

Trends in the Diffusion of Misinformation on Social Media

Hunt Allcott, *New York University, Microsoft Research, and NBER**

Matthew Gentzkow, *Stanford University and NBER*

Chuan Yu, *Stanford University*

October 2018

Abstract

In recent years, there has been widespread concern that misinformation on social media is damaging societies and democratic institutions. In response, social media platforms have announced actions to limit the spread of false content. We measure trends in the diffusion of content from 570 fake news websites and 10,240 fake news stories on Facebook and Twitter between January 2015 and July 2018. User interactions with false content rose steadily on both Facebook and Twitter through the end of 2016. Since then, however, interactions with false content have fallen sharply on Facebook while continuing to rise on Twitter, with the ratio of Facebook engagements to Twitter shares decreasing by 60 percent. In comparison, interactions with other news, business, or culture sites have followed similar trends on both platforms. Our results suggest that Facebook's efforts to limit the diffusion of misinformation after the 2016 election may have had a meaningful impact.

*E-mail: hunt.allcott@nyu.edu, gentzkow@stanford.edu, chuanyu@stanford.edu. We thank the Stanford Institute for Economic Policy Research (SIEPR), the Stanford Cyber Initiative, the Toulouse Network for Information Technology, the Knight Foundation, and the Alfred P. Sloan Foundation for generous financial support. We thank David Lazer, Brendan Nyhan, David Rand, David Rothschild, Jesse Shapiro, Nils Wernerfelt, and seminar participants at Facebook for helpful comments and suggestions. We also thank our dedicated research assistants for their contributions to this project.

1 Introduction

Although the political process has a long history of misinformation and popular misperceptions, misinformation on social media has caused widespread alarm in recent years (Flynn et al. 2017; Lazer et al. 2018). A substantial number of U.S. adults were exposed to false stories prior to the 2016 election, and post-election surveys suggest that many people who read these stories believed them to be true (Allcott and Gentzkow 2017; Guess et al. 2018). Many argue that false stories played a major role in the 2016 election (for example, Parkinson 2016; Gunther et al. 2018), and in the ongoing political divisions and crises that have followed it (for example, Spohr 2017; Azzimonti and Fernandes 2018). In response, Facebook and other social media companies have made a range of algorithmic and policy changes to limit the spread of false content. In the appendix, we list twelve announcements by Facebook and five by Twitter aimed at reducing the circulation of misinformation on their platforms since the 2016 election.

Evidence on whether these efforts have been effective—or how the scale of the misinformation problem is evolving more broadly—remains limited.¹ A recent study argues that false stories remain a problem on Facebook even after changes to the platform’s news feed algorithm in early 2018 (Newswhip 2018). Many articles that have been rated as false by major fact-checking organizations have not been flagged in Facebook’s system, and two major fake news sites have seen little or no decline in Facebook engagements since early 2016 (Funke 2018). Facebook’s now-discontinued strategy of flagging inaccurate stories as “Disputed” has been shown to modestly lower the perceived accuracy of flagged headlines (Blair et al. 2017), though some research suggests that the presence of warnings can cause untagged false stories to be seen as more accurate (Pennycook and Rand 2017). Media commentators have argued that efforts to fight misinformation through fact-checking are “not working” (Levin 2017) and that misinformation overall is “becoming unstoppable” (Ghosh and Scott 2018).

In this paper, we present new evidence on the volume of misinformation circulated on social media from January 2015 to July 2018. We assemble a list of 570 sites identified as sources of false stories in a set of five previous studies and online lists. We refer to these collectively as *fake news sites*. We measure the volume of Facebook engagements and Twitter shares for all stories on these sites by month. As points of comparison, we also measure the same outcomes for stories on (i) a set of major news sites; (ii) a set of small news sites not identified as producing misinformation; and (iii) a set of sites covering business and culture topics.

The results show that interactions with the fake news sites in our database rose steadily on both Facebook and Twitter from early 2015 to the months just after the 2016 election. Interactions then declined by more than half on Facebook, while they continued to rise on Twitter. The ratio of

¹Lazer et al. (2018) write, “There is little research focused on fake news and no comprehensive data-collection system to provide a dynamic understanding of how pervasive systems of fake news provision are evolving . . . researchers need to conduct a rigorous, ongoing audit of how the major platforms filter information.”

