

NBER WORKING PAPER SERIES

SOCIAL MEDIA AND FAKE NEWS IN THE 2016 ELECTION

Hunt Allcott
Matthew Gentzkow

Working Paper 23089
<http://www.nber.org/papers/w23089>

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
January 2017

We are grateful to Chuan Yu and Nano Barahona for research assistance, and we thank Stanford University for financial support. Our survey was determined to be exempt from human subjects review by the NYU and Stanford Institutional Review Boards. The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research.

At least one co-author has disclosed a financial relationship of potential relevance for this research. Further information is available online at <http://www.nber.org/papers/w23089.ack>

NBER working papers are circulated for discussion and comment purposes. They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.

© 2017 by Hunt Allcott and Matthew Gentzkow. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

Social Media and Fake News in the 2016 Election
Hunt Allcott and Matthew Gentzkow
NBER Working Paper No. 23089
January 2017
JEL No. C52,C53,D7,H0,J60

ABSTRACT

We present new evidence on the role of false stories circulated on social media prior to the 2016 US presidential election. Drawing on audience data, archives of fact-checking websites, and results from a new online survey, we find: (i) social media was an important but not dominant source of news in the run-up to the election, with 14 percent of Americans calling social media their “most important” source of election news; (ii) of the known false news stories that appeared in the three months before the election, those favoring Trump were shared a total of 30 million times on Facebook, while those favoring Clinton were shared eight million times; (iii) the average American saw and remembered 0.92 pro-Trump fake news stories and 0.23 pro-Clinton fake news stories, with just over half of those who recalled seeing fake news stories believing them; (iv) for fake news to have changed the outcome of the election, a single fake article would need to have had the same persuasive effect as 36 television campaign ads.

Hunt Allcott
Department of Economics
New York University
19 W. 4th Street, 6th Floor
New York, NY 10012
and NBER
hunt.allcott@nyu.edu

Matthew Gentzkow
Department of Economics
Stanford University
579 Serra Mall
Stanford, CA 94305
and NBER
gentzkow@stanford.edu

1 Introduction

American democracy has been repeatedly buffeted by changes in media technology. In the 19th century, cheap newsprint and better presses allowed partisan newspapers to dramatically expand their reach. Many have argued that the effectiveness of the press as a check on power was significantly compromised as a result (Kaplan 2002). In the 20th century, as radio and then television became the dominant platforms, observers worried they would reduce substantive policy debates to sound bites, privilege charismatic or “telegenic” candidates over those who might have more ability to lead but are less polished, and concentrate power in the hands of a few large corporations (Lang and Lang 2002; Bagdikian 1983).¹ At the turn of the 21st century, the gradual shift of media consumption online prompted a new set of concerns, among them that excess diversity of viewpoints would make it easier for like-minded citizens to form “echo chambers” or “filter bubbles” where they would be insulated from contrary perspectives (Sunstein 2001a, 2001b, 2009; Pariser 2011).

Following the 2016 presidential election, the focus of concern has shifted to social media. Social media platforms, such as Facebook, have a dramatically different structure than any previous media technology. Content can be relayed among users with no significant third party filtering, fact checking, or editorial judgement, and an individual user with no track record or reputation can in some cases reach as many readers as Fox News, CNN or the *New York Times*. Among the most prominent concerns has been the impact of false or misleading information – “fake news,” as it has been dubbed in the public discussion. Recent evidence shows that: (i) 62 percent of U.S. adults get news on social media (Pew 2016a); (ii) the most popular fake news stories were more widely shared on Facebook than the most popular mainstream news stories (Silverman 2016); (iii) many people who see fake news stories report that they believe them (Silverman and Singer-Vine 2016); and (iv) the most discussed fake news stories tended to favor Donald Trump over Hillary Clinton (Silverman 2016). Putting these facts together, a number of analysts and commentators have suggested that Donald Trump would not have been elected president were it not for the influence of fake news spread through social media.²

¹After the first televised presidential debate, between Kennedy and Nixon in 1960, historian Henry Steele Commager (1960) wrote that he hoped “TV debates will be eliminated from future presidential campaigns. [...] The present formula of TV debate is designed to corrupt the public judgment and, eventually, the whole political process.”

²See, for example, Parkinson (2016) and Read (2016). In an interview with the *Washington Post*, fake news writer and promoter Paul Horner said that “I think Trump is in the White House because of me. His followers don’t fact-check anything – they’ll post everything, believe anything. His campaign manager posted my story about a protester getting paid \$3,500 as fact. Like, I made that up. I posted a fake ad on Craigslist” (Dewey 2016).

In this paper, we present new evidence on the role of social media and fake news. We begin by quantifying the importance of social media relative to other information sources in the 2016 election. Previous published work, new web browsing data, and a 1200-person post-election online survey we conducted for this paper during the week of November 28, 2016 are consistent in suggesting that social media was an important but not dominant source of information. Our survey suggests that 14 percent of Americans viewed social media as their “most important” source of election news.

Second, we collect a database of election stories that have been categorized as fake by leading fact checking websites in the three months before the election. For each story, we record the number of times it has been shared on social media, and whether its assertion favored Hilary Clinton or Donald Trump. We confirm that fake news was both widely shared and tilted in favor of Donald Trump, with pro-Trump (i.e. anti-Clinton) fake stories in our database shared a total of 30 million times, and pro-Clinton (i.e. anti-Trump) fake stories shared a total of 7.6 million times.

Third, we use the results of our new survey to estimate the share of Americans who saw and believed each of a set of 14 fake news headlines. We present each respondent with a random subset of the headlines, asking for each: (i) “Do you recall seeing this reported or discussed prior to the election?” and (ii) “At the time of the election, would your best guess have been that this statement was true?” To address survey misreporting, we also include a set of *placebo* fake news headlines – untrue headlines that we invented and that never actually circulated. This approach mirrors the use of placebo drugs as controls in clinical trials. Consistent with a similar survey carried out the same week as ours (Silverman and Singer-Vine 2016), about 15 percent of U.S. adults report that they recall seeing the average fake news headline. About 8 percent report seeing and believing it. However, these numbers are statistically identical for our placebo headlines, suggesting that the raw responses could overstate true exposure by an order of magnitude.

Using the difference between actual and placebo stories as a measure of true recall, we estimate that 1.2 percent of people recall seeing the average story. Projecting these per-article exposure rates to the universe of fake news in our database under the assumption that exposure is proportional to Facebook shares, our point estimate suggests that the average voting-age American saw and remembered about 0.92 pro-Trump fake stories and 0.23 pro-Clinton fake stories in the run-up to the election. Our confidence intervals rule out that the average voting-age American saw, remembered, *and believed* more than 0.71 pro-Trump fake stories and 0.18 pro-Clinton fake stories.

Finally, to benchmark the possible effect of fake news on the election, we combine our ex-

posure estimates with a simple voter persuasion model in the spirit of DellaVigna and Kaplan (2007). We assume that each fake news story someone sees and remembers has a fixed probability f of inducing her to vote for the candidate the story favors, assuming she were not voting for that candidate already. We assume that these effects cumulate additively, so that a voter who would otherwise vote for Clinton but sees three pro-Trump articles and one pro-Clinton article, for example, would have probability $(3 - 1)f = 2f$ of shifting her vote to Trump. We also assume that the effect of fake news on vote shares is uniform across the country. Rather than estimate the persuasion rate, we benchmark it against recent estimates of the persuasive effects of television campaign advertisements from Spenkuch and Toniatti (2016). We estimate that in order for fake news to have changed the election result, the average fake story would need to have $f \approx 0.0073$, making it about as persuasive as 36 television campaign ads. The precise answer is sensitive to the specific assumptions we make, but across a range of alternative specifications and robustness checks we consider, the minimum necessary persuasion rate is the equivalent of 13 ads.

We emphasize a number of important caveats. First, our database of fake news includes only those stories archived on prominent fact-checking websites. The effect of fake stories not archived on these sites is thus omitted from our results. Second, we rely on our post-election online survey to gauge exposure to fake stories. Survey-based recall measures in other contexts have been shown to be inflated due to recall error and misreporting (Prior 2009); we attempt to correct for this using placebo stories, but doing so requires assuming that the false recall rate is the same for our placebo stories and the true stories. Moreover, our survey captures recall of fake stories as of three weeks after the election. Any effect of stories that respondents saw but forgot will be omitted from our results. Finally, our calibrated election model relies on a number of other strong assumptions, including a persuasion rate of fake news that is the same for all stories and voters and that cumulates linearly.

Our paper connects to several areas of existing literature. First, many papers measure factual beliefs and reasons for disagreement with mainstream or scientific consensus on issues such as global warming and vaccine effectiveness, including Bullock et al. (2015), Prior et al. (2015), Malka et al. (2009), Nyhan et al. (2013), and Nyhan et al. (2014). Second, there is a literature measuring media exposure, including Guess (2015), Price and Zaller (1993), Prior (2009), and others. Third, papers in political science and economics measure the persuasive effects of media and political advertising, including Bartels (1993), DellaVigna and Kaplan (2007), Enikolopov et al. (2011), Gerber, et al. (2011), Martin and Yurukoglu (2016), Spenkuch and Toniatti (2016), and

others; see DellaVigna and Gentzkow (2010) and Napoli (2014) for overviews. Fourth, we build on papers measuring ideological segregation and how this relates to news consumption, including Bakshy et al. (2015), Gentzkow and Shapiro (2011), Flaxman et al. (2016), and others.

Section 2 details the data. Section 3 discusses the importance of social media as a source of election news, Section 4 describes the quantity and slant of fake news that circulated during the election, and Section 5 presents the results of our fake news survey. Section 6 benchmarks whether fake news could have impacted the presidential election outcome, and Section 7 concludes.

2 Data

2.1 Web browsing data

One way to measure the importance of social media as a news source is to measure the share of traffic on news websites that come from social media vs. other sources. Each time a user visits a webpage, that user has either navigated directly (for example, by typing `www.wsj.com` into a browser) or has been referred from some other site. Major referral sources include social media (for example, clicking on a link in the Facebook news feed) and search engines (for example, searching for “Pope endorsed Trump?” on Google and clicking on a search result).

We use two data sources on referrals to websites covering U.S. news. The first is Alexa (`alexa.com`), which gathers traffic data from browser extensions used by a sample of “millions of internet users,” as well as directly from websites that use Alexa measurement services. We obtain Alexa data for late October through late November 2016. The second source is comScore, which gathers data from another panel of approximately two million internet users. comScore provides longer historical data and is more established, but it has less coverage of lower-traffic websites. In the body of the paper, we focus on Alexa data, but we present parallel results using comScore data in the Appendix, and our conclusions are not sensitive to which data source we use. Both data sources are limited in that they do not include mobile browsers and do not include news articles viewed within social media sites, for example, when people see headlines in their news feeds but do not click through to the news site that originated the article.

