random order). To the best of their ability, telephone interviewers recorded individuals’ responses verbatim. Gallup asked these questions to a random sample of 500 adults (ages 18 and higher, living in all 50 U.S. states and the District of Columbia, approximately 60% cellphone and 40% landline) each day between July 11 and November 7, 2016. The margin of sampling error at the weekly level is ± 2 percentage points at the 95% confidence level. For each interview response, the dataset also includes the gender, age, location, and ideological leaning of the respondent. We refer the reader to [9] for additional details about the survey.

5 OVERALL PREVALENCE OF FAKE NEWS PRODUCER CONTENT

A number of investigative efforts suggest wide spread of fake and hyper-partisan news consumption during the 2016 election. For instance, a *Buzzfeed News* study [64] showed that *top* hyper-partisan news articles attracted more attention (measured in terms of total number of shares, reactions, and comments) than *top* mainstream news articles on Facebook by examining top Facebook stories of the election season. Here, we extend this knowledge in three important ways: First, we examine the aggregate popularity (measured in terms of number of shares) of all content shared by traditional and fake news publishers, as opposed to focusing on the top few most commonly shared articles. Second, given that our dataset is balanced over time due to random tweet sampling at the daily level, our results are not dominated by few articles that were in the public eye for a short amount of time. Third, in this study we use data from Twitter, which has important differences in dynamics of information sharing compared to Facebook³. As a result, we extend the understanding of the prevalence of fake news in 2016 Presidential election cycle to a new social media platform.

We find that traditional news articles were shared by 2.8 users on average while this number was 2.88 for articles published by fake news outlets—the small difference observed is not statistically significant ($p \approx 0.08$). Similarly, a traditional news article was shared on 4.7 days on average while this number is 5.5 for articles published by fake news producers ($p < 0.001$ for the difference in means). The distributions are highly skewed (≈ 14 standard deviation for the number of shares and ≈ 33 for the number of days). Median of the distributions are the same (1 for both number of shares and number of days). These measures point to comparable popularity *per article* produced by fake and traditional news publishers⁴. Yet, because the

³Since the writing of this study, two other studies ([10, 31]) also examined Twitter content, albeit through different data collection mechanisms. We find similar qualitative findings in fake news prevalence despite different data collection choices.

⁴Note, however, that our data only include articles with *at least one* share.

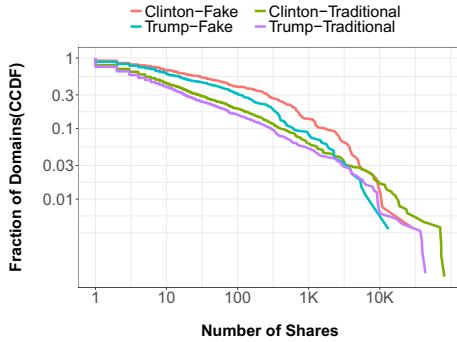


Figure 1: Popularity of Fake and Traditional News Domains for Clinton and Trump Related Tweets.

volume of traditional news production was higher, collectively traditional news accounted for significantly more shares—the number of users sharing traditional news content outnumbered those that share content produced by fake news outlets roughly one to four.

We further investigate how this popularity is concentrated across different publishers and present the results of our findings in Figure 1. The x-axis denotes the popularity of a given domain (separately among tweets mentioning Trump and Clinton) and the y-axis denotes the fraction of domains of the corresponding type (e.g. a fake news domain in Clinton tweets is given by the red curve) with at least that much popularity. Here popularity is defined as the number of tweets that included a url from the corresponding domain. This figure shows that there are a larger fraction of fake news publishers that were able to garner *reasonable* attention compared to traditional news sites. For instance, 68% and 61% of fake news publishers had popularity strictly less than 100 for Trump and Clinton respectively. These numbers are 84% and 81% for traditional news publishers. In contrast, traditional news coverage dominance was mostly driven by a small number of big outlets that drove most of the traffic. For instance, there are no fake news publishers that single-handedly garnered more than 30K shares in our dataset while there are a number of traditional news sites (e.g. New York Times, Fox News). This highlights important differences in production and consumption of fake and traditional news publishers—popularity was more concentrated in the head for traditional news compared to fake news.

6 CHANGES IN FAKE NEWS PRODUCTION AND CONSUMPTION

Here we investigate the relationship between fake news production/consumption and the favorability of the two candidates. Our goal is to determine whether the prevalence of content produced by fake news outlets was correlated with campaign dynamics—following a pattern similar to the relative favorability of candidates and changes to this measure. The favorability of Hillary Clinton on a given day i , denoted by $Fav_{Clinton,i}$, is computed using the Gallup interview dataset as:

$$Fav_{Clinton,i} = \frac{\sum_{r \in R_i} I_{r,Clinton} * wr}{\sum_{r \in R_i} wr} \quad (1)$$

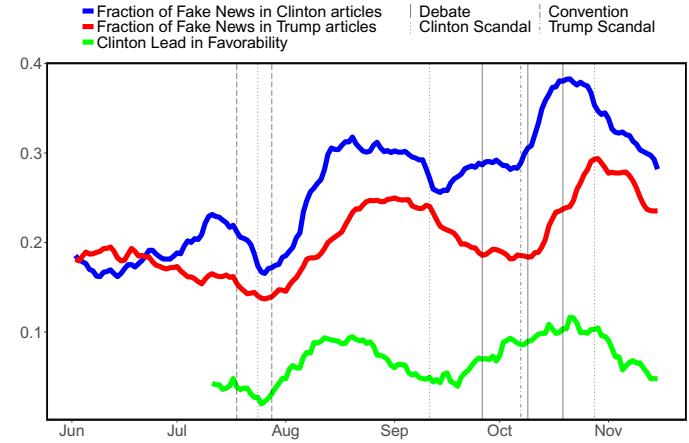


Figure 2: Fraction of article shares that are from fake news producers and Clinton net favorability over time. The error bars give standard errors.

where R_i is the set of respondents on day i , $I_{r,Clinton}$ is an indicator that is equal to 1 if respondent r is favorable towards Hillary Clinton, and w_r is the weight used for post-stratification to correctly estimate sentiment across the United States.⁵ A similar method is used to identify Donald Trump favorability ($Fav_{Trump,i}$) at the daily level. Next, Clinton lead in favorability on day i is computed as $Fav_{\Delta Clinton,i} = Fav_{Clinton,i} - Fav_{Trump,i}$.

We next compute the fake news ratio on a daily level. First, we define 4 sets of tweets for each day i : $T_{Clinton,i}$, $T_{Trump,i}$, $F_{Clinton,i}$, and $F_{Trump,i}$ that correspond to the set of tweets that include (i) link to a traditional publisher and the word “Clinton”, (ii) link to a traditional publisher and “Trump”, (iii) link to a fake news publisher and “Clinton”, and (iv) link to a fake news publisher and “Trump”. Next, fake news content fraction for Clinton and Trump tweets are measured as $f_{Clinton,i} = \frac{|F_{Clinton,i}|}{|F_{Clinton,i}| + |T_{Clinton,i}|}$ and $f_{Trump,i} = \frac{|F_{Trump,i}|}{|F_{Trump,i}| + |T_{Trump,i}|}$ respectively.

The prevalence of fake news content ($f_{Clinton,i}$ and $f_{Trump,i}$) and net Clinton favorability ($Fav_{\Delta Clinton,i}$) over time are provided in Figure 2. The blue curve in Figure 2 denotes the prominence of articles produced by fake news producers in tweets about Clinton. Since $f_{Clinton,i}$ can fluctuate from day-to-day, we smooth out the curve by using and plotting a moving average of 15 days centered at i . The red curve gives a similar measure for the fraction of fake news content in Trump tweets ($f_{Trump,i}$). The green curve denotes the net Clinton favorability over time ($Fav_{\Delta Clinton,i}$). This time series is smoothed out to give the 15-day running average for consistency.