Facebook engagements to Twitter shares was roughly steady at around 40:1 from the beginning of our period to late 2016, then fell to approximately 15:1 by the end of our sample period. In contrast, interactions with major news sites, small news sites, and business and culture sites have all remained relatively stable over time, and have followed similar trends on Facebook and Twitter both before and after the 2016 election. While this evidence is far from definitive and is subject to the important caveats discussed below, we see it as consistent with the view that the overall magnitude of the misinformation problem may have declined, at least temporarily, and that efforts by Facebook following the 2016 election to limit the diffusion of misinformation may have had a meaningful impact.

Our results also reveal that the absolute level of interaction with misinformation remains high, and that Facebook continues to play a particularly important role in its diffusion. In the period around the election, fake news sites received almost as many Facebook engagements as the 38 major news sites in our sample. Even after the post-election decline, Facebook engagements with fake news sites still average roughly 70 million per month.

This research demonstrates how novel data on social media usage can be used to understand important questions in political science around media exposure and social media platforms' content moderation practices. Parallel work released soon after our working paper finds broadly similar results (Resnick et al. 2018).

2 Data

We compile a list of sites producing false stories by combining five previous lists: (i) an academic paper by Grinberg et al. (2018, 490 sites); (ii) PolitiFact's article titled "PolitiFact's guide to fake news websites and what they peddle" (Gillin 2017, 325 sites); (iii) three articles by BuzzFeed on fake news (Silverman 2016; Silverman et al. 2017a; Silverman et al. 2017b; 223 sites); (iv) an academic paper by Guess et al. (2018, 92 sites); and (v) FactCheck's article titled "Websites that post fake and satirical stories" (Schaedel 2017, 61 sites). The two lists from academic papers originally derive from subsets of the other three, plus Snopes.com, another independent fact-checking site, and lists assembled by blogger Brayton (2016) and media studies scholar Zimdars (2016). The union of these five lists is our set of fake news sites.

PolitiFact and FactCheck work directly with Facebook to evaluate the veracity of stories flagged by Facebook users as potentially false. Thus, these lists comprise fake news sites that Facebook is likely to be aware are fake. As a result, our results may be weighted toward diffusion of *misinformation that Facebook is aware of*, and may not fully capture trends in *misinformation that Facebook is not aware of*. It is difficult to assess how large this latter group might be. Our list almost certainly includes the most important providers of false stories, as Facebook users can flag any and all questionable articles for review. On the other hand, the list likely excludes a large tail

of web domains that are small and/or active for only a short period.

Combining these five lists yields a total of 673 unique sites. We report in the appendix the names and original lists of 50 largest sites in terms of total Facebook engagements plus Twitter shares during the sample period. In our robustness checks, we consider alternative rules for selecting the set of sites.

Our sets of comparison sites are defined based on category-level web traffic rankings from Alexa (www.alexa.com). Alexa measures web traffic using its global traffic panel, a sample of millions of Internet users who have installed browser extensions allowing their browsing data to be recorded, plus data from websites that use Alexa to measure their traffic. It then ranks sites based on a combined measure of unique visitors and pageviews. We define *major news sites* to be the top 100 sites in Alexa's News category. We define *small news sites* to be the sites ranked 401-500 in the News category. We define *business and culture sites* to be the top 50 sites in each of the Arts, Business, Health, Recreation, and Sports categories. For each of these groups, we omit from our sample government websites, databases, sites that do not mainly produce news or similar content, international sites whose audiences are primarily outside the U.S., and sites that are included in our list of fake news sites. Our final sample includes 38 major news sites, 78 small news sites, and 54 business and culture sites.

We gather monthly Facebook engagements and Twitter shares of all articles published on these sites from January 2015 to July 2018 from BuzzSumo (www.buzzsumo.com). BuzzSumo is a commercial content database that tracks the volume of user interactions with internet content on Facebook, Twitter, and other social media platforms, using data available from the platforms' application programming interfaces (APIs). We use BuzzSumo's data on total Facebook engagements and total Twitter shares by originating website and month. Facebook engagements are defined as the sum of shares, comments, and reactions such as "likes." (Ideally we would measure exposure to fake articles using data on views, but such data are not publicly available.) We have data for 570 out of 673 fake news sites in our list and all sites in the comparison groups. We sum the monthly Facebook engagements and Twitter shares of articles from all sites in each category and then average by quarter.