2.2 Fake news database

We define “fake news” as *news stories that have no factual basis but are presented as facts*. By “news stories,” we mean stories that originated in social media or the news media; this excludes

false statements originated by political candidates or major political figures. By “presented as facts,” we exclude websites that are well-known to be satire, such as the Onion. There are two remaining challenges. First, we need some objective measure of “no factual basis.” Second, for some of our calculations below, we would like a comprehensive database of fake news articles.

We address these two challenges by using three outside lists of fake news articles posted in the three months before the November election. First, we use Snopes (snopes.com), which calls itself “the definitive Internet reference source for urban legends, folklore, myths, rumors, and misinformation.” Users can submit stories for fact-checking on the Snopes website, and Snopes researchers rate each story as “true,” “false,” or “mixture.” Snopes is non-partisan and makes money only from advertising from Google AdSense, not from donors or sponsors. We scraped all stories dated between August 1st and November 7th, 2016 from <http://www.snopes.com/tag/donald-trump/> and <http://www.snopes.com/tag/hillary-clinton/>.

Second, we use PolitiFact (politifact.com), a fact-checking website that rates the accuracy of political claims. Like Snopes, PolitiFact has a research staff that rates the accuracy of media articles. PolitiFact is run by editors and reporters from the *Tampa Bay Times*, and also like Snopes, PolitiFact is non-partisan and is funded by advertisements. PolitiFact also receives funding from foundations such as the Gates, Ford, and Knight foundations. We scraped all stories dated between August 1st and November 7th, 2016 from <http://www.politifact.com/truth-o-meter/elections/2016/president-united-states/>. Most of these stories are fact checks of statements made by presidential candidates, which we drop, but some are fake news headlines. We use fake news headlines that PolitiFact rated as “Pants on Fire” or “False.”

Third, we use a list of 21 major fake news articles that appeared between August 1st and November 7th, 2016 that was compiled by Craig Silverman, an editor at BuzzFeed (buzzfeed.com). Of these 21 articles, 12 appear in our Snopes database. Nine were rated as “false,” and the other three were rated “mixture,” “unproven,” and “mostly false.” Combining these three lists, we have a database of 156 fake news articles.

We match these articles to data on Facebook shares from BuzzSumo (buzzsumo.com), an online content database that links to the Facebook API and records the number of shares for individual URLs. Individual fake news stories in our database typically occur on multiple URLs – for example, the false story that “the Pope endorsed Donald Trump” was reported independently by a number of different news websites, with different specific URLs for each story. For each story in our fake news database, we searched relevant key words on BuzzSumo, and recorded the number

of Facebook shares for every URL that had been shared more than 1000 times. While BuzzSumo does have shares from other social media sites such as Twitter, we do not record shares on these other sites because the number of Facebook shares is orders of magnitude larger. As we carried out these searches in early December 2016, the number of shares includes several post-election weeks, and thus may overstate the number of pre-election shares. We also gather the number of Facebook shares of the fact-check articles from Snopes.³

Finally, we gather a list of domains that publish large volumes of fake news, based on a list of largely conservative sites compiled by Melissa Zimdars (2016) and a list of largely liberal sites compiled by Ed Brayton (2016).

2.3 Post-election survey

During the week of November 28th, 2016, we conducted an online survey of 1208 U.S. adults aged 18 and over using the SurveyMonkey platform. The sample was drawn from SurveyMonkey's Audience Panel, an opt-in panel recruited from the more than 30 million people who complete SurveyMonkey surveys every month.⁴

The survey consisted of four sections. First, we acquired consent to participate and a promise to provide best answers, which we describe below. Those who did not agree were disqualified from the survey. Second, we asked a series of demographic questions, including political affiliation before the 2016 campaign, vote in the 2016 presidential election, education, and race/ethnicity. Third, we asked about 2016 election news consumption, including time spent on reading, watching, or listening to election news in general and on social media in particular, and the most important source of news and information about the 2016 election. Fourth, we gave 15 news headlines about the 2016 election. For each headline, we asked, "Do you recall seeing this reported or discussed prior to the election?" and "At the time of the election, would your best guess have been that this statement was true?" We also received age and income categories, gender, and Census division from profiling questions that respondents had completed when they first started taking surveys on the Audience panel. The survey instrument can be accessed from <https://www.surveymonkey.com/r/RSYD75P>.

³Some rumors from Snopes were images shared on social media with no specific origin URL, so we do not have Facebook shares of the false "article." In these cases, we impute the Facebook shares of false articles from the Facebook shares of the corresponding Snopes' fact-check articles using a log-log regression, based on the sample of stories for which we have both variables; the R^2 of this regression is 0.17.

⁴For more information, see <https://www.surveymonkey.com/mp/audience/>.

Each respondent's 15 news headlines were randomly selected from a list of 30 news headlines, six from each of five categories. Within each category, our list contains an equal split of pro-Clinton and pro-Trump assertions, so 15 of the 30 articles favored Clinton, and the other 15 favored Trump. The first category contains six fake news stories mentioned in three mainstream media articles (one in the *New York Times*, one in the *Wall Street Journal*, and one in BuzzFeed) discussing fake news during the week of November 14th, 2016. The second category contains the four most recent pre-election headlines from each of Snopes and PolitiFact deemed to be unambiguously false. We refer to these two categories individually as "Big Fake" and "Small Fake" respectively, or collectively as "Fake." The third category contains the most recent six major election stories from the *Guardian's* election timeline. We refer to these as "Big True" stories. The fourth category contains the two most recent pre-election headlines from each of Snopes and PolitiFact deemed to be unambiguously true. We refer to these as "Small True" stories. Our headlines in these four categories appeared on or before November 7th.

The fifth and final category contains invented "Placebo" fake news headlines. We invented three damaging fake headlines that could apply to either Clinton or Trump, then randomized whether a survey respondent saw the pro-Clinton or pro-Trump version. We experimented with several alternative Placebo headlines during a pilot survey, and we chose these three because the data showed them to be approximately equally believable as the "Small Fake" stories. We confirmed using Google searches that none of the Placebo stories had appeared in actual fake news articles. Table 1 presents the exact text of the headlines presented in the survey.

There are two standard concerns with this type of online survey. The first is that respondents might rush through the surveys, providing low-quality and unreliable answers. To reduce panelists' incentive to rush through surveys just for the reward, SurveyMonkey rewards panelists with only non-cash incentives – charitable donations and sweepstakes entries. The median respondent spent about 6 minutes and 30 seconds on our survey, which in our own pilot testing was more than sufficient to carefully think through the questionnaire. On the recommendation of survey design experts, we also added the following question at the beginning of the survey:

We care about the quality of our data. In order for us to get the most accurate measures of your knowledge and opinions, it is important that you thoughtfully provide your best answers to each question in this survey.

Do you commit to thoughtfully provide your best answers to each question in this survey?

- *I will provide my best answers.*
- *I will not provide my best answers.*
- *I can't promise either way.*

95 percent of respondents answered “I will provide my best answers.” Those who did not were disqualified from the survey.

The second concern is that online survey respondents are not nationally representative. Adding to this problem, the Audience Panel is so large that it can provide a complete set of survey responses within hours, but respondent demographics vary by time of day. For example, we launched a small pilot survey one evening at 9:12 PM Pacific Time. We had 60 responses by 10:00 PM, of whom 40 percent lived on the West Coast, and Clinton supporters outweighed Trump supporters by more than a 2-1 margin. We therefore launched our final survey in batches at different times of day over a four-day period.

To improve national representativeness, we re-weight the sample to match the nationwide adult population on 10 observable characteristics that we hypothesized might be correlated with survey responses. Table 2 presents summary statistics for these variables. Unsurprisingly, the unweighted sample we received was not representative of the national population: our sample is disproportionately well-educated, female, and Caucasian, and relies relatively heavily on the web and social media for news. We re-weight the sample in column 1 to match population means in column 2, using the entropy weighting procedure of Hainmueller (2012). All results reported below are weighted.

By construction, the mean weight is one. As diagnostics, the standard deviation of our sample weights is 1.4, the maximum weight is 20.4, 2.3 percent of weights are larger than 5, and 0.25 percent of weights (three observations) are larger than 10. In our unweighted data, Clinton received 15 percentage points more votes than Trump, while in our weighted data, she received 6 percentage points more. The latter margin is statistically indistinguishable from the predictions of most pre-election polls.

3 How important were social media as a news source in the 2016 election?

3.1 Survey data

Pew (2016b) is a frequently cited point of reference on the overall importance of social media. As noted above, 62 percent of US adults in this survey get news from social media. Of these, however, only 18 percent do so “often.” 26 percent do so “sometimes,” and 18 percent do so “never.” By comparison, the shares who “often” get news from local television, national broadcast television, and cable television are 46 percent, 30 percent, and 31 percent respectively. Moreover, only 34 percent of web-using adults trust the information they get from social media “some” or “a lot.” By contrast, this share is 76 percent for national news organizations and 82 percent for local news organizations.

The results of our post-election survey are broadly consistent with this picture. In the month before the 2016 election, our respondents report spending 66 minutes per day reading, watching, or listening to election news. (Again, these and all other survey results are weighted for national representativeness.) Of this, 25 minutes (38 percent) were on social media. Our survey then asked, “which of these sources was your most important source of news and information about the 2016 election?” The word “important” was designed to elicit a combination of consumption frequency and trust in information. Figure 1 plots responses. In order, the four most common responses are cable TV, network TV, websites, and local TV. Social media are the fifth most common response, with 14 percent of U.S. adults listing them as their most “important” news source.

Taken together, these results suggest that social media has become an important but not dominant source of political news and information. Television remains more important by a large margin.

3.2 Web referrals data

To provide another view of the importance of social media, figure 2 breaks down the source of web traffic for news websites into four sources: direct navigation, plus referrals from social media, search engines, and other links. We present results both for the top 690 U.S. news sites and for 65 fake news sites. For the top 690 sites, social media referrals represent only about 10 percent of total traffic. By contrast, fake news websites rely on social media for a much higher share of their traffic. Appendix figure 3 presents parallel results for major U.S. news websites using the comScore data; the share of traffic to mainstream news sites from social media is slightly lower

than in the Alexa data. We again emphasize that these figures exclude mobile browsing and do not capture news viewed directly on social media, for example reading on Facebook or Twitter news feeds.

4 How much fake news circulated in the run-up to the election?

Figure 3 presents the key statistics from our fake news database. We record 41 pro-Clinton and 115 pro-Trump articles, which were shared a total of 7.6 million and 30.3 million times, respectively. There are thus about three times more fake pro-Trump articles than pro-Clinton articles, and the average pro-Trump article was shared more on Facebook than the average pro-Clinton article.