Figure 2 provides a number of important insights. First, tweets about Clinton had a higher fraction of shares from fake news publishers compared to tweets about Trump. Past work shows that ideological bias is commonly exercised by criticizing the other side as opposed to championing for the preferred party [15]. Therefore,

⁵The computation of the weight are described here: <https://www.gallup.com/185462/gallup-daily-work.aspx>.

our finding suggests that fake news domains were commonly publishing negative content about Hillary Clinton. This finding is also in agreement with [64] that claims that fake news content commonly had an anti-Hillary Clinton bias. We further substantiated this conjecture by performing tone analysis of content published by fake and traditional news producers that are shared in Clinton tweets using Lexicoder [74].⁶

Figure 2 also shows an increasing popularity of fake news producers over time—more so for Clinton tweets but also for Trump tweets. While the average fraction of fake news content is about 20% before the conventions, it lies at $\approx 30\%$ after the nominations and reaches $\approx 40\%$ right before the election. For both time series, we observe two peaks. The first upward trending starts after the conventions and the second starts after the second presidential debate. Fake news content popularity is ultimately a function of both production and consumption. This trend can be due to an increase in production of fake news content for ideological reasons to drive public opinion or due to consumption that is a result of changes in public opinion, or a combination of the two.

Finally, and most intriguingly, net Clinton favorability (blue curve) and the prevalence of fake news content (red and green curves) are highly correlated (pearson correlation of 0.8). Intrigued by this correlation, we next perform cross-correlation analysis [70] to investigate whether changes in net Clinton favorability *lead* or *lag* changes in fake news content ratio. CCF or cross correlation function is defined as the set of correlations between the two time series s_{t+l}^1 and s_t^2 for lags $l = 0, \pm 1, \pm 2, \dots$. A negative value for l represents a correlation between the s^1 -series at a time *before* t and the s^2 -series at time t . If, for example, the lag $l = -3$, then the cross correlation value would give the correlation between s_{t+3}^1 and s_t^2 .

The results of the cross-correlation analysis reveals that the highest correlation of 0.85 is observed for $l = 4$, indicating that fake news production/consumption changes *lag* changes in net Clinton favorability. We caution the reader that this analysis is only correlational. While the results suggest a strong relationship between production of fake news content and changes in candidate favorability, this should not be taken as evidence for the *effect* of fake news in the outcome of the 2016 election. If anything, our results suggest that fake news production and consumption responded to changes in the 2016 election campaigns as opposed to affecting them. We further caution that the causal mechanism is unknown. These observed patterns can be due to production (with outlets producing more damaging content for Hillary Clinton to combat improvements in her favorability numbers) or due to consumption, or rather a combination of the two. It is also entirely possible for both candidate favorability patterns and the fake news production and consumption to be responding to a third signal—with the fake news response being delayed longer. Future research is needed to

⁶Lexicoder is a Java-based, multi-platform software for automated content analysis of text. Lexicoder provides two dictionaries—negative and positive words—designed to capture the sentiment of political texts. This method has been used by various studies to examine tone in news coverage [28, 67]. Using the dictionaries, we identify the set of positive and negative words in each article in our dataset and measure the net positivity as #positive - #negative words. Next, we measure the average net positivity of articles shared from fake and traditional news outlets separately. Our findings indeed show that articles published by fake news outlets and shared in tweets mentioning Clinton are on average more negative compared to articles published by traditional news outlets. The average tone is -1.3 and 2.2 respectively for fake and traditional news content.

Table 1: List of Topics, Most Weighted Keywords, and F1 Scores inferred through the methodology described in [11].

Topic	Important Words	F1
Abortion	woman abort life plan_parenthood issu punish femal	0.87
Benghazi	attack benghazi libya committe report secretari secur	0.75
Candidate	medic doctor releas report mental suffer letter pneumonia	0.75
Health		
Climate	climat coal environment industri land administr regul power	0.89
Wall Street	speech wall_street talk ask question issu transcript releas	0.82
Diversity	commun lgbt issu group equal woman discrimin anti marriag	0.78
Economy	trade job china deal compani manufactur econom	0.79
Election	sander berni primari voter percent poll voter cruz	0.77
Email	email depart investig inform server classifi comey secretari	0.84
Immigration	immigr border mexico wall illeg deport mexican build latino	0.85
Middle East	muslim islam israel isi terror terrorist attack unit syria obama	0.76
Religion	christian evangel church faith religi leader pastor religion pope	0.78
Russia	russia russian putin intellig hack offici govern vladimir_putin	0.76
National Security	iran china nuclear polici foreign deal unit world nato secur	0.78
Sex Scandals	woman accus alleg rape husband sexual claim sexual_assault	0.82

identify the causal story—our analysis highlights the importance of carrying out such analysis.

7 FAKE AND TRADITIONAL NEWS TOPICS

Here our goal is to determine the similarities and differences between fake and traditional news coverage leading up to the 2016 election through topic modeling.

The topics covered by traditional and fake news media and shared by Twitter users are inferred through the methodology proposed by [11]. This methodology involves four high level steps: 1) identifying candidate topics and topic-words through vanilla LDA [8] where the number of topics is chosen through topic coherence scores [51], 2) remapping topics and topic words to a smaller subset of coherent topics and associated seed words through expert manual labeling, 3) employing guided LDA [38] using these seed word lists, and 4) collapsing topics that share the same human readable category into a single topic and removing topics without coherence (using top words per topic). This approach produces 19 distinct topics. Each article is then assigned to a single topic according to the maximum probability of its topic distribution. These article-topic pairs are next evaluated manually through Amazon Mechanical Turk. The results reveal that 15 out of the 19 topics to have high accuracy ($F1 \geq 0.75$). Table 1 provides a summary of the 15 coherent topics identified by [11]. We refer the reader to [11] for more details about the methodology.

Having identified the topic of each article in our sample and filtered out articles that do not map to one of the 15 coherent topics, we next examine the similarities and differences between the agenda setting functions of fake and traditional news outlets. In Figure 3(a), we summarize the frequency of each topic separately for fake and traditional news outlets among the set of articles that map to one of the 15 topics. The topics are ordered as a function of the difference in coverage by fake and traditional news outlets. For instance, the top-most topic—*election*—accounts for approximately 48% of traditional news shares while it only accounts for 24% of fake news shares. The difference of 24% signifies to what degree this topic was more central to traditional news production and