In practice, the 570 "fake news sites" on our list carry some combination of true news and clickbait in addition to misleading and false content. To more precisely focus attention on the latter, we gather a list of specific URLs spreading misinformation. We scrape all claims on the fact-checking site Snopes.com that are classified as "false" or "mostly false." In late 2015, Snopes began to provide permanent URLs for the sources of these false claims through a web archiving site, archive.is. We collect all these URLs for articles published in 2016 or later, yielding an intermediate sample of 1,535 article URLs. We then extract keywords from the titles of these articles, and we capture all articles in the BuzzSumo database published in 2016 or later that contain these keywords and have at least 100 Facebook engagements or 10 Twitter shares, manually

screening out those that are not in fact spreading the false claims. This yields a final sample of 10,240 false stories URLs.

3 Results

Figure 1 shows trends in the number of Facebook engagements and Twitter shares of stories from each category of site. Interactions for major news sites, small news sites, and business and culture sites have remained relatively stable during the past two years, and follow similar trends on Facebook and Twitter. Both platforms show a modest upward trend for major news and small news sites, and a modest downward trend for business and culture sites. In contrast, interactions with fake news have changed more dramatically over time, and these changes are very different on the two platforms. Fake news interactions increased steadily on both platforms from the beginning of 2015 up to the 2016 election. Following the election, however, Facebook engagements fell sharply (declining by more than 50 percent), while Twitter shares continued to increase.

Figure 2 shows our main result: trends in the ratios of Facebook engagements to Twitter shares. The ratios have been relatively stable for major news, small news, and business and culture sites. For fake news sites, however, the ratio has declined sharply, from around 45:1 during the election to around 15:1 two years later.

While these results suggest that the circulation of misinformation on Facebook has declined, it is important to emphasize that the absolute quantity of interactions with misinformation on both platforms remains large, and that Facebook in particular has played an outsized role in its diffusion. Figure 1 shows that Facebook engagements fell from a peak of roughly 200 million per month at the end of 2016 to roughly 70 million per month at the end of our sample period. As a point of comparison, the 38 major news sites in the top left panel—including the New York Times, Wall Street Journal, CNN, Fox News, etc.—typically garner about 200-250 million Facebook engagements per month. On Twitter, shares of false content have been in the 4-6 million per month range since the end of 2016, compared to roughly 20 million per month for the major news sites.

Figure 3 presents results for our list of false stories URLs. Since the number of URLs we capture starts close to zero in 2016 and grows from month to month, there is a steeper increase in Facebook and Twitter interactions with these URLs than that in the site-level analysis. Similar to the site-level analysis, the ratio of Facebook engagements to Twitter shares has declined by half or more after the 2016 election.

3.1 Interpretation and Robustness Checks

Our evidence is subject to many important caveats and must be interpreted with caution. This is particularly true for the raw trends in interactions. While we have attempted to make our database of false stories as comprehensive as possible, it is likely far from complete, and many factors could

generate selection biases that vary over time. The raw decline in Facebook engagements may partly reflect the under-sampling of sites that could have entered or gained popularity later in our sample period, as well as efforts by producers of misinformation to evade detection on Facebook by changing their domain names. It may also reflect changes over time in demand for highly partisan political content that would have existed absent efforts to fight misinformation, and could reverse in the future, for example in the run-up to future elections. Actions by policymakers and civil society organizations to improve media literacy could have also affected the observed trends, independently of actions by the platforms (Zubrzycki 2017, Strauss 2018).

We see the comparison of Facebook engagements to Twitter shares as potentially more informative. If the design of these platforms and the behavior of their users were stable over time, we might expect sample selection biases or demand changes to have similar proportional effects, and thus leave the ratio of Facebook engagements to Twitter shares roughly unchanged. For example, we might expect producers changing domain names to evade detection to produce similar declines in our measured interactions on both platforms. The fact that Facebook engagements and Twitter shares follow similar trends prior to late 2016 and for the non-fake-news sites in our data, but diverge sharply for fake news sites following the election, suggests that some factor has slowed the relative diffusion of misinformation on Facebook. The suite of policy and algorithmic changes made by Facebook following the election seems like a plausible candidate.