One measure of the comprehensiveness of the three lists is to look at the overlap between the lists of stories from Snopes, PolitiFact, and BuzzFeed. Snopes is our largest list, including 138 of our total of 156 articles. As a benchmark, 12 of the 21 articles in the BuzzFeed list appear in Snopes, and 4 of the 13 articles in the PolitiFact appear in Snopes. The lack of perfect overlap shows that none of these lists is comprehensive and suggests that there may be fake news articles that are omitted from our database.

5 How many people heard and believed fake election news?

5.1 Graphical results

Our post-election survey gave respondents 15 news headlines – three headlines randomly selected from each of the five categories detailed in Section 2.3 – and asked if they recalled seeing the headline (“Do you recall seeing this reported or discussed prior to the election?”) and if they believed it (“At the time of the election, would your best guess have been that this statement was true?”).

Figure 4 presents the share of respondents that recall seeing (left bar) and recall seeing and believe (right bar) headlines, averaging responses across all the headlines within each of our main categories. Rates of both seeing and believing are much higher for true than fake stories, and they are higher for the “Big True” headlines (the major headlines leading up to the election) than for the “Small True” headlines (the more minor fact-checked headlines that we gathered from Snopes and PolitiFact). 15 percent of people report seeing the Fake stories, and 8 percent report seeing and believing them. These shares are broadly consistent with the results of a separate survey

conducted during the same week by Silverman and Singer-Vine (2016); for a set of five fake news stories, they find that the share of respondents who have heard them ranges from 10 to 22 percent.⁵ However, figure 4 also shows that Placebo fake news articles, which never actually circulated, are approximately equally likely to be recalled and believed as the Fake news articles that did actually circulate. This clearly shows that there is a meaningful rate of false recall. If this false recall rate is similar for Fake and Placebo articles, this suggests that the raw responses significantly overstate the circulation of Fake news articles, and that the true circulation of Fake news articles was quite low. Appendix figures 1 and 2 present results for each of the 30 individual headlines.

Figure 5 presents the share of people in our post-election survey who believed the stories in each category, broken down by responses to whether they had heard the headlines. This figure demonstrates that there is a strong correlation between seeing and believing articles. More than half of people who recalled seeing the average Fake or Placebo headline believed it, compared to less than 10 percent who did not recall seeing it. This correlation could arise because seeing an article directly causes people to believe it, and/or because people assess the likelihood that they saw an headline based on how plausible they think it is *ex post*. Since nobody actually saw the Placebo headlines, we know that at least some of this correlation is driven by the latter mechanism.

Two key results from these figures – false recall and connection between recall and plausibility – motivate the simple model of survey response below.

5.2 A simple model of survey response

In Section 6, we want to know two parameters: the share of population that was truly exposed to the average fake news article in our survey, and the share that was truly exposed and believed the average fake news article. Since the above finding of false recall means that true exposure is not directly observed, it is helpful to formalize a simple model of survey response to understand how these two parameters can be inferred.

We assume that the probability that survey respondent i reports seeing (S_{ia}) or believing (B_{ia}) article a is some weakly increasing function G of true exposure $E_{ia} \in \{0, 1\}$ and the plausibility P_{ia} that the respondent assigns to the article. For $Y \in \{S, B\}$, this means that

$$\Pr(Y_{ia} = 1) = G_Y(\beta_Y E_{ia}, \gamma_Y P_{ia}), \quad (1)$$

⁵The authors do not report shares hearing and believing, but the overall share who rate the stories as “very accurate” ranges from 28 to 49 percent.

with $\beta_Y, \gamma_Y \geq 0$. Larger β_S implies better memory, $\beta_B > 0$ if exposure *per se* causes people to believe articles, $\gamma_S > 0$ if respondents consider an article’s plausibility when trying to recall whether they saw it in the media, and $\gamma_B > 0$ simply reflects that more plausible articles are more likely to be believed. We define $M_{ia} \in \{0, 1\}$ as false memory – that is, $M_{ia} = 1$ when $S_{ia} = 1$ but $E_{ia} = 0$.

There are two types of articles, $t \in \{f, p\}$ for Fake and Placebo, and we denote the sets of articles as \mathcal{F} for Fake and \mathcal{P} for Placebo. By construction, the Placebo article exposure rate is zero: $E_{ia} = 0, \forall a \in \mathcal{P}$. Using \mathbb{E} to denote the expectation taken over both individuals and articles, the empirical fact that $\mathbb{E}[S_{ia}|a \in \mathcal{P}] > 0$ demonstrates that $\mathbb{E}[M_{ia}|a \in \mathcal{P}] > 0$. The empirical fact that seeing and believing are correlated for Placebo articles is explained by $\gamma_H, \gamma_B > 0$, i.e. plausibility P_{ia} affects both seeing and believing.

Consider the following two assumptions.

Assumption 1: People do not forget articles if they were actually exposed:

$$S_{ia} = 1 \text{ if } E_{ia} = 1. \tag{2}$$

Assumption 2: For the set of people who misremember seeing articles, plausibility is independent of article type:

$$P_{ia} \perp t, \forall i, a \text{ s.t. } M_{ia} = 1. \tag{3}$$

In essence, Assumption 2 is that Fake and Placebo articles are equally plausible.

We constructed the survey so that these assumptions would be credible. We implemented the survey soon after the election to minimize forgetting and false recall, although we will explore the implications of forgetting in Section 6. Assumption 2 is not directly testable because misremembering is unobserved. However, figure 5 shows an approximate test of Assumption 2 if true exposure rates are small. Specifically, for the share of people who say they were exposed to the article, we see that Fake and Placebo articles are approximately equally likely to be believed. This is approximately a test of Assumption 2 since all people who recalled seeing Placebo headlines are misremembering, as are almost all people who recalled seeing Fake headlines (for small exposure rates). More broadly, Assumption 2 is likely to hold by design because we wrote the Placebo headlines, and refined them in the pilot, to ensure that they were approximately equally plausible as the Fake headlines.

These two assumptions allow us to infer rates of both true exposure as well as true expo-

sure and believing. Under assumptions 1 and 2, it is straightforward to show that $\mathbb{E}[E_{ia}|a \in \mathcal{F}] = \mathbb{E}[S_{ia}|a \in \mathcal{F}] - \mathbb{E}[S_{ia}|a \in \mathcal{P}]$ and $\mathbb{E}[E_{ia}B_{ia}|a \in \mathcal{F}] = \mathbb{E}[S_{ia}B_{ia}|a \in \mathcal{F}] - \mathbb{E}[S_{ia}B_{ia}|a \in \mathcal{P}]$. In words, subtracting the reported rates for Placebo articles from the reported rates for Fake articles gives the true rates for Fake articles. Intuitively, this is the case because Placebo headlines that are calibrated to be equally-plausible provide a control for false recall.

5.3 Regression results

Table 3 presents estimates of average rates of seeing as well as seeing and believing. The left three columns present estimates for seeing, while the right three columns report estimates for seeing and believing. Within each trio of columns, the first column uses only the sample of Fake headlines, the second uses only the sample of Placebo headlines, and the third reports the difference between the first two columns. In this table and all tables below, standard errors are robust and clustered by respondent, and observations are weighted for national representativeness.

Columns 1 and 4 show that 15.3 or 7.9 percent of U.S. adults, respectively, recall seeing or recall seeing and believed the average Fake news story in our survey. This seems to suggest significant circulations of fake news. However, columns 2 and 5 show that 14.1 or 8.3 percent of U.S. adults, respectively, recall seeing or recall seeing and believed the average Placebo news story. This benchmarks the amount of noise and recall error in the data. Columns 3 and 6 show that the differences are not statistically significantly different from zero, meaning that people were no more likely to recall seeing actual Fake headlines than they were to falsely recall Placebo headlines. The point estimate for true exposure in column 3 is 1.2 percent, and the 95 percent confidence interval allows us to exclude values greater than 2.88 percent. The point estimate for the rate of true exposure and belief is actually negative, and the 95 percent confidence interval allows us to exclude values greater than 0.9 percent. We use these values in our election calibration in Section 6.

We can compare these results to the number of times these stories were shared on Facebook. According to our database, the average Fake headline in our survey was shared 386,000 times. Given the U.S. voting age population for 2015 of 248 million people (Federal Register 2016), this implies that the number of shares per person is $0.386/248 \approx 0.0016$, or about 0.16 percent. Our point estimate thus implies that $1.2/0.16 \approx 7.5$ people see and remember a story for each time it is shared.

What types of people are more likely to see or see and believe fake news? We can answer this question by regressing indicators Y_{ia} (representing either recalling seeing S_{ia} or believing B_{ia}) on a

vector \mathbf{X}_i of individual characteristics in a linear probability model regression:

$$Y_{ia} = \boldsymbol{\alpha}_1 \mathbf{X}_i + \alpha_0 + \varepsilon_{ia}. \quad (4)$$

Table 4 presents results. The first three columns are the same as in table 3. The last two columns use believing, instead of seeing and believing, as the dependent variable. In all columns, \mathbf{X}_i includes the reported variables plus additional demographic controls: income, age, race, and gender.

Column 1 includes only Fake news articles. People who report that social media are their most important source of election news, as well as social media users who have ideologically “segregated networks” (with three-quarters or more of their friends preferring the same presidential candidate), are more likely to recall seeing Fake headlines. Compared to both Republicans and independents (the omitted political affiliation category), Democrats are less likely to recall seeing Fake news.

Column 2, however, shows a similar pattern of results for Placebo articles. Indeed, column 3 shows that the differences in $\boldsymbol{\beta}_1$ coefficients between Fake and Placebo headlines are all statistically indistinguishable from zero. In the context of our model of survey response, this suggests that differences in *recalled* exposure – for example, that heavy social media users are more likely to recall seeing fake news articles – are not primarily driven by differences in *true* exposure. If differences in recalled exposure were driven by differences in true exposure, then social media users would be statistically more likely to recall seeing Fake articles relative to Placebo articles. Instead, it appears that the correlations in column 1 of table 4 are driven by differences in the way that people assess the plausibility of fake headlines – for example, that heavy social media users are more likely to find fake news plausible. Our model suggests that differences in recalled exposure then arise because people falsely recall seeing articles that they find plausible. The one near-exception is that social media users with segregated networks are marginally insignificantly more likely to recall seeing Fake headlines compared to Placebo headlines, with a point estimate of 3.4 percentage points.

Columns 4 and 5 present comparable regressions with the believed indicator B_{ia} as the dependent variable. In our model, these are not direct tests of the perceived plausibility of a headline, because believing is also directly affected by true exposure. However, since true exposure appears to be quite low and appears not to differ significantly by the \mathbf{X}_i covariates, these columns are closely approximate tests. Consistent with our argument above, we see that social media users

are more likely to believe both Fake and Placebo articles. In both columns 4 and 5, people with high school or less education are more likely to believe Fake and Placebo headlines than people with more education, and Republicans are more likely to believe Fake and Placebo headlines than Democrats or independents.