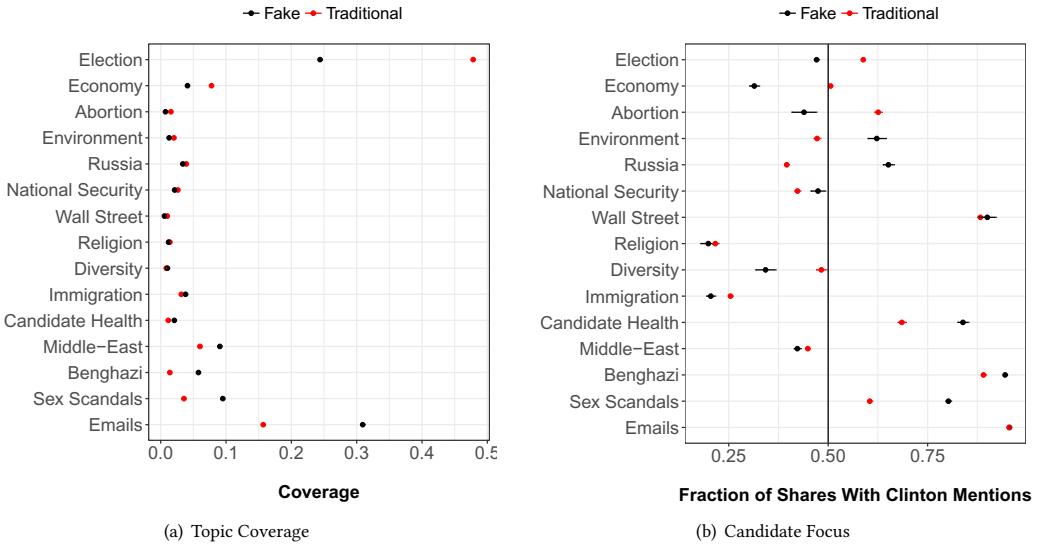


Figure 3: Topical Analysis of Traditional and Fake News Media Shared on Twitter. The topics are ordered from most uniquely associated with traditional news (top) to most uniquely associated with fake news producers (bottom)

consumption compared to fake news production and consumption. In contrast, the bottom-most topic—*emails*—accounts for 31% of fake news shares while it accounts for 16% of traditional news shares. The difference of -15% indicates to what degree emails was a central theme of the fake news coverage and attention.

A number of important patterns emerge here. First, we see that election-related—or in other words horse-race—topic was most significantly overly represented in traditional news coverage. This finding is in line with established research that highlights the significance of the horse-race in news coverage [37] leading up to elections. Here, we show that fake news production diverges from traditional news in this context. Other topics more unique to traditional news outlets include policy focused topics such as the economy, women, and the environment. In comparison, topics that relate to Clinton scandals (emails, sex scandals—commonly about Bill Clinton—, the Benghazi scandal, candidate health issues) were most distinctly covered in articles published by fake news outlets and shared on Twitter. These findings highlight important differences in agenda setting and coverage by these two types of sources.

Despite these clear differences, important similarities emerge when one focuses on the coverage in absolute numbers. For instance, even though emails was a more important topic for fake news outlets compared to traditional news, this topic still made up a significant fraction (16%) of traditional news shares—making emails a more important topic for traditional news than the economy or other policy issues. We further emphasize that no Trump related scandal was common enough to be identified as a coherent topic. Overall, we find that traditional news coverage was extremely light on policy and the focus on a single scandal was disproportionately applied to Clinton. This finding is in line with related work [55] which suggests that traditional news outlets failed to inform the electorate properly leading up to the 2016 U.S. presidential election.

Next, we also investigate whether the coverage of topics were more particularly focused on Clinton or Trump and how that pattern differed across topics. To measure the Clinton focus of a given topic, we compute the fraction of times articles on that topic were shared in Clinton tweets (as opposed to Trump tweets)—separately for fake and traditional news outlets. The results are presented in Figure 3(b). The topics are ordered in the same order as Figure 3(a) for ease of viewing. The x-axis denotes the Clinton focus of the given topic. For instance, the first row shows that approximately 48% of shares of fake news sources on the election were mentioning Clinton, while approximately 54% of shares of traditional news sources on the same topic mentioned Clinton.

Topics with $x > 0.5$ for both fake and traditional news media are consistently Clinton scandals (emails, sex scandals—commonly about Bill Clinton—, Benghazi, candidate health problems, Clinton Wall Street speech scandal). Furthermore, when a topic can relate to scandals of either candidate (e.g. health problems and sex scandals⁷), the coverage of fake news producers focused significantly more on Clinton compared to Trump. The opposite is observed for policy coverage—fake news coverage had a stronger Trump emphasis when covering topics such as abortion and the economy. This suggests that the fake news outlets had a higher tendency to cover Trump policy positions, conditional on covering a given policy topic. These findings further support the claim that fake news production was consistently pro-Trump and anti-Clinton.

8 FAKE NEWS EXPOSURE

Did the information Gallup interview respondents retain about the two candidates better align with traditional or fake news coverage? To answer this question, we compute a measure of fake-ness—the

⁷Both Hillary Clinton and Trump had health issues leading up to the election. Similarly, sex scandals could be referring to Donald Trump or Bill Clinton.

degree to which the given word is unique to fake news coverage—for each unique word uttered by respondents when asked about what they have seen, heard, or read about Clinton (and separately for Trump). Next, we aggregate this information across all interviews to (i.) examine the fake-ness of popular phrases uttered in interviews by individuals with different party alignments, and (ii.) to build a mixed-effects model to characterize fake-ness of interview responses as a function of the attributes of the interviewee (e.g. party, gender, age).

Measuring word fake-ness. We adopt the log-odds-ratio technique with informative Dirichlet priors [48] to compute the fake-ness measure for each word for each candidate. This method has the advantage of de-emphasizing fluctuation of rare words by leveraging the prior information obtained from the background corpus, and has been shown to outperform other methods such as PMI and TF-IDF [40, 48, 50].

Using the log-odds-ratio with informed Dirichlet priors, the usage difference for word w among two corpora is computed as:

$$\delta_w^{(i-j)} = \log \frac{y_w^i + \alpha_w}{n^i + \alpha_0 - y_w^i - \alpha_w} - \log \frac{y_w^j + \alpha_w}{n^j + \alpha_0 - y_w^j - \alpha_w} \quad (2)$$

where n^i (resp. n^j) is the size of corpus i (resp. j), y_w^i (resp. y_w^j) is the word count of w in corpus i (resp. j), α_0 is the size of the background corpus, and α_w is the word count of w in the background corpus. The variance of this measure is:

$$\sigma^2(\delta_w^{(i-j)}) \approx \frac{1}{(y_w^i + \alpha_w)} + \frac{1}{(y_w^j + \alpha_w)} \quad (3)$$

and the Z -score is:

$$Z = \frac{\delta_w^{(i-j)}}{\sqrt{\sigma^2(\delta_w^{(i-j)})}} \quad (4)$$

We measure the fake-ness of a word w (through its Z -score) separately for Clinton and Trump. We do this since the usage of the same phrase when uttered for the two candidates has divergent implications.⁸ Here we describe the process to compute fake-ness scores for each unique lemmatized word $w \in W_{int, Clinton}$ uttered when interviewees were asked what they heard, read or saw about Clinton: We first create three documents that correspond to corpus i, j , and background corpus described above. Corpus i (fake news corpus) includes all $w \in W_{int, Clinton}$ weighted by the number of times w is observed in the body of an article shared in a Clinton tweet in our dataset⁹. Similarly, corpus j (traditional news corpus) includes all $w \in W_{int, Clinton}$ weighted by the number of times w is observed in the body of an article produced by traditional news outlets and shared on Twitter with a Clinton mention. The background corpus is the concatenation of the two documents i and j . Next, we apply equations (2-4) to compute the fake-ness scores of each word when uttered for Clinton. We refer to this measure

⁸Consider for instance, the word “foundation”. This word, when uttered for Clinton, is likely to refer to the Clinton foundation and the relevant scandals that were damaging for her [39]. The same phrase when uttered about Trump is likely to refer to the Trump foundation and the scandals that were damaging to him [26].

⁹If a word w is seen three times in a fake news article that is shared 10 times and twice in another article that is shared 20 times, w will be repeated $3 * 10 + 2 * 20 = 70$ times in corpus i .

as $f(w_{Clinton, fake})$. The same method is applied next to compute the fake-ness of each word for Trump ($f(w_{Trump, fake})$).