However, even the relative comparison of the platforms is only suggestive. Both Facebook and Twitter have made changes to their platforms, and so this measure at best captures the relative effect of the former compared to the latter. Engagements on Facebook affect sharing on Twitter and vice versa. The selection of stories into our database could for various reasons differentially favor the kinds of stories likely to be shared on one platform or the other, and this selection could vary over time. Demand changes need not have the same proportional effect on the two platforms. Some of these factors would tend to attenuate changes in the Facebook-Twitter ratio, leading our results to be conservative, but others could produce a spurious decrease over time.

We report a number of robustness checks in the appendix, most of which are designed to address concerns about selection into our sample of sites. First, we restrict to sites that are identified as fake news sites by at least two or three of our original five lists, which leaves 116 and 19 sites, respectively. Second, given that people might disagree with any one particular study's list of fake news sites, we run five additional analyses, each excluding fake news sites identified exclusively by one of our five lists. Third, we focus on sites that started active operations after November 2016, sites that were still in active operation as of July 2018, and sites that were in active operation from August 2015 to July 2018, which leaves 226, 215, and 82 sites respectively. (Active operation is defined to be a global traffic rank reported by Alexa of at least one million.) Fourth, we exclude the five largest sites in terms of total interactions to ensure the trend is not driven solely by outliers. We also look at sites in the first decile and sites in the bottom nine deciles separately to see if the


trend holds for both large sites and small sites. Fifth, Grinberg et al. (2018) provide three lists of sites classified by different likelihoods to publish misinformation. We look at each of these lists separately. Our main qualitative conclusions remain consistent across these checks, though the exact size and shape of the trends vary. Sixth, we present an alternative comparison group: a small set of politically focused sites such as Politico and The Hill. These sites do see a decline in engagements on Facebook relative to Twitter, but it mainly occurred in late-2015. Finally, we show that results look similar when using only the count of Facebook shares instead of engagements, which includes shares, comments, and reactions such as “likes.”