Ideological segregation in news sources and social networks could imply that Democrats are more likely to be exposed to pro-Clinton fake news, and Republicans are more likely to be exposed to pro-Trump fake news. Furthermore, for various reasons, Democrats might be more likely to *believe* pro-Clinton fake news, and vice-versa. Table 5 tests this in our post-election survey data, by re-running equation (4) using only political affiliation indicators for \mathbf{X}_i and interacting with indicators for whether article a is pro-Clinton or pro-Trump. The three political affiliation indicators and pro-Clinton vs. pro-Trump indicators both saturate the regression, so there is no constant. The samples and columns are the same as table 4.

Comparing the magnitudes of the different coefficients in column 1 suggests that fake news exposure might be ideologically segregated: Republicans are more likely than independents, and independents in turn more likely than Democrats, to report seeing pro-Trump headlines, although for pro-Clinton headlines, the differences are less stark. In column 2, however, we see very similar results for Placebo headlines, and column 3 shows that five of the six coefficients do not differ for Fake relative to Placebo. In the context of our model, we interpret these results similarly to the results of table 4: differences across people in *recalled* exposure seem to be primarily driven by differences in perceived plausibility, and less by differences in true exposure. There may still be differences in true exposure, but this would need to be documented with web browsing data instead of our survey recall measures.

Are there indeed differences in believing rates that might generate these differences in false recall? Columns 4 and 5 indeed show dramatic differences: Republicans are four to eight times as likely as Democrats to report believing pro-Trump headlines, and Democrats are 50 to 100 percent more likely than Republicans to believe pro-Clinton headlines. Appendix table 1 repeats these regressions in the subsample of social media users with ideologically segregated networks. The relative ordering of coefficients is similar, but the magnitudes are considerably larger: social media users with segregated networks are, as we saw above, more likely to report seeing and believing fake news, and relatively more likely to report seeing and believing fake news that favors their candidates.

6 What impact might fake news have had on the election?

Did fake news shift the election from a Clinton victory to a Trump victory? We formalize a simple extension of DellaVigna and Kaplan’s (2007) persuasion rate model to benchmark the possible impact of fake news on Trump’s margin of victory in the election.

Because we do not have sufficient power to estimate fake news exposure separately by state, we imagine that in a counterfactual world without fake news, Trump’s margin of error would change by the same amount in all states. Under this assumption, for Clinton to have won the election, Trump’s margin of victory would have to decrease by $\Delta m^* \approx 0.51\%$ of the voting age population, which would shift Michigan, Pennsylvania, and Wisconsin into Clinton wins and deliver the Electoral College.⁶ Thus, the core question is whether fake news could have increased Trump’s margin of error by more than 0.51 percent of the voting age population.

6.1 Framework

There are three possible candidates for president, $k \in \{C, T, N\}$, where C and T refer to Clinton and Trump, respectively, and N refers to any alternative, including both third-party candidates and not voting. We define the vote share of candidate k to be the total number of voters choosing k divided by the voting age population. Consider two potential outcomes: \tilde{v}_k denotes vote share in a counterfactual world without fake news, and v_k denotes vote share in an observed world with fake news. The effect of fake news on k ’s vote share is $\Delta v_k = v_k - \tilde{v}_k$. We refer to a voter’s “type” as the candidate whom she would vote for in the absence of fake news, and for notational simplicity we also index types by k .

Fake news persuades some voters to vote for a different candidate. Define e_T and e_C , respectively, as the exposure rate to pro-Trump and pro-Clinton articles, in units of articles viewed per voting-age person. Let $e = e_T - e_C$ be the “net” pro-Trump exposure rate; figure 3 shows that $e > 0$. Define f as the persuasion rate per net pro-Trump article: the proportion of type C or N voters that switch to voting for Trump for each unit increase in e .⁷ In this formulation, the presence of fake news increases Trump’s vote share from type N and C voters, does not shift any type N voters to

⁶This is the margin of victory as a percent of the total voting-age population, which will be the relevant statistic for our model. In other settings, one might consider the margin of victory as a percent of the total number of votes or of the total number of Clinton plus Trump votes.

⁷The assumption that persuasion depends on $e_T - e_C$ simplifies the model. Alternatively, we could assume that persuasion depends on the net number of pro-Trump articles that *each individual* has seen, and integrate over individuals assuming stochastic arrival of article views to individuals.

Clinton, and does not affect type T voters at all. The effect of fake news on N and C vote shares is

$$\Delta v_k = -\tilde{v}_k ef, k \in \{N, C\}. \quad (5)$$

Substituting $\tilde{v}_k = v_k - \Delta v_k$ and simplifying, we have

$$\Delta v_k = -v_k \frac{ef}{1-ef}, k \in \{N, C\}. \quad (6)$$

In our model, Trump gets all the voters that Clinton and Neither lost as a result of fake news: $\Delta v_T = -\Delta v_C - \Delta v_N$. Then, the effect on Trump's margin of victory relative to Clinton is $\Delta m = \Delta v_T - \Delta v_C = -\Delta v_N - 2\Delta v_C$: every type- N voter persuaded increases Trump's vote share, and every type- C voter persuaded does double duty by reducing Clinton's vote share and increasing Trump's vote share.

Substituting equation (6) into the equation for Δm and simplifying gives:

$$\Delta m = (v_N + 2v_C) \frac{ef}{1-ef}. \quad (7)$$

The persuasion rate necessary to change the vote margin by Δm^* is:

$$f^* = \frac{\Delta m^*}{e(v_N + 2v_C + \Delta m^*)}. \quad (8)$$

From the count of the U.S. voting-age population and the vote counts for Clinton and Trump, we have $v_N = 0.480$ and $v_C = 0.266$, which implies $v_N + 2v_C \approx 1$. Furthermore, ef will be much closer to zero than to one, so $\Delta m \approx ef$. Roughly speaking, then, our question hinges on whether the per-capita exposure to pro-Trump fake news articles times the persuasion rate is larger than Δm^* , or 0.51 percent.

6.2 Empirical calibration

As a reference point for persuasion rates, we use Spenkuch and Toniatti's (2016) estimate of the effect of television campaign advertising. Using data from the 2004 and 2008 presidential elections, and exploiting discontinuous changes in advertising exposure at television market boundaries for identification, they estimate that increasing the difference between Republican and Democratic campaign ad exposures per voting age adult by 10 increases the difference between the vote shares of the Republican and Democratic candidate (i.e., the vote margin) by $\Delta m = 0.00175$, or 0.175

percentage points (table 8, column 9). Plugging in this value along with the average of $v_N + 2v_C$ for 2004 and 2008 into equation 8 yields a persuasion rate of $f = 0.0002$.

We judge this to be the most credible and comparable estimate of the effect of television ads in the literature. It aligns reasonably well with the results of a highly-cited study by Huber and Arceneaux (2007), whose results imply persuasion rates of $f = 0.0003$ for Republican ads and $f = 0.0002$ for Democratic ads. Two other studies find larger persuasive effects. Gordon and Hartmann (2013) use a difference-in-difference design and ad price instruments to estimate ad effects in the 2000 and 2004 elections, and find results that imply a persuasion rate of $f = 0.002$. Gerber *et al.* (2011) use a randomized experiment to estimate the effects of television ads in a Texas gubernatorial election on vote intentions as reported in a survey, and find results that imply a persuasion rate of $f = 0.008$. The larger estimate in the first study may reflect differences in the research design: Spenkuch and Toniatti (2016) show that using a panel specification rather than their preferred market-boundary natural experiment increases the estimated effects by an order of magnitude. The larger estimate in the second study may reflect differential responsiveness of survey measures as compared to actual votes, as well as the fact that Gerber *et al.* (2011) measure outcomes within a week after the ad exposure. Studies from other domains also find larger persuasion rates. Gerber and Green (2000), for example, use a randomized experiment to estimate effects of get-out-the-vote mailings on voter turnout of $f = 0.005$.⁸

Our primary approach to calculating e is to calibrate it with a combination of our post-election survey data and our fake news database. The average Fake article in our post-election survey was shared 0.386 million times on Facebook, and the point estimate in column 3 of table 3 is that about 1.2 percent of the population was actually exposed to the average Fake article. This gives $\frac{0.012}{0.386 \text{ million}} \approx 0.03$ exposures per million Facebook shares. As shown in figure 3, there were 30.3 and 7.6 million pro-Trump and pro-Clinton fake news shares on Facebook, respectively, for articles appearing in the three months leading up to the election. Multiplying these gives that the average U.S. adult heard $\hat{e}_T \approx 0.03 \times 30.3 \approx 0.92$ pro-Trump and $\hat{e}_C \approx 0.03 \times 7.6 \approx 0.23$ pro-Clinton fake news articles. Subtracting the Clinton from the Trump totals gives 22.7 million net pro-Trump fake news shares, which gives exposure $\hat{e} \approx 0.03 \times (30.3 - 7.6) \approx 0.69$ net pro-Trump fake news articles. If, in order to be persuaded to change their votes, voters must believe a fake news article instead of just being exposed to it, this is an upper bound on persuasion given our survey results

⁸DellaVigna and Gentzkow (2011) report $f = 0.01$ for 1-3 GOTV cards mailed in Gerber and Green (2000), and the treatment group was sent an average of two cards.

showing that many people do not report believing fake news.

Figure 6 presents results of our primary calibration. The y-axis presents the impact on Trump’s margin of victory Δm , highlighting the $\Delta m^* \approx 0.51$ percent to convert Wisconsin and thus change the election outcome. The x-axis modifies the persuasion rate per fake news impression as a multiple of the Spenkuch and Toniatti (2016) estimate of $f_{ad} = 0.0002$ per campaign ad. For example, a multiplier of one (near the left side of the graph) means that one fake news article has the same effect on votes as one campaign ad, whereas a multiplier of 60 (near the right side of the graph) means that one fake news article has the same effect as 60 campaign ads. The upward-sloping line presents our calibration using $\hat{e} \approx 0.69$. The required persuasion rate is $f^* \approx 0.0073$, i.e. that a single news article would need to convince 0.73 percent of voters who saw the article and would otherwise not vote for Trump to vote for him. For fake news to have changed the election outcome, one fake news article would need to be 36 times as persuasive as one political ad.

6.3 Alternative assumptions

Several important assumptions underlie our primary estimate of the exposure rate e . We now explore the implications of alternative assumptions.

Assumption 1 in Section 5 was that survey respondents always recalled seeing fake news articles that they had actually been exposed to. To address the possibility that survey respondents forgot fake news articles, one set of alternative assumptions is that survey respondents remember share $\beta_H \in [0, 1]$ of the fake news stories that they were exposed to, so true exposure rates are \hat{e}/β_H . Rows 2 and 3 of table 6 present alternative results with $\beta_H = 0.75$ and $\beta_H = 0.5$, respectively.