This method assigns a value of 68.2 for the word *email* when used for Hillary Clinton—meaning that the word is more frequently observed in fake news articles and is therefore associated with positive *fake-ness*. The word *debat* (“debate” lemmatized), on the other hand is more frequently observed in traditional news articles and has a *fake-ness* measure of -48.5. Table 2 provides an ordered list of the top-20 fake words (words with the highest Z -score) and the list of the top-20 traditional news words (words with the lowest Z -score) for Trump and Clinton. This table highlights the effectiveness of this methodology. For instance, words relating to Clinton scandals (e.g. Huma, Wikileaks, corrupt) have high fake-ness measures (in line with [64]).

Table 2: Words that are most unique to fake vs. traditional articles for Trump and Clinton tweets.

News Outlet	Candidate	Most unique words
Fake	Trump	advertis, video, islam, continu, pleas, twitter, liber, obama, media, illeg, america, leav, establish, advanc, cruz, share, lie, support, christian, alien voter, campaign, republican, compani, casino, percent, photo, former, citi, iowa, advis, univers, hotel, tower, regist, month, busi, presidenti, chariti, build
	Trump	voter, campaign, republican, compani, casino, percent, photo, former, citi, iowa, advis, univers, hotel, tower, regist, month, busi, presidenti, chariti, build
Fake	Clinton	video, foundat, lawsuit, document, twitter, corrupt, huma, wikileak, lie, die, file, sex, court, death, report, email, content, forget, advertis, healthi sander, percent, voter, campaign, democrat, photo, iowa, candid, republican, presidenti, primari, support, hampshire, senat, nomin, debat, poll, race, biden, gop
	Clinton	video, foundat, lawsuit, document, twitter, corrupt, huma, wikileak, lie, die, file, sex, court, death, report, email, content, forget, advertis, healthi sander, percent, voter, campaign, democrat, photo, iowa, candid, republican, presidenti, primari, support, hampshire, senat, nomin, debat, poll, race, biden, gop

Fake-ness of popular interview phrases. Next, we compute the fake-ness of top- k words¹⁰ uttered in Gallup surveys by a set of respondents R for each candidate c as follows:

$$f_{c,k,R}^G = \frac{\sum_{w \in G_{c,k,R}} f(w_{c,fake}) * n_{c,w}}{\sum_{w \in G_{c,k}} n_{c,w}} \quad (5)$$

where $G_{c,k}$ are the set of top- k words uttered by respondents R for candidate c in the Gallup interviews and $n_{c,w}$ is the number of times word w occurred in responses from R for candidate c . In other words, when computing the fake-ness of top- k words uttered by a group of respondents R , we weigh the estimate by how frequently each word is uttered by this group. For instance, *email* occurs 2.5 times more frequently than *debate* when considering all respondents. Therefore, “*email*” will account 2.5 times more heavily compared to “*debate*”. Next, we inspect $f_{c,k,R}^G$ for both candidates separately for (i.) all respondents, (ii.) only Democrats, and (iii.) only Republicans, with k ranging between 1 and 100. As a result, we are able to compare and contrast the alignment between fake news coverage and information retained by different groups. The result of this analysis is summarized in Figure 4. The first panel in Figure 4 characterizes the fake-ness of responses from the three groups

¹⁰Stop words are removed for this analysis.

when asked about what they had read, seen or heard about Hillary Clinton recently. The solid, dotted, and dashed lines correspond to all respondents, Democrats, and Republicans respectively. The second panel presents the same information for the responses about Donald Trump.

A number of observations are worth noting from Figure 4. First, the top-100 words the interviewees uttered about Hillary Clinton consistently resemble fake news coverage more than traditional news coverage (with y -values greater than 0). Second, there is a clear divide between the Republican and Democrat interviewees. Popular Republican leaning interviewee response words are consistently in line with fake news coverage while the popular words used by Democrat leaning interviewees are more in line with traditional news coverage. It is important to note that our findings are purely correlational. While suggestive, one cannot conclude that the fake news producers and content shared by them *resulted in* the topics individuals remember. Indeed, the result we observe can very well be due to fake news domains focusing more on topics that were more memorable (discriminately for the Republican leaning voters). Indeed, *email* provides a good example here. While this word had a positive fake-ness, i.e. more central to fake news coverage, an interviewee could very well have heard about it from traditional news. While we cannot establish causality using this observational data without a more careful causal inference framework, the alignment between the fake news agenda and retained information extracted from Gallup interviews is worthy of further exploration.

The results presented for Donald Trump (second panel) are rather distinct from what we observed for Hillary Clinton. The most frequent words for all interviews were more in line with traditional news coverage compared to the first panel. The fake-ness becomes positive only beyond the most popular topics and compared to Hillary Clinton analysis, the magnitude is small. In addition, the difference between the Republican and Democrat leaning interviewees is much smaller—even hard to detect—and the directionality of the results is largely consistent between these two groups. What explains the relatively low degree to which survey responses align with fake news content about Donald Trump? One conjecture is as follows: Past research claims that fake and hyper-partisan sites were predominantly pro- Donald Trump and anti- Hillary Clinton [64]. Our findings suggest that they showed their ideological leaning not by writing positive articles about Donald Trump but by writing critical articles about Hillary Clinton. Such behavior could lead to a lack of alignment in this topic

Measuring fake-ness of interview responses as a function of demographic characteristics. Figure 4 demonstrates that party affiliation had a strong correlation with the fake-ness of the content interviewees remembered about each candidate. We investigate to what extend party affiliation and other demographic characteristics such as gender, age, race, and geographic area matters through statistical analysis.

To that end, we first measure the fake-ness of an interview response—or rather the fake-ness of the words uttered in it—as a function of the demographic characteristics of the interviewee

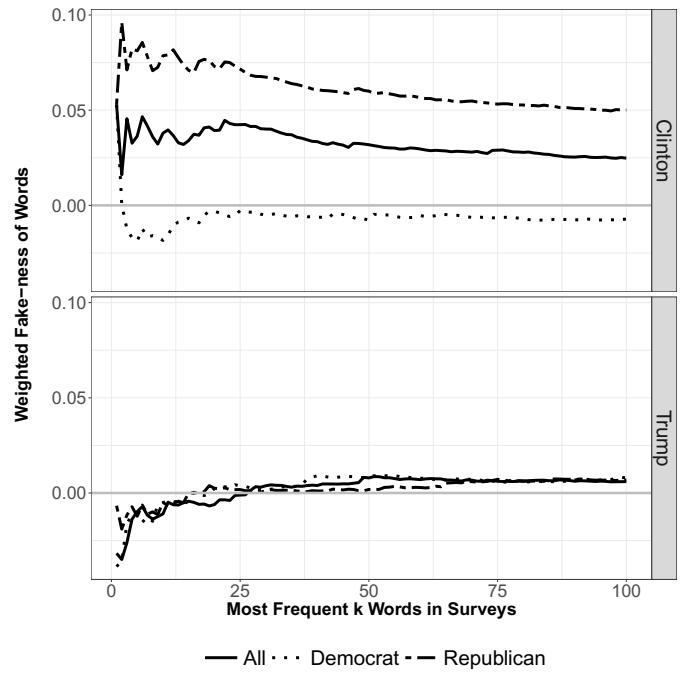


Figure 4: Fake-ness of the top-100 words uttered about Hillary Clinton and Donald Trump in interviews with respondents with party affiliations

and temporal patterns. In particular, we fit the following random-intercept multi-level regression model:

$$f_{i,j} = \beta_0 + \beta_1 party_i + \beta_2 age_i + gender_i + \beta_3 race_i + \beta_4 location_i + \beta_5 months_i + \alpha_i + \epsilon_{i,j} \quad (6)$$

$$\epsilon_{i,j} \sim N(0, \sigma_e^2) \quad (7)$$

$$\alpha_i \sim N(0, \sigma_u^2) \quad (8)$$

where $f_{i,j}$ is the fake-ness of the j -th word uttered by interviewee i , $party_i$ gives the party of interviewee i (Democrat, Independent, or Republican), age_i denotes the age (categorical scale, interview responses coded as 18-29, 30-49, 50-64, or 65+), $gender_i$ denotes the gender of interviewee i , $race_i$ gives the race (Asian, Hispanic, Non-Hispanic White, and Non-Hispanic Black¹¹), $location_i$ denotes the region where the interviewee resides (Southeast, New England, East Central/Great Lakes, Middle Atlantic, West Central, Rocky Mountain, Pacific or Southwest) and $months_i$ denotes the timing of the interview as the time between the interview and the election. This model allows grouping of all words uttered in the same response as opposed to treating them like independent measurements. We fit $f_{i,j}$ separately for response about Clinton and Trump to identify demographic and temporal characteristics that determine how aligned retained information for each candidate was with fake news coverage. Furthermore, since our fake-ness measurement of words that are more commonly observed in the news dataset is less noisy, we perform weighted generalized linear regression where each datapoint ($f_{i,j}$) is weighted by the frequency

¹¹A fifth category of “Other” is filtered out.

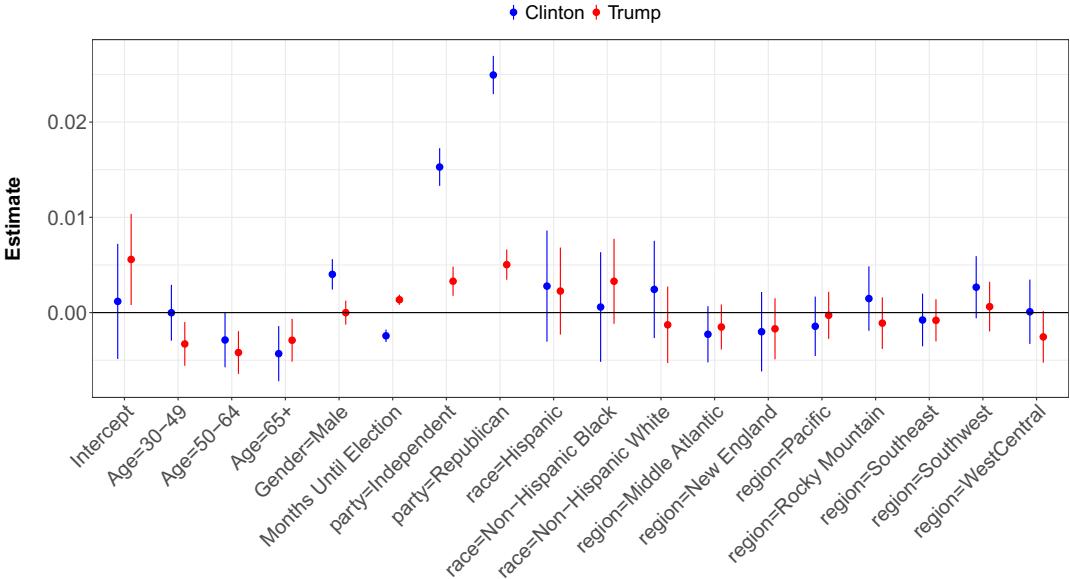


Figure 5: Coefficient estimates from linear mixed-effect models for fake-ness of recalled information about Clinton and Trump shared in interview responses as a function of interviewee characteristics. Base class for gender is female; race is Asian; region is East Central/Great Lakes. Data points are colored to according to the candidate the interviewee was listing retained information about. Each data point gives the coefficient estimate. Error bars mark standard errors of the estimates.

with which we observe the given word in our news sample (across fake and real news combined, in log-scale).

The estimated coefficients are provided in Figure 5. Statistically significant effect sizes are identified for age, gender, party, and timing. The effect directions are largely consistent across the Trump and Clinton responses. However, the effect sizes are significantly larger for Clinton responses. The exception to the consistency between the models is observed for timing patterns. The fake-ness of interview responses increases over time for retained information about Clinton while it decreases for Trump. This difference in temporal patterns can readily be explained by the *increasing* emphasis fake news producers placed on Clinton—observed through various analysis presented earlier in the paper (e.g. Figure 2). Focusing on effects with consistent direction and divergent effect sizes, we see that the Republicans (and to a lesser extent, Independents), younger and male interviewees had a higher tendency to remember information better aligned with fake news publishers. The effects were larger for Clinton. Geographical area and race are not statistically significant. The patterns observed for age are consistent with [3] that claims that fake news commonly spread on social media and older voters were less likely to be exposed to such content. Increased tendency of male respondents to share words with positive fake-ness measures are in line with recent research that showed significant gender effects in views of the two candidates. This highlights the importance of future research to better understand what characteristics are predictive of consuming and retaining information more aligned with fake news publishers.

9 LIMITATIONS

It is worthwhile to note a number of shortcomings of the analysis provided here. First, identifying fake, hyper-partisan, and otherwise fake news at the domain (as opposed to article) level, while supported by related work [45], has certain shortcomings. First, domains classified as fake at times might publish reliable news. Second, even mainstream news sites exhibit ideological bias and at times can disseminate misinformation. Future work that increases the accuracy of fake news classification—both at the article and domain levels—will be crucial for forming a richer understanding. This open problem is something that platforms are increasingly facing. The perfect solution is yet to be identified. This is evidenced in how platforms such as Facebook are alternating between publisher [19] and article [18] level solutions.

Second, content analysis performed in this paper is at the word level. Future studies that employ other content analysis methods and provide a rich qualitative analysis of open-ended Gallup responses can provide further insights about the alignment between fake news coverage and retained information about the candidates. Third, much of our analysis is based on Twitter. While Twitter is one of the popular social media sites used for sharing and consuming the news, it certainly is not the only one [29] and its population is not representative of the U.S. population. Furthermore, we can only examine articles shared at least once in our Twitter sample. Consequently, we cannot disentangle the effects of fake news production and consumption. Future work that examines *all* articles published by these producers can help us determine the effects of consumption and production separately.

Most notably, the analyses presented in this paper are purely correlational. Our findings unearth interesting patterns—such as the relationship between favorability of candidates and fake news consumption or the alignment between the information retained by voters and the content of fake news articles. However, such patterns are only suggestive. Finally, we do not consider the effects of fake news on readers' voting or other relevant attitudes. For instance, our study cannot help determine whether the spread of fake news determined the outcome of the 2016 U.S. Presidential election.

10 DISCUSSION AND IMPLICATIONS

Given the growing pace with which fake news spreads and the societal implications of such spread, it is crucial to examine fake news publisher content and understand its appeal. Motivated by this challenge, here we study the overall prevalence and focus of fake news content leading up to the 2016 U.S. Presidential election and examine how fake news production and consumption related to election campaign dynamics using a rich dataset that includes daily sample of random tweets mentioning the two presidential candidates, content of external urls shared in those tweets and nationally representative interviews conducted by Gallup at the daily level. Consequently, this paper extends our understanding fake news production and consumption and highlights various implications for platform owners and journalists.