Figure 1: Engagement on Facebook and Twitter

Panel A: Facebook Engagements

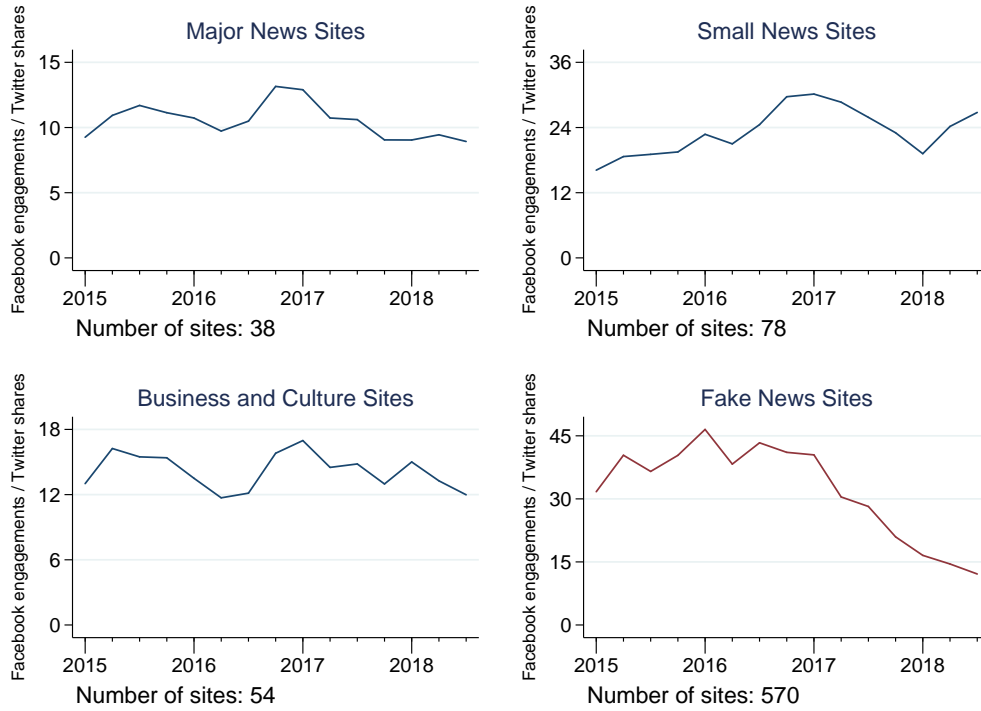


Panel B: Twitter Shares




Notes: This figure shows monthly Facebook engagements and Twitter shares of all articles published on sites in different categories averaged by quarter. Data comes from BuzzSumo. Major News Sites include 38 sites selected from the top 100 sites in Alexa’s News category. Small News Sites include 78 sites selected from the sites ranking 401-500 in the News category. Business and Culture Sites include 54 sites selected from the top 50 sites in each of the Arts, Business, Health, Recreation, and Sports categories. Fake News Sites include 570 sites assembled from five lists. The complete lists can be found in the appendix.

Figure 2: Relative Engagement on Facebook



Notes: This figure shows the monthly ratio of Facebook engagements to Twitter shares of all articles published on sites in different categories averaged by quarter. Data comes from BuzzSumo. Major News Sites include 38 sites selected from the top 100 sites in Alexa’s News category. Small News Sites include 78 sites selected from the sites ranking 401-500 in the News category. Business and Culture Sites include 54 sites selected from the top 50 sites in each of the Arts, Business, Health, Recreation, and Sports categories. Fake News Sites include 570 sites assembled from five lists. The complete lists can be found in the appendix.

Figure 3: Engagement on Facebook and Twitter for Fake News URLs



Notes: This figure shows Facebook engagements, Twitter shares, and the ratio of Facebook engagements to Twitter shares of a set of 10,240 URLs spreading misinformation. See the Data section for details on how the set of URLs is constructed. We sum the Facebook engagements and Twitter shares of all URLs by month and average by quarter.

References

- Allcott H, Gentzkow M (2017) Social media and fake news in the 2016 election. *Journal of Economic Perspectives* 31(2): 211–236.
- Azzimonti M, Fernandes M (2018) Social media networks, fake news, and polarization. *NBER Working Paper No. 24462*.
- Blair S, et al. (2017) Real solutions for fake news? Measuring the effectiveness of general warnings and fact-check tags in reducing belief in false stories on social media. *Working Paper*. Available at <https://www.dartmouth.edu/~nyhan/fake-news-solutions.pdf>. Accessed September 6, 2018.
- Brayton E (2016) Please stop sharing links to these sites. *Patheos*. Available at <http://www.patheos.com/blogs/dispatches/2016/09/18/please-stop-sharing-links-to-these-sites/>. Accessed September 5, 2018.
- Flynn DJ, Nyhan B, Reifler J (2017) The nature and origins of misperceptions: Understanding false and unsupported beliefs about politics. *Political Psychology* 38(S1): 127-50.
- Funke D (2018) Fact-checkers have debunked this fake news site 80 times. It’s still publishing on Facebook. *Poynter.org*. Available at <https://www.poynter.org/news/fact-checkers-have-debunked-fake-news-site-80-times-its-still-publishing-facebook>. Accessed September 4, 2018.
- Ghosh, D, Scott B (2018) Disinformation is becoming unstoppable. *Time*. Available at <http://time.com/5112847/facebook-fake-news-unstoppable/>. Accessed September 2, 2018.
- Gillin J (2017) Politifact’s guide to fake news websites and what they peddle. *PolitiFact*. Available at <http://www.politifact.com/punditfact/article/2017/apr/20/politifacts-guide-fake-news-websites-and-what-they/>. Accessed September 3, 2018.
- Grinberg N, Joseph K, Friedland L, Swire-Thompson B, Lazer D (2018) Fake news on Twitter during the 2016 U.S. presidential election. *Working Paper*. Available from the authors.
- Guess A, Nyhan B, Reifler J (2018) Selective exposure to misinformation: evidence from the consumption of fake news during the 2016 US presidential campaign. *Working Paper. European Research Council*. Available at <https://www.dartmouth.edu/~nyhan/fake-news-2016.pdf>. Accessed September 3, 2018.
- Gunther R, Beck PA, Nisbet EC (2018) Fake news may have contributed to Trump’s 2016 victory. Available at <https://assets.documentcloud.org/documents/4429952/Fake-News-May-Have-Contributed-to-Trump-s-2016.pdf>. Accessed September 3, 2018.
- Lazer DM, Baum MA, Benkler Y, Berinsky AJ, Greenhill KM, Menczer F, Metzger MJ, Nyhan B, Pennycook G, Rothschild D, Schudson M (2018). The science of fake news. *Science* 359(6380): 1094-1096.
- Levin S (2017) Facebook promised to tackle fake news. But the evidence shows it’s not working.