Our fake news database is not a comprehensive list of all fake news articles, which – like forgotten exposures – causes us to underestimate total exposure. Rows 4 and 5 consider the alternative assumptions that the total fake news shares (and thus exposure rates) are 50 and 100 percent larger than in our database.

In our calculations, we have allowed fake news to be persuasive even if people don’t believe it. We can alternatively measure “exposure” as the combination of being exposed to and believing an article, which will reduce our assumed e . Row 6 implements this by calculating \hat{e} based on the upper bound of the 95 percent confidence interval from column 6 of table 3, reflecting estimates for “recalled seeing and believed” instead of just “recalled seeing.” The average American saw, remembered, and believed $\hat{e}_T \approx 0.71$ pro-Trump and $\hat{e}_C \approx 0.18$ pro-Clinton fake news articles, giving $e \approx 0.53$. This bound reduces our assumed exposure rate by about one-quarter.

A secondary way of assessing the exposure rate e is to use the Facebook shares gathered in the fake news database. From figure 3, there were 23 million more Facebook shares of pro-Trump than pro-Clinton articles. Using the fact that there are 248 million U.S. adults aged 18 and older (Federal Register 2016) and the rule of thumb that one Facebook share causes P people to see an article, the average U.S. adult saw an estimated $ne = \frac{23 \cdot P}{248}$ fake news articles in the run-up to the election. Rows 7 and 8, respectively, show results for $P = 10$ and $P = 20$, respectively. This gives exposure rates of 0.93 and 1.85, respectively, and required multipliers of 27 and 13.

A final extension is to allow exposure rates to pro-Clinton vs. pro-Trump fake news to differ by voter type. Bakshy, Messing, and Adamic (2015), for example, show that on Facebook, liberals are more likely to be exposed to liberal news, and conservatives to conservative news, largely due to their friend networks' sharing patterns. If type C or N voters are similarly exposed to a higher ratio of pro-Clinton to pro-Trump stories than average, this will reduce the impact of fake news on their votes.

To extend our model, define e_{Tk} and e_{Ck} as exposure rates of type- k voters to pro-Trump and pro-Clinton news, respectively, and define $e_k \equiv e_{Tk} - e_{Ck}$ as type k 's net pro-Trump exposure rate. Bakshy, Messing, and Adamic (2015, figure 3B) find that liberals' exposure to conservative news on Facebook is 49 percent of what it would be if exposure were random. On the basis of this result, we assume that type C voters are exposed to 49 percent of the pro-Trump fake news, and 150 percent of the pro-Clinton fake news, than the average person. This significantly down-weights exposure to pro-Trump fake articles and up-weights exposure to pro-Clinton fake articles. Repeating the above calculation of exposure rates gives $\hat{e}_C \approx 0.49 \times 0.03 \times 30.3 - 1.5 \times 0.03 \times 7.6 \approx 0.1$ net pro-Trump articles, about seven times less than in our initial calibration. We assume that e_N is unchanged. Using modified versions of equations (7) and (8), we find $f^* \approx 0.013$, which gives $f^*/f_{ad} = 66$.

Across all these alternative assumptions, the minimum required exposure rate is $f \approx 0.0027$, which is 13 times the effect of a television campaign ad.

7 Conclusion

As a concluding note, we observe that rumors, conspiracy theories, and other cousins of fake news are not new to the social media era. Figure 7 considers 14 conspiracy theories with political implications that have circulated over the past half-century. Using polling data compiled by the

American Enterprise Institute (2013), we plot the share of people who believed each statement is true, from polls conducted in the listed year. These conspiracy theories are slightly different than most of the fake news we study, in the sense that many fake news articles can be traced back to a single person who invented the article without any facts to back it up, whereas some conspiracy theories could in principle be true and often have no unique origin. Notwithstanding, they are an interesting historical benchmark.

For example, during the 2016 campaign, Donald Trump and various online media outlets reopened speculation that Bill Clinton aide Vince Foster had been murdered, whereas the five investigations into his death had concluded that it was a suicide (Kessler 2016). Four official investigations were completed between 1993 and 1995, all of which concluded that his death was a suicide (Kessler 2016). In 1997, independent counsel Kenneth Starr released a fifth report on the matter, concluding that “In sum, based on all of the available evidence, which is considerable, the [Office of Independent Counsel] agrees with the conclusion reached by every official entity that has examined the issue: Mr. Foster committed suicide” (Kessler 2016). Figure 7 illustrates how divergent conclusions on factual issues predates the social media era: as of 1995, 20 percent of Americans reported believing that Foster had been murdered.

In summary, our data suggest that social media were not the most important source of election news, and even the most widely circulated fake news stories were seen by only a small fraction of Americans. For fake news to have changed the outcome of the election, a single fake news story would need to have convinced about 0.7 percent of Clinton voters and non-voters who saw it to shift their votes to Trump, a persuasion rate equivalent to seeing 36 television campaign ads.

References

- Alexa. <http://www.alexa.com/>.
- American Enterprise Institute. 2013. Public opinion on conspiracy theories. https://www.aei.org/wp-content/uploads/2013/11/-public-opinion-on-conspiracy-theories_181649218739.pdf.
- American National Election Studies. 2012. Times series cumulative data file. http://www.electionstudies.org/studypages/anes_timeseries_cdf/anes_timeseries_cdf.htm.
- Bagdikian, Ben. 1983. *The Media Monopoly*. Beacon Press.
- Bakshy, Eytan, Solomon Messing, and Lada A. Adamic. 2015. Exposure to ideologically diverse news and opinion on Facebook. *Science* 348(6239): 1130-1132.
- Bartels, Larry M. 1993. Messages received: the political impact of media exposure. *American Political Science Review* 87(2): 267-285.
- Brayton, Ed. 2016. Please stop sharing links to these sites. <http://www.patheos.com/blogs/dispatches/2016/09/18/please-stop-sharing-links-to-these-sites/>.
- Bullock, John G., Alan S. Gerber, Seth J. Hill, and Gregory A. Huber. 2015. Partisan bias in factual beliefs about politics. *Quarterly Journal of Political Science* 10(4): 519-578.
- BuzzFeed News. Election content engagement. <https://docs.google.com/spreadsheets/d/1ysnzawW6pDGBEqbXqeYuzWa7Rx2mQUIp6CXUUUk4jIk/edit#gid=1756764129>.
- BuzzSumo. <http://buzzsumo.com/>.
- Commager, Henry Steele. 1960. Washington would have lost a TV debate. *New York Times Magazine* VI-13(Oct 30): 79-80.
- ComScore. <http://www.comscore.com/>.
- DellaVigna, Stefano, and Ethan Kaplan. 2007. The Fox News effect: media bias and voting. *The Quarterly Journal of Economics* 122(3): 1187-1234.
- DellaVigna, Stefano, and Matthew Gentzkow. 2010. Persuasion: empirical evidence. *Annual Review of Economics* 2(1): 643-669.
- Dewey, Caitlin. 2016. Facebook fake-news writer: “I think Donald Trump is in the White House because of me.” <https://www.washingtonpost.com/news/the-intersect/wp/2016/11/17/facebook-fake-news-writer-i-think-donald-trump-is-in-the-white-house-because-of-me/>.
- Enikolopov, Ruben, Maria Petrova, and Ekaterina Zhuravskaya. 2011. Media and political persuasion: evidence from Russia. *The American Economic Review* 101(7): 3253-3285.
- Federal Register. 2016. Estimates of the voting age population for 2015. <https://www.federalregister.gov/documents/2016/02/04/2016-02019/estimates-of-the-voting-age-population-for-2015>.
- Flaxman, Seth, Sharad Goel, and Justin Rao. 2016. Filter bubbles, echo chambers, and online news consumption. *Public Opinion Quarterly* 80(1): 298-320.
- Gentzkow, Matthew, and Jesse M. Shapiro. 2011. Ideological segregation online and offline. *The*

- Quarterly Journal of Economics 126(4): 1799-1839.
- Gerber, Alan S., and Donald P. Green. 2000. The effects of canvassing, telephone calls, and direct mail on voter turnout: a field experiment. *American Political Science Review* 94(3): 653-663.
- Gerber, Alan S., James G. Gimpel, Donald P. Green, and Daron R. Shaw. 2011. How large and long-lasting are the persuasive effects of televised campaign ads? Results from a randomized field experiment. *American Political Science Review* 105(1): 135-150.
- Gordon, Brett R., and Wesley R. Hartmann. 2013. Advertising effects in presidential elections. *Marketing Science* 32(1): 13-35.
- Guess, Andrew M. 2015. Measure for measure: an experimental test of online political media exposure. *Political Analysis* 23(1): 59-75.
- Hainmueller, Jens. 2012. Entropy balancing for causal effects: a multivariate reweighting method to produce balanced samples in observational studies. *Political Analysis* 20(1): 25-46.
- Kaplan, Richard L. 2002. *Politics and the American Press: The Rise of Objectivity, 1865-1920*. Cambridge University Press.
- Kessler, Glenn. 2016. No, Donald Trump, there's nothing 'fishy' about Vince Foster's suicide. https://www.washingtonpost.com/news/fact-checker/wp/2016/05/25/no-donald-trump-theres-nothing-fishy-about-vince-fosters-suicide/?utm_term=.4a9b7f611134.
- Lang, Kurt and Gladys Engel Lang. 2002. *Television and Politics*. Transaction Publishers.
- Malka, Ariel, Jon A. Krosnick, and Gary Langer. 2009. The association of knowledge with concern about global warming: trusted information sources shape public thinking. *Risk Analysis* 29(5): 633-647.
- Martin, Gregory J., and Ali Yurukoglu. 2016. Bias in cable news: persuasion and polarization. Working paper.
- Napoli, Philip M. 2014. Measuring media impact: an overview of the field. <https://learcenter.org/pdf/measuringmedia.pdf>.
- Nyhan, Brendan, Jason Reifler, and Peter A. Ubel. 2013. The hazards of correcting myths about health care reform. *Medical Care* 51(2): 127-132.
- Nyhan, Brendan, Jason Reifler, Sean Richey, and Gary L. Freed. 2014. Effective messages in vaccine promotion: a randomized trial. *Pediatrics* 133(4): 835-842.
- Parkinson, Hannah Jane. 2016. Click and elect: how fake news helped Donald Trump win a real election. <https://www.theguardian.com/commentisfree/2016/nov/14/fake-news-donald-trump-election-alt-right-social-media-tech-companies>.
- Pariser, Eli. 2011. *The Filter Bubble: What the Internet is Hiding from You*. Penguin UK.
- Pew Center. 2016a. News use across social media platforms 2016. <http://www.journalism.org/2016/05/26/news-use-across-social-media-platforms-2016>.
- Pew Center. 2016b. News use across social media platforms 2016. <http://assets.pewresearch.org/w>

p-content/uploads/sites/13/2016/05/PJ_2016.05.26_social-media-and-news_FINAL-1.pdf.