First, despite initial fears raised by investigative journalism [64], we observe that fake news producer prevalence was limited compared to traditional news producers. Despite this positive finding, however, the prevalence of fake news producers was far from negligible. This highlights the importance of technology and design mechanisms that filter out unreliable content or provide means for the consumers to easily inspect the quality of a piece of information. Platform owners are already taking notice. For instance, Facebook has [19] announced in early 2018 that they plan to begin prioritizing "trustworthy" media outlets in users' feeds. Yet, this is a serious technological challenge that requires further research. Indeed, a few months later, Facebook announced that they started labeling content at the article (or video/photo) level [18].

Second, we show that the prevalence of fake news content quickly increased over the course of the campaign. This highlights the importance of tracking mechanisms that keep platform owners in the know early enough for the platform to respond to such challenges. Encouragingly, companies are beginning to reach out to academia to help them define important metrics to keep track of for combatting misinformation [36]. Our study highlights the importance of keeping track of prevalence of fake news content in a continuous manner.

Third, we present the first study to examine the relationship between the favorability of presidential candidates and the production and consumption of fake and traditional news. We show that the patterns in prevalence of fake news and favorability of the candidates are consistent. Indeed, we observe that increases in production and consumption of fake news *follows* increase in relative favorability of Clinton. This contribution highlights the importance of studying these two phenomena in conjunction. Yet, one must be cautious here. Our findings do not determine a *causal* relationship between these two trends.

Fourth, through semi-supervised machine learning techniques, we identify similarities and differences in agenda setting and coverage by fake and traditional news outlets. To the best of our knowledge, this is the first study to examines the similarities and differences between the agenda setting functions of fake and traditional news media leading up to the 2016 U.S. Presidential election. Some of our findings here suggest a failing by traditional news outlets, as evidenced by the similarity between the popular topics for fake and traditional news domains. Through the inspection of traditional news content, we observed that these outlets had a rather weak focus on policy and covered scandals of the two candidates unevenly. The behavior of traditional news outlets might be partially explained through the pressure of competing with fake news domains. Regardless, this brings about questions that journalists need to tackle to provide richer and more useful information to voters in the future. Our findings highlight the potential of tools that reveal to journalists similarities and differences between their coverage and the coverage of fake publishers.

Finally, we also showed that information voters retained about Clinton—especially by Republican voters—were more in line with fake news content than traditional news. Our study also highlights the relationship between demographic characteristics and fake-ness of the information people remembered about presidential candidates. The heterogeneity of vulnerability of different demographic groups will be important to keep in mind while building solutions to combat misinformation. This perhaps highlights the importance of diversity of social connections. Individuals with a diverse ego-network can be exposed to more diverse information and being exposed to diverse information can help one question the reliability of fake content. However, challenges still exist given the confirmation bias tendency. In fact, research is inconclusive. Diversity seems to play a positive role under some circumstances [16] while being negative in others [5]—highlighting the importance of context in which people are exposed to diverse ideas.

Where do we go from here? If fake news corrodes our democracy, what can and should be done about it? Social media platforms, while initially skeptical about their role [43], have taken steps to identify fake news content and warn their consumers [49, 61], mostly relying on independent organizations for fact checking. A similar step has been taken by search companies that provide news aggregation tools [72]. These are undoubtedly important steps towards a more robust news ecosystem. We believe further research on identifying and combating such problematic content is vital to the health of our democracy. We need to build the tools of strategy for traditional news organizations to stay relevant and compete with this growing industry and tools of literacy for news consumers to distinguish credible news content from unreliable ones.

ACKNOWLEDGMENTS

This research was partly supported by Michigan Institute for Data Science (MIDAS) at the University of Michigan and the National Science Foundation (Grant IIS-1815875). The author would like to thank Leticia Bode, Jonathan M. Ladd, Frank Newport, Josh Pasek, Lisa O. Singh, Stuart N. Soroka, and Michael W. Traugott for helpful comments and data.

REFERENCES

[1] Norah Abokhodair, Daisy Yoo, and David W. McDonald. 2015. Dissecting a Social Botnet: Growth, Content and Influence in Twitter. In *Proceedings of the 18th ACM Conference on Computer Supported Cooperative Work & Social Computing (CSCW '15)*. ACM, New York, NY, USA, 839–851. <https://doi.org/10.1145/2675133.2675208>

[2] Inc. Alexa Internet. 1996. Alexa—About Us. <http://www.alexa.com/about>.

[3] Hunt Allcott and Matthew Gentzkow. 2017. *Social media and fake news in the 2016 election*. Technical Report. National Bureau of Economic Research.

[4] Hunt Allcott, Matthew Gentzkow, and Chuan Yu. 2019. *Trends in the diffusion of misinformation on social media*. Technical Report. National Bureau of Economic Research.

[5] Christopher A Bail, Lisa Argyle, Taylor Brown, John Bumpuss, Haohan Chen, M.B. F 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. <https://doi.org/10.31235/osf.io/4yux>

[6] Michael Barbaro. 2016. Donald Trump Clung to ‘Birther’ Lie for Years, and Still Isn’t Apologetic. <https://www.nytimes.com/2016/09/17/us/politics/donald-trump-obama-birther.html>.

[7] Alessandro Bessi and Emilio Ferrara. 2016. Social bots distort the 2016 U.S. Presidential election online discussion. *First Monday* 21, 11 (2016). <http://firstmonday.org/ojs/index.php/fm/article/view/7090>

[8] David M Blei, Andrew Y Ng, and Michael I Jordan. 2003. Latent dirichlet allocation. *Journal of machine Learning research* 3, Jan (2003), 993–1022.

[9] Leticia Bode, Ceren Budak, Jonathan M. Ladd, Frank Newport, Josh Pasek, Lisa O. Singh, Stuart N. Soroka, and Michael W. Traugott. Forthcoming. *Words That Matter: How the News and Social Media Shaped the 2016 Presidential Campaign*. Brookings Institution Press, Washington, D.C.

[10] Alexandre Bovet and Hernán A Makse. 2019. Influence of fake news in Twitter during the 2016 US presidential election. *Nature communications* 10, 1 (2019), 7.

[11] Lia Bozarth, Aparajita Saraf, and Ceren Budak. 2019. Higher Ground? How Groundtruth Labeling Impacts Our Understanding of the Spread of Fake News During the 2016 Election. <https://www.ssrn.com/abstract=3340173>. SSRN (2019).

[12] Ed Brayton. 2016. Please Stop Sharing Links to These Sites.

[13] Phil Brown and Jessica Minty. 2006. Media coverage & charitable giving after the 2004 tsunami. (2006).

[14] Ceren Budak, Divyakant Agrawal, and Amr El Abbadi. 2011. Limiting the Spread of Misinformation in Social Networks. In *Proceedings of the 20th International Conference on World Wide Web (WWW '11)*. ACM, New York, NY, USA, 665–674. <https://doi.org/10.1145/1963405.1963499>

[15] Ceren Budak, Sharad Goel, and Justin M Rao. 2016. Fair and balanced? quantifying media bias through crowdsourced content analysis. *Public Opinion Quarterly* 80, S1 (2016), 250–271.

[16] Ceren Budak and Duncan J Watts. 2015. Dissecting the spirit of Gezi: Influence vs. selection in the Occupy Gezi movement. *Sociological Science* 2 (2015), 370–397.

[17] Carlos Castillo, Mohammed El-Haddad, Jürgen Pfeffer, and Matt Stempeck. 2014. Characterizing the Life Cycle of Online News Stories Using Social Media Reactions. In *Proceedings of the 17th ACM Conference on Computer Supported Cooperative Work & Social Computing (CSCW '14)*. ACM, New York, NY, USA, 211–223. <https://doi.org/10.1145/2531602.2531623>

[18] Josh Constine. 2018. Facebook starts fact checking photos/videos, blocks millions of fake accounts per day. <https://techcrunch.com/2018/03/29/facebook-fact-check-photos/>.