- The Guardian*. Available at <https://www.theguardian.com/technology/2017/may/16/facebook-fake-news-tools-not-working>. Accessed September 3, 2018.
- NewsWhip (2018) Navigating the Facebook algorithm change: 2018 report. Available at <http://go.newswhip.com/rs/647-QQK-704/images/FacebookAlgorithmMarch18.pdf>. Accessed September 3, 2018.
- Parkinson HJ (2016) Click and elect: how fake news helped Donald Trump win a real election. *The Guardian*. Available at <https://www.theguardian.com/commentisfree/2016/nov/14/fake-news-donald-trump-election-alt-right-social-media-tech-companies>. Accessed September 2, 2018.
- Pennycook G, Rand DG (2017) The implied truth effect: attaching warnings to a subset of fake news stories increases perceived accuracy of stories without warnings. *Working Paper*. Available at <http://dx.doi.org/10.2139/ssrn.3035384>. Accessed September 3, 2018.
- Resnick P, Ovadya A, Gilchrist G (2018) Iffy Quotient: A Platform Health Metric for Misinformation. Available at <https://csmr.umich.edu/>.
- Schaedel S (2017) Websites that post fake and satirical stories. *FactCheck*. Available at <https://www.factcheck.org/2017/07/websites-post-fake-satirical-stories>. Accessed September 3, 2018.
- Silverman C (2016) Here are 50 of the biggest fake news hits on Facebook from 2016. *BuzzFeed News*. Available at <https://www.buzzfeednews.com/article/craigsilverman/top-fake-news-of-2016>. Accessed September 3, 2018.
- Silverman C, Lytvynenko J, Pham S (2017a) These are 50 of the biggest fake news hits on Facebook in 2017. *BuzzFeed News*. Available at <https://www.buzzfeednews.com/article/craigsilverman/these-are-50-of-the-biggest-fake-news-hits-on-facebook-in> Accessed September 3, 2018.
- Silverman C, Singer-Vine J, Vo LT (2017b) In spite of the crackdown, fake news publishers are still earning money from major ad networks. *BuzzFeed News*. Available at <https://www.buzzfeednews.com/article/craigsilverman/fake-news-real-ads>. Accessed September 3, 2018.
- Spohr D (2017) Fake news and ideological polarization: filter bubbles and selective exposure on social media. *Business Information Review* 34(3): 150–160.
- Strauss V (2018) The News Literacy Project takes on ‘fake’ news — and business is better than ever. *The Washington Post*. Available at <https://www.washingtonpost.com/news/answer-sheet/wp/2018/03/27/not-sure-whats-real-or-fake-anymore-the-news-literacy-project-teaches-kids-how-to-tell-the-difference-and-its-growing-faster-than-ever/>. Accessed September 3, 2018.
- Zimdars M (2016) False, misleading, clickbait-y, and satirical “news” sources. Available at <http://d279m997dpfwgl.cloudfront.net/wp/2016/11/Resource-False-Misleading-Clickbaity-and-Satirical-%E2%80%9CNews%E2%80%9DSources-1.pdf>. Accessed September 5, 2018.
- Zubrzycki J (2017) More states take on media literacy in schools. *Education Week*. Available

at http://blogs.edweek.org/edweek/curriculum/2017/07/media_literacy_laws.html. Accessed September 3, 2018.