PolitiFact. <http://www.politifact.com/truth-o-meter/elections/2016/president-united-states/>.

Price, Vincent, and John Zaller. 1993. Who gets the news? Alternative measures of news reception and their implications for research. *Public Opinion Quarterly* 57(2): 133-164.

Prior, Markus. 2009. The immensely inflated news audience: assessing bias in self-reported news exposure. *Public Opinion Quarterly* 73(1): 130-143.

Prior, Markus, Gaurav Sood, and Kabir Khanna. 2015. You cannot be serious: the impact of accuracy incentives on partisan bias in reports of economic perceptions. *Quarterly Journal of Political Science* 10(4): 489-518.

Read, Max. 2016. Donald Trump won because of Facebook. <http://nymag.com/selectall/2016/11/donald-trump-won-because-of-facebook.html>.

Silverman, Craig. 2016. This analysis shows how fake election news stories outperformed real news on Facebook. <https://www.buzzfeed.com/craigsilverman/viral-fake-election-news-outperformed-real-news-on-facebook>.

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

Spenkuch, Jörg L., and David Toniatti. 2016. Political advertising and election outcomes. Working paper.

Sunstein, Cass R. 2001a. *Echo chambers: Bush v. Gore, impeachment, and beyond*. Princeton University Press.

Sunstein, Cass R. 2001b. *Republic.com*. Princeton University Press.

Sunstein, Cass R. 2009. *Republic.com 2.0*. Princeton University Press.

Zimdars, Melissa. 2016. False, misleading, clickbait-y, and satirical “news” sources. <http://d279m997dpfwgl.cloudfront.net/wp/2016/11/Resource-False-Misleading-Clickbait-y-and-Satirical-%E2%80%9CNews%E2%80%9D-Sources-1.pdf>.

Table 1: News headlines used in the post-election survey

(1) Article text	(2) True/false	(3) Article favors
<u>“Big Fake” news headlines covered in New York Times, Wall Street Journal, and BuzzFeed after the election</u>		
Pope Francis endorsed Donald Trump.	FALSE	Trump
An FBI agent connected to Hillary Clinton’s email disclosures murdered his wife and shot himself.	FALSE	Trump
The Clinton Foundation bought \$137 million in illegal arms.	FALSE	Trump
Mike Pence said that “Michelle Obama is the most vulgar First Lady we’ve ever had.”	FALSE	Clinton
In May 2016, Ireland announced that it was officially accepting Americans requesting political asylum from a Donald Trump presidency.	FALSE	Clinton
Celebrity RuPaul said that Donald Trump mistook him for a woman and groped him at a party in 1995.	FALSE	Clinton
<u>“Small Fake” and “Small True” headlines from PolitiFact</u>		
At the beginning of November, the FBI uncovered evidence of a pedophile sex ring run under the guise of the Clinton Foundation.	FALSE	Trump
Under Donald Trump’s tax plan, it is projected that 51% of single parents would see their taxes go up.	TRUE	Clinton
At a rally a few days before the election, President Obama screamed at a protester who supported Donald Trump.	FALSE	Trump
FBI Director James Comey’s October 28th letter about new developments in the investigation of Hillary Clinton’s emails went only to Republican members of Congress, and not to Democrats.	FALSE	Clinton
A Republican congressman helped broker a deal for Donald Trump to buy a taxpayer-owned building in order to build the Trump International Hotel in Washington, D.C.	FALSE	Clinton
Repeated requests for additional security in Benghazi were routinely denied by Hillary Clinton’s State Department.	TRUE	Trump
<u>“Small Fake” and “Small True” headlines from Snopes, Hillary Clinton tag</u>		
The Clinton campaign secretly paid musicians Beyonce and Jay Z \$62 million to appear at a rally in support of Hillary Clinton.	FALSE	Trump
Hillary Clinton’s first name was spelled with an extra “i” (“Hilliary,” with the word “liar” in the middle) on election ballots printed for use in Lonoke County, Arkansas.	TRUE	Clinton
An email written by Hillary Clinton aide Huma Abedin to her brother revealed that she is a radical Muslim.	FALSE	Trump
<u>“Small Fake” and “Small True” headlines from Snopes, Donald Trump tag</u>		
Donald Trump threatened to deport Puerto Rican Broadway star Lin-Manuel Miranda, not realizing that Puerto Rico is a U.S. territory and Puerto Ricans are U.S. citizens.	FALSE	Clinton
Wikileaks was caught by Newsweek fabricating emails with the intent of damaging Hillary Clinton’s campaign.	FALSE	Clinton
Donald Trump and his campaign donated food and supplies to Hurricane Matthew victims in North Carolina.	TRUE	Trump
<u>“Placebo” headlines that we invented</u>		
Leaked documents reveal that the Clinton campaign planned a scheme to offer to drive Republican voters to the polls but then take them to the wrong place.	FALSE	Trump
Leaked documents reveal that the Trump campaign planned a scheme to offer to drive Democratic voters to the polls but then take them to the wrong place.	FALSE	Clinton
FBI Director James Comey was secretly communicating with Hillary Clinton about when to release results of the FBI investigation into Clinton’s private email server.	FALSE	Trump
FBI Director James Comey was secretly communicating with Donald Trump about when to release results of the FBI investigation into Clinton’s private email server.	FALSE	Clinton
Clinton Foundation staff were found guilty of diverting funds to buy alcohol for expensive parties in the Caribbean.	FALSE	Trump
Trump Foundation staff were found guilty of diverting funds to buy alcohol for expensive parties in the Caribbean.	FALSE	Clinton
<u>“Big True” headlines from the Guardian’s election timeline</u>		
Hillary Clinton said that “you could put half of Trump’s supporters into what I call the basket of deplorables.”	TRUE	Trump
At the 9/11 memorial ceremony, Hillary Clinton stumbled and had to be helped into a van.	TRUE	Trump
At the third presidential debate, Donald Trump refused to say whether he would concede the election if he lost.	TRUE	Clinton
On October 28th, the FBI director alerted members of Congress that it had discovered new emails relevant to its investigation of Hillary Clinton’s personal server.	TRUE	Trump
The musicians Beyonce and Jay Z appeared at a rally in support of Hillary Clinton.	TRUE	Clinton
Two days before the election, the FBI director told Congress that a newer batch of emails linked to Hillary Clinton’s private email server did not change his conclusion that Clinton should face no charges over her handling of classified information.	TRUE	Clinton

Notes: This table presents the 30 news articles used in the post-election survey. Each respondent received a randomly selected 15 of these stories, stratified to receive three from each of the five major categories listed.

Table 2: Post-election survey summary statistics

	(1) Survey sample	(2) U.S. adult population
Household income (000s)	72.73	76.16
College graduate	0.44	0.27
High school or less	0.27	0.42
Male	0.35	0.49
Age	45.88	47.15
Caucasian	0.79	0.62
Democrat	0.35	0.37
Republican	0.24	0.29
Web news consumption frequency	2.34	1.58
Social media news consumption frequency	1.88	1.24

Notes: This table presents demographic data and summary statistics for the post-election survey and the U.S. adult population. News consumption frequency is coded as 3 (often), 2 (sometimes), 1 (rarely), and 0 (never). National average income, education, gender, age, and race are from the U.S. Census and are relevant for the U.S. population aged 18 and over. National party affiliation data are from the American National Election Studies' 2012 Time Series Study. National news consumption frequencies are from the Pew Center (2016b).

Table 3: Rates of seeing and believing fake news relative to placebo fake news

	(1)	(2)	(3)	(4)	(5)	(6)
	Fake	Recall seeing		Recall seeing and believed		
		Placebo	Fake-Placebo	Fake	Placebo	Fake-Placebo
Share of population	0.153*** (0.009)	0.141*** (0.011)	0.012 (0.009)	0.079*** (0.007)	0.083*** (0.009)	-0.005 (0.007)
N	8,456	3,624	12,080	8,456	3,624	12,080
95 pct confidence bound	.171	.1632	.0288	.0924	.1012	.009

Notes: This table presents the share people who recall seeing (columns 1-3) or recall seeing and believed (columns 4-6) news headlines. Columns 1 and 4 include only Fake headlines, columns 2 and 5 include only Placebo headlines, and columns 3 and 6 present differences between the previous two columns. Observations are weighted for national representativeness. Standard errors are robust and clustered by survey respondent. *, **, ***: statistically significant from zero with 90, 95, and 99 percent confidence, respectively.

Table 4: Associations of hearing and believing fake news with observables

	(1)	(2)	(3)	(4)	(5)
		Recall seeing		Believed	
	Fake	Placebo	Fake-Placebo	Fake	Placebo
Social most important	0.093*** (0.033)	0.074* (0.039)	0.019 (0.030)	0.085*** (0.031)	0.100** (0.042)
Use social media	0.031 (0.021)	0.050* (0.026)	-0.018 (0.022)	-0.008 (0.021)	0.004 (0.031)
Segregated network	0.036* (0.022)	0.002 (0.027)	0.034 (0.021)	0.032 (0.021)	-0.020 (0.025)
College graduate	0.021 (0.019)	0.018 (0.021)	0.002 (0.020)	0.001 (0.017)	0.030 (0.022)
High school or less	0.021 (0.021)	0.056** (0.025)	-0.035 (0.022)	0.038* (0.020)	0.045* (0.027)
Democrat	-0.049*** (0.018)	-0.024 (0.023)	-0.024 (0.020)	0.010 (0.018)	-0.023 (0.023)
Republican	0.017 (0.025)	0.023 (0.029)	-0.006 (0.022)	0.075*** (0.023)	0.076** (0.032)
Undecided	0.002 (0.018)	0.032 (0.024)	-0.030 (0.020)	0.005 (0.017)	-0.040* (0.023)
N	8,456	3,624	12,080	8,456	3,624

Notes: This table presents estimates of equation (4), a linear probability model regression of indicator variables for recall (columns 1-3) or belief (columns 4 and 5) of news headlines on individual characteristics. Columns 1 and 4 include only Fake headlines, while columns 2 and 5 include only Placebo headlines. Column 3 presents differences between the previous two columns. All columns include additional demographic controls: income, age, race, and gender. “Segregated network” is an indicator variable for whether the respondent uses social media and reports that at least 75 percent of social media friends preferred the same political candidate. “Undecided” is an indicator variable for whether the respondent decided which candidate to vote for less than three months before the election. Observations are weighted for national representativeness. Standard errors are robust and clustered by survey respondent. *, **, ***: statistically significant from zero with 90, 95, and 99 percent confidence, respectively.