[19] Rob Crilly. 2018. Facebook to start ‘trust ratings’ for media outlets as it fights back against fake news. <http://www.telegraph.co.uk/news/2018/01/19/facebook-start-trust-ratings-media-outlets-fights-back-against/>.

[20] Clayton Allen Davis, Onur Varol, Emilio Ferrara, Alessandro Flammini, and Filippo Menczer. 2016. BotOrNot: A System to Evaluate Social Bots. In *Proceedings of the 25th International Conference Companion on World Wide Web (WWW '16 Companion)*. International World Wide Web Conferences Steering Committee, Republic and Canton of Geneva, Switzerland, 273–274. <https://doi.org/10.1145/2872518.2889302>

[21] Stefano DellaVigna and Ethan Kaplan. 2007. The Fox News effect: media bias and voting. *The Quarterly Journal of Economics* 122, 3 (2007), 1187–1234.

[22] Caitlin Dewey. 2016. Facebook Fake-News Writer: ‘I Think Donald Trump is in the White House Because of Me’.

[23] Susanna Dilliplane. 2011. All the News You Want to Hear: The Impact of Partisan News Exposure on Political Participation. *Public Opinion Quarterly* 75, 2 (2011), 287–316. <https://doi.org/10.1093/poq/nfr006>

[24] Seth Flaxman, Sharad Goel, and Justin M Rao. 2016. Filter bubbles, echo chambers, and online news consumption. *Public Opinion Quarterly* 80, S1 (2016), 298–320.

[25] Jack Fuller. 2010. *What is happening to news: The information explosion and the crisis in journalism*. University of Chicago Press, Chicago.

[26] Lauren Gambino. 2016. Trump Foundation: the allegations against Donald’s charity, explained. <https://www.theguardian.com/us-news/2016/oct/05/donald-trump-foundation-allegations-charity-new-york>.

[27] Joshua Gillin. 2017. PolitiFact’s Guide to Fake News Websites and What They Peddle. <http://bit.ly/2o8kj3b>.

[28] Sandra González-Bailón and Georgios Paltoglou. 2015. Signals of public opinion in online communication: A comparison of methods and data sources. *The ANNALS of the American Academy of Political and Social Science* 659, 1 (2015), 95–107.

[29] Jeffrey Gottfried and Elisa Shearer. 2016. News use across social media platforms 2016. <http://www.journalism.org/2016/05/26/news-use-across-social-media-platforms-2016>.

[30] N Grinberg, K Joseph, L Friedland, B Swire-Thompson, and D Lazer. 2018. *Fake news on Twitter during the 2016 US presidential election*. Technical Report. Working Paper. Available from the authors.

[31] Nir Grinberg, Kenneth Joseph, Lisa Friedland, Briony Swire-Thompson, and David Lazer. 2019. Fake news on Twitter during the 2016 US presidential election. *Science* 363, 6425 (2019), 374–378.

[32] Andrew Guess, Brendan Nyhan, and Jason Reifler. 2018. Selective Exposure to Misinformation: Evidence from the consumption of fake news during the 2016 US presidential campaign. , 14 pages.

[33] Aditi Gupta, Hemank Lamba, Ponnurangam Kumaraguru, and Anupam Joshi. 2013. Faking Sandy: Characterizing and Identifying Fake Images on Twitter During Hurricane Sandy. In *Proceedings of the 22Nd International Conference on World Wide Web (WWW '13 Companion)*. ACM, New York, NY, USA, 729–736. <https://doi.org/10.1145/2487788.2488033>

[34] ANDREW HIGGINS, MIKE McINTIRE, and GABRIEL J.X. DANCE. 2016. Inside a Fake News Sausage Factory: ?This Is All About Income? <https://www.nytimes.com/2016/11/25/world/europe/fake-news-donald-trump-hillary-clinton-georgia.html>.

[35] Pik-Mai Hui, Chengcheng Shao, Alessandro Flammini, Filippo Menczer, and Giovanni Luca Ciampaglia. 2018. The Hoaxy Misinformation and Fact-Checking Diffusion Network. In *Proceedings of the Twelfth International Conference on Web and Social Media, ICWSM 2018, June 25–28, 2018*. AAAI Press, Stanford, California, USA, 528–530. <https://aaai.org/ocs/index.php/ICWSM/ICWSM18/paper/view/17851>

[36] Twitter Inc. 2018. Twitter Health Metrics Proposal Submission. https://blog.twitter.com/official/en_us/topics/company/2018/twitter-health-metrics-proposal-submission.html.

[37] Shanto Iyengar, Helmut Norpoth, and Kyu S Hahn. 2004. Consumer demand for election news: The horserace sells. *Journal of Politics* 66, 1 (2004), 157–175.

[38] Jagadeesh Jagarlamudi, Hal Daumé, III, and Raghavendra Udupa. 2012. Incorporating Lexical Priors into Topic Models. In *Proceedings of the 13th Conference of the European Chapter of the Association for Computational Linguistics (EACL '12)*. Association for Computational Linguistics, Stroudsburg, PA, USA, 204–213. <http://dl.acm.org/citation.cfm?id=2380816.2380844>

[39] Elise Jordan. 2016. Hillary Clinton’s Reaction to Her Foundation Scandal Is Disastrous. <http://time.com/4466911/hillary-clinton-email-excuses/>.

[40] Dan Jurafsky, Victor Chahuneau, Bryan R Routledge, and Noah A Smith. 2014. Narrative framing of consumer sentiment in online restaurant reviews. *First Monday* 19, 4 (2014).

[41] Richard L Kaplan. 2002. *Politics and the American press: The rise of objectivity, 1865–1920*. Cambridge University Press, Cambridge.

[42] Brian C. Keegan. 2018. The Dynamics of Peer-Produced Political Information During the 2016 U.S. Presidential Campaign. In *Proceedings of the Twelfth International Conference on Web and Social Media, ICWSM 2018, June 25–28, 2018*. AAAI Press, Stanford, California, USA, 632–635. <https://aaai.org/ocs/index.php/ICWSM/ICWSM18/paper/view/17800>

[43] Kira Kokalitcheva. 2016. Mark Zuckerberg Says Fake News on Facebook Affecting the Election Is a ‘Crazy Idea’. <http://fortune.com/2016/11/11/facebook-election-fake-news-mark-zuckerberg/>.

[44] Haewoon Kwak, Jisun An, Joni Salminen, Soon-Gyo Jung, and Bernard J. Jansen. 2018. What We Read, What We Search: Media Attention and Public Attention Among 193 Countries. In *Proceedings of the 2018 World Wide Web Conference (WWW '18)*. International World Wide Web Conferences Steering Committee, Republic and Canton of Geneva, Switzerland, 893–902. <https://doi.org/10.1145/3178876.3186137>

[45] David M. J. Lazer, 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 (2018), 1094–1096. <https://doi.org/10.1126/science.aoa2998> arXiv:<http://science.sciencemag.org/content/359/6380/1094.full.pdf>

[46] Jim Maddock, Kate Starbird, Haneen J. Al-Hassani, Daniel E. Sandoval, Mania Ostrand, and Robert M. Mason. 2015. Characterizing Online Rumoring Behavior Using Multi-Dimensional Signatures. In *Proceedings of the 18th ACM Conference on Computer Supported Cooperative Work & Social Computing (CSCW '15)*. ACM, New York, NY, USA, 228–241. <https://doi.org/10.1145/2675133.2675280>

[47] Alice Marwick and Rebecca Lewis. 2017. Media Manipulation and Disinformation Online.