Table 5: Ideological alignment and perceptions of fake news

	(1)	(2)	(3)	(4)	(5)
	Fake	Recall seeing		Believed	
		Placebo	Fake-Placebo	Fake	Placebo
Pro-Trump × Republican	0.213*** (0.024)	0.233*** (0.035)	-0.020 (0.031)	0.251*** (0.026)	0.329*** (0.043)
Pro-Trump × Independent	0.165*** (0.016)	0.154*** (0.022)	0.010 (0.019)	0.108*** (0.013)	0.171*** (0.024)
Pro-Trump × Democrat	0.105*** (0.014)	0.080*** (0.016)	0.025 (0.019)	0.060*** (0.011)	0.043*** (0.010)
Pro-Clinton × Republican	0.146*** (0.024)	0.096*** (0.021)	0.050*** (0.019)	0.168*** (0.027)	0.106*** (0.030)
Pro-Clinton × Independent	0.159*** (0.017)	0.136*** (0.021)	0.023 (0.020)	0.165*** (0.015)	0.134*** (0.018)
Pro-Clinton × Democrat	0.143*** (0.016)	0.159*** (0.024)	-0.016 (0.022)	0.246*** (0.025)	0.214*** (0.026)
N	8,456	3,624	12,080	8,456	3,624

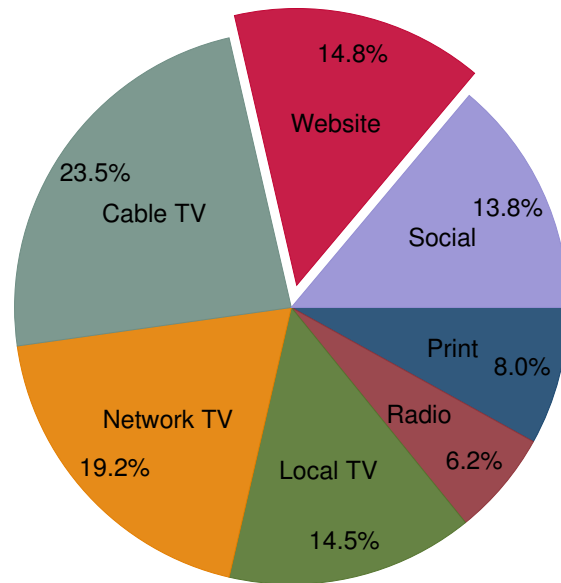
Notes: This table presents linear probability model regressions of indicator variables for recall (columns 1-3) or belief (columns 4 and 5) of news headlines on the interaction of political affiliation indicators and pro-Clinton or pro-Trump headline indicators. There is no omitted category, so the coefficients represent shares of the population. Columns 1 and 4 include only Fake headlines, while columns 2 and 5 include only Placebo headlines. Column 3 presents differences between the previous two columns. Observations are weighted for national representativeness. Standard errors are robust and clustered by survey respondent. *, **, ***: statistically significant from zero with 90, 95, and 99 percent confidence, respectively.

Table 6: Possible impact of fake news on the 2016 election

Row	Scenario	(1) Net pro-Trump exposure rate (e) (articles/person)	(2) Required persuasion rate (f^*)	(3) Required multiplier f^*/f_{ad}
1	Primary: Survey exposure rates	0.69	0.0073	36
2	75% of stories remembered	0.91	0.0055	27
3	50% of stories remembered	1.37	0.0036	18
4	50% more fake news shares	1.03	0.0049	24
5	100% more fake news shares	1.83	0.0027	14
6	Use survey “saw and believed”	0.53	0.0094	47
7	Exposure = 10x shares from database	0.93	0.0054	27
8	Exposure = 20x shares from database	1.85	0.0027	13
9	Heterogeneous exposure by type	$e_C \approx 0.1, e_N \approx 0.69$	0.0132	66

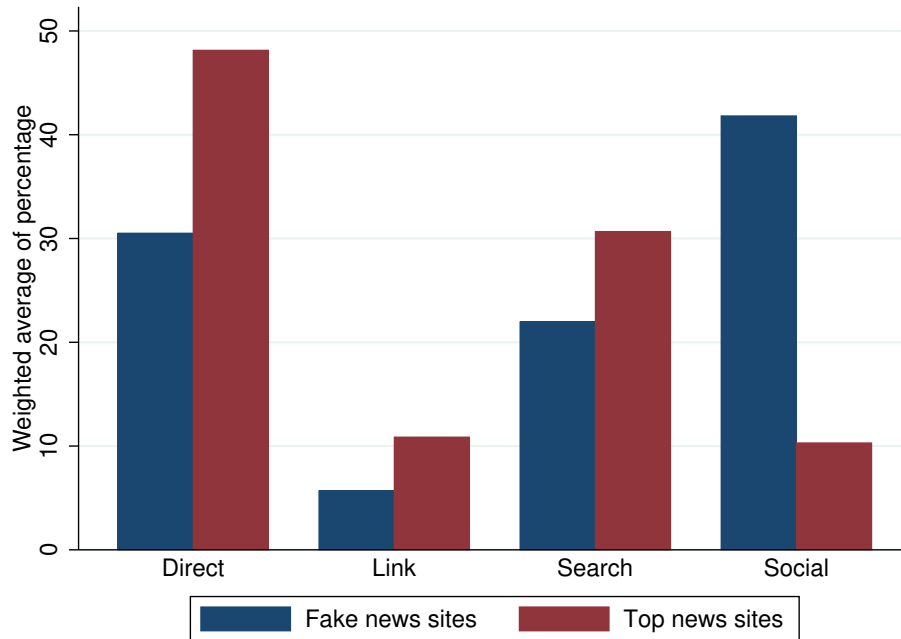
Notes: This table benchmarks the possible effects of fake news on Trump’s margin of victory in the 2016 election. Column 2 is the required persuasion rate for fake news to have changed the election outcome, as calculated from equation (8). Column 3 is the required multiplier by which fake news articles must be more persuasive than television campaign ads ($f_{ad} = 0.0002$) for fake news to have changed the election outcome.

Figure 1: Most important source of 2016 election news



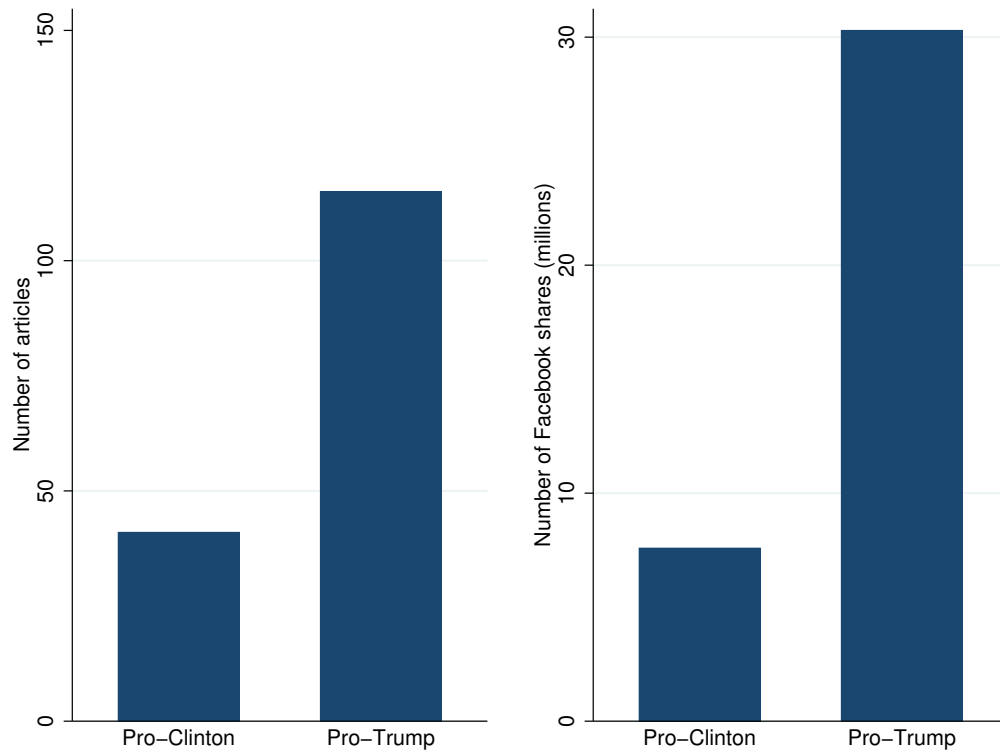
Notes: Our post-election survey asked, “which of these sources was your most important source of news and information about the 2016 election?” This figure plots responses. Observations are weighted for national representativeness.

Figure 2: Share of visits to U.S. news websites by source



Notes: This figure presents the share of traffic from different sources for the top 690 U.S. news websites and for 65 fake news websites. Sites are weighted by number of monthly visits. Data are from Alexa.

Figure 3: Number of articles and Facebook shares in our fake news database



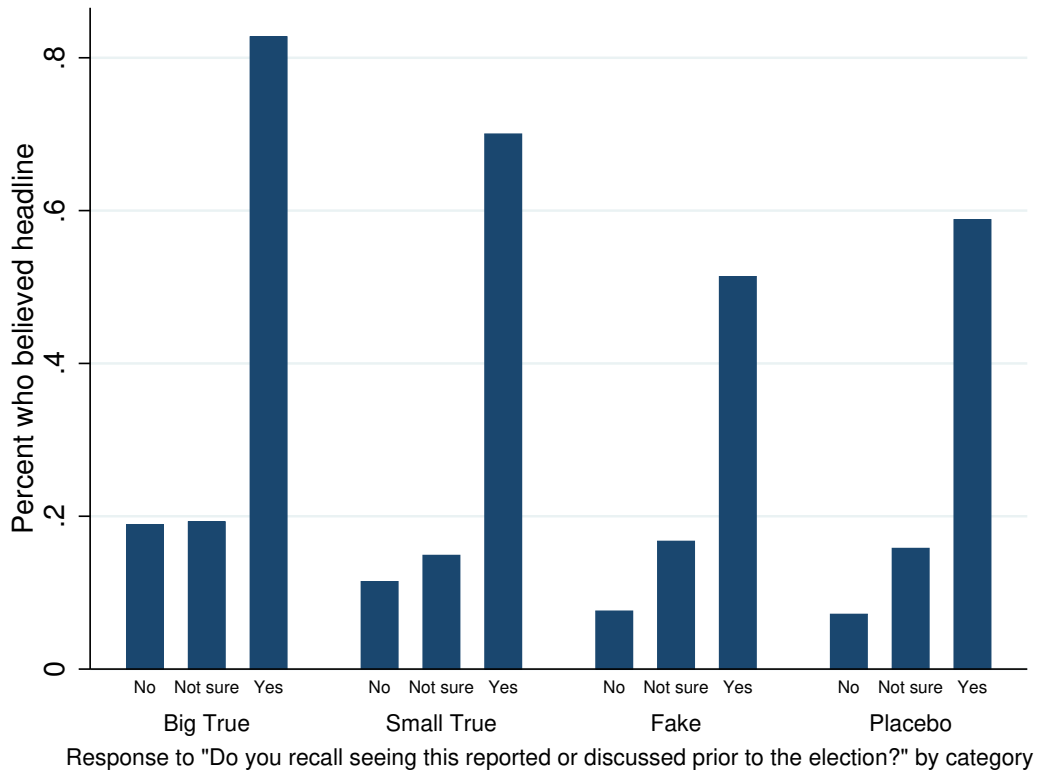
Notes: This figure presents the total number of articles and total number of Facebook shares for pro-Trump and pro-Clinton articles in our fake news database. The database comprises all election-related fake news articles that appeared in the three months before the election on Snopes, PolitiFact, and BuzzFeed; see Section 4 for details.