[48] Burt L Monroe, Michael P Colaresi, and Kevin M Quinn. 2008. Fightin’words: Lexical feature selection and evaluation for identifying the content of political

conflict. *Political Analysis* 16, 4 (2008), 372–403.

[49] Adam Mosseri. 2016. News feed fyi: Addressing hoaxes and fake news. <http://newsroom.fb.com/news/2016/12/news-feed-fyi-addressing-hoaxes-and-fake-news>.

[50] Supun Nakandala, Giovanni Luca Ciampaglia, Norman Makoto Su, and Yong-Yeol Ahn. 2016. Gendered Conversation in a Social Game-Streaming Platform. *CoRR* abs/1611.06459 (2016). arXiv:1611.06459 <http://arxiv.org/abs/1611.06459>

[51] David Newman, Jey Han Lau, Karl Grieser, and Timothy Baldwin. 2010. Automatic Evaluation of Topic Coherence. In *Human Language Technologies: The 2010 Annual Conference of the North American Chapter of the Association for Computational Linguistics (HLT '10)*. Association for Computational Linguistics, Stroudsburg, PA, USA, 100–108. <http://dl.acm.org/citation.cfm?id=1857999.1858011>

[52] A. Conrad Nied, Leo Stewart, Emma Spiro, and Kate Starbird. 2017. Alternative Narratives of Crisis Events: Communities and Social Botnets Engaged on Social Media. In *Companion of the 2017 ACM Conference on Computer Supported Cooperative Work and Social Computing (CSCW '17 Companion)*. ACM, New York, NY, USA, 263–266. <https://doi.org/10.1145/3022198.3026307>

[53] Hannah Jane Parkinson. 2016. Click and elect: how fake news helped Donald Trump win a real election.

[54] Josh Pasek, Lisa O. Singh, Yifang Wei, Stuart N. Soroka, Jonathan M. Ladd, Michael W. Traugott, Ceren Budak, Leficia Bode, and Frank Newport. 2018. Attention to Campaign Events: Do Twitter and Self-Report Metrics Tell the Same Story? *Forthcoming in BIGSURV18: Big Data Meets Survey Science Edited Volume* (2018).

[55] Thomas E Patterson. 2016. *News Coverage of the 2016 General Election: How the Press Failed the Voters*. Technical Report. Harvard University, John F. Kennedy School of Government.

[56] Paul Resnick, Samuel Carton, Souneil Park, Yuncheng Shen, and Nicole Zeffer. 2014. Rumorlens: A system for analyzing the impact of rumors and corrections in social media.

[57] Paul Resnick, R. Kelly Garrett, Travis Kriplean, Sean A. Munson, and Natalie Jo- mini Stroud. 2013. Bursting Your (Filter) Bubble: Strategies for Promoting Di- verse Exposure. In *Proceedings of the 2013 Conference on Computer Supported Cooperative Work Companion (CSCW '13)*. ACM, New York, NY, USA, 95–100. <https://doi.org/10.1145/2441955.2441981>

[58] Filipe Nunes Ribeiro, Lucas Henrique, Fabricio Benevenuto, Abhijnan Chakraborty, Jubi Kulshrestha, Mahmoudreza Babaei, and Krishna P. Gummadi. 2018. Media Bias Monitor: Quantifying Biases of Social Media News Outlets at Large-Scale. In *Proceedings of the Twelfth International Conference on Web and Social Media, ICWSM 2018, June 25–28, 2018*. AAAI Press, Stanford, California, USA, 290–299. <https://aaai.org/ocs/index.php/ICWSM/ICWSM18/paper/view/17878>

[59] S Schaedel. 2018. Websites that post fake and satirical stories. FactCheck. <https://www.factcheck.org/2017/07/websites-post-fake-satirical-stories>.

[60] Chengcheng Shao, Giovanni Luca Ciampaglia, Alessandro Flammini, and Filippo Menczer. 2016. Hoaxy: A Platform for Tracking Online Misinformation. In *Proceedings of the 25th International Conference Companion on World Wide Web (WWW '16 Companion)*. International World Wide Web Conferences Steering Committee, Republic and Canton of Geneva, Switzerland, 745–750. <https://doi.org/10.1145/2872518.2890098>

[61] Mallory Shelbourne. 2017. Facebook rolls out feature to combat fake news. <http://thehill.com/policy/technology/322427-facebook-begins-to-introduce-new-feature-to-combat-fake-news>.

[62] Ignacio Siles and Pablo J Boczkowski. 2012. Making sense of the newspaper crisis: A critical assessment of existing research and an agenda for future work. *New Media & Society* 14, 8 (2012), 1375–1394.

[63] Craig Silverman. 2016. Here are 50 of the biggest fake news hits on Facebook from 2016. <https://www.buzzfeed.com/craigsilverman/top-fake-news-of-2016>.

[64] C Silverman. 2016. This analysis shows how fake election news stories outper- formed real news on Facebook. <https://www.buzzfeed.com/craigsilverman/viral-fake-election-news-outperformed-real-news-on-facebook>.

[65] Craig Silverman, Jane Lytvynenko, and Scott Pham. 2017. These are 50 of the biggest fake news hits on Facebook in 2017. <https://www.buzzfeednews.com/article/craigsilverman/these-are-50-of-the-biggest-fake-news-hits-on-facebook-in>.

[66] Craig Silverman and Jeremy Singer-Vine. 2016. Most Americans who see fake news believe it, new survey says. <https://www.buzzfeed.com/craigsilverman/fake-newssurvey>.

[67] Stuart N Soroka. 2012. The gatekeeping function: Distributions of information in media and the real world. *The Journal of Politics* 74, 2 (2012), 514–528.

[68] Kate Starbird. 2017. Examining the Alternative Media Ecosystem Through the Production of Alternative Narratives of Mass Shooting Events on Twitter. In *ICWSM*. AAAI Press, Montréal, Québec, Canada, 230–239.

[69] Art Swift. 2016. Americans' trust in mass media sinks to new low.

[70] William N Venables and Brian D Ripley. 2013. *Modern applied statistics with S-PLUS*. Springer Science & Business Media, Berlin, Germany.

[71] Soroush Vosoughi, Deb Roy, and Sinan Aral. 2018. The spread of true and false news online. *Science* 359, 6380 (2018), 1146–1151. <https://doi.org/10.1126/science.aap9559> arXiv:<http://science.sciencemag.org/content/359/6380/1146.full.pdf>

[72] Nick Wingfield, Mike Isaac, and Katie Benner. 2016. Google and Facebook take aim at fake news sites. *The New York Times, November* 14 (2016), 2016.

[73] Donghee Yvette Wohin, Casey Fiesler, Libby Hemphill, Munmun De Choudhury, and J. Nathan Matias. 2017. How to Handle Online Risks?: Discussing Content Curation and Moderation in Social Media. In *Proceedings of the 2017 CHI Conference Extended Abstracts on Human Factors in Computing Systems (CHI EA '17)*. ACM, New York, NY, USA, 1271–1276. <https://doi.org/10.1145/3027063.3051141>

[74] Lori Young and Stuart Soroka. 2012. Lexicoder sentiment dictionary.

[75] Melissa Zimdars. 2016. False, Misleading, Clickbait-y, and Satiri- cal "News" Sources. https://docs.google.com/document/u/1/d/10eA5-mCZLSS4MQY5QGb5ewC3VAL6pLkT53V_81ZytM/mobilebasic.