Figure 4: Percent of U.S. adult population that heard or believed election news, by category



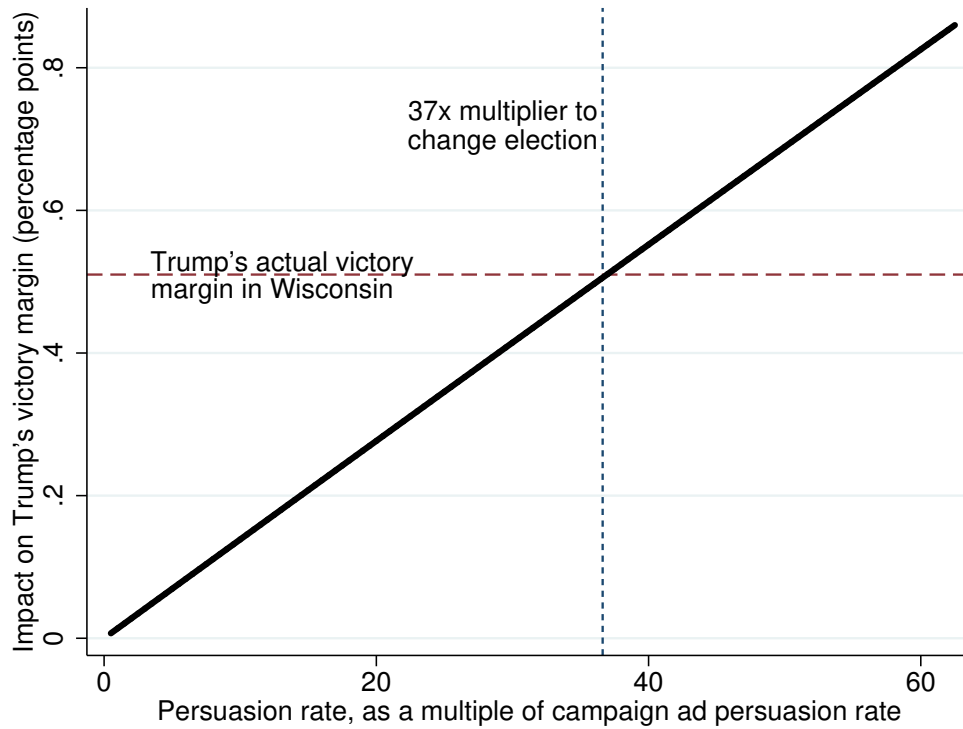
Notes: In our post-election survey, we presented 15 headlines. For each headline, the survey asked whether respondents had heard the headline (“Do you recall seeing this reported or discussed before the election?”) and whether they believed it (“At the time of the election, would your best guess have been that this statement was true?”). The left bars present the share of respondents who had heard the headlines in each category, and the right bars present the share of respondents who heard and believed the headlines. Observations are weighted for national representativeness.

Figure 5: Share who believe news by whether they heard news, by category



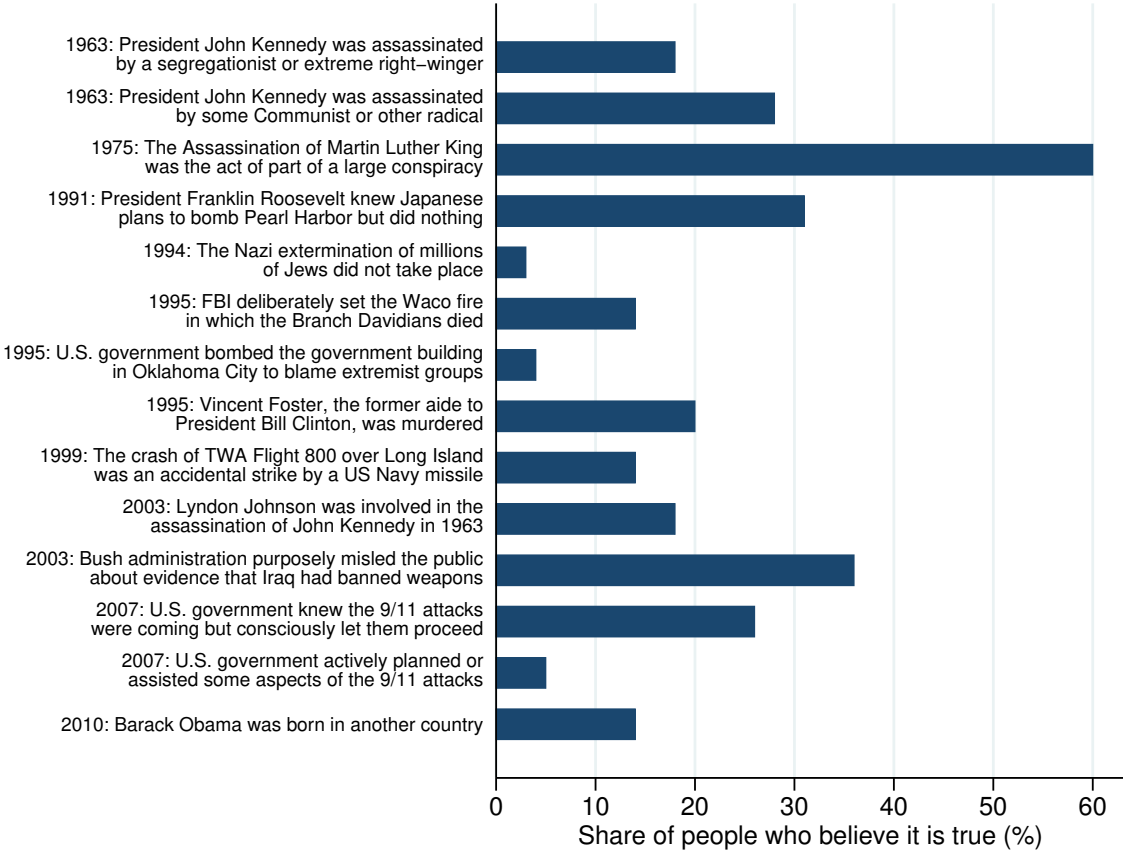
Notes: In our post-election survey, we presented 15 headlines. For each headline, the survey asked whether respondents had heard the headline (“Do you recall seeing this reported or discussed before the election?”) and whether they believed it (“At the time of the election, would your best guess have been that this statement was true?”). This figure presents the share of people who believed the headlines in each category, broken down by responses to whether they had heard each headline. Observations are weighted for national representativeness.

Figure 6: Benchmarking the impact of fake news on the 2016 presidential election



Notes: The x-axis is the persuasion rate per fake news article, as a multiple of the persuasion rate of political campaign ads, $f_{ad} = 0.0002$. The upward-sloping line on this figure is calculated by applying equation (7) to the persuasion rate implied on the x-axis, assuming that Americans were exposed to $\hat{e} \approx 0.69$ more pro-Trump than pro-Clinton fake news articles. Trump's actual victory margin was approximately 0.51 percent: the amount by which his vote margins in all states would have to decrease to shift Michigan, Pennsylvania, and Wisconsin to Clinton victories and therefore change the outcome of the Electoral College vote.

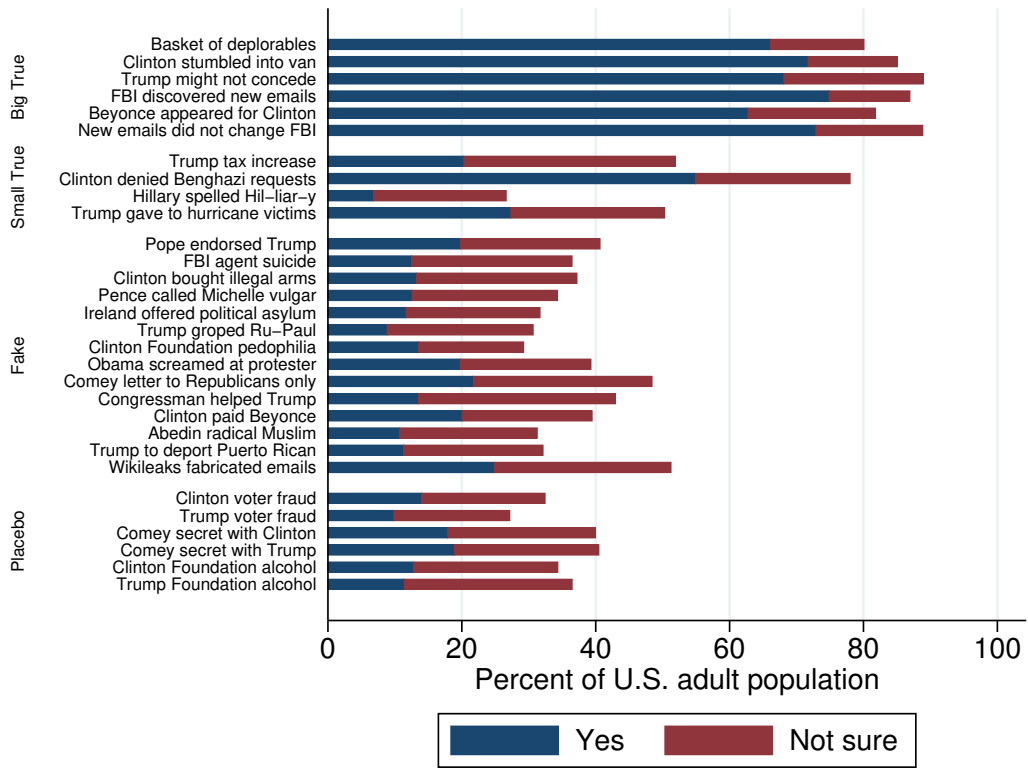
Figure 7: Share of Americans believing historical partisan conspiracy theories



Notes: From polling data compiled by the American Enterprise Institute (2013), we selected all conspiracy theories with political implications. This figure plots the share of people who report believing the statement listed, using opinion polls from the date listed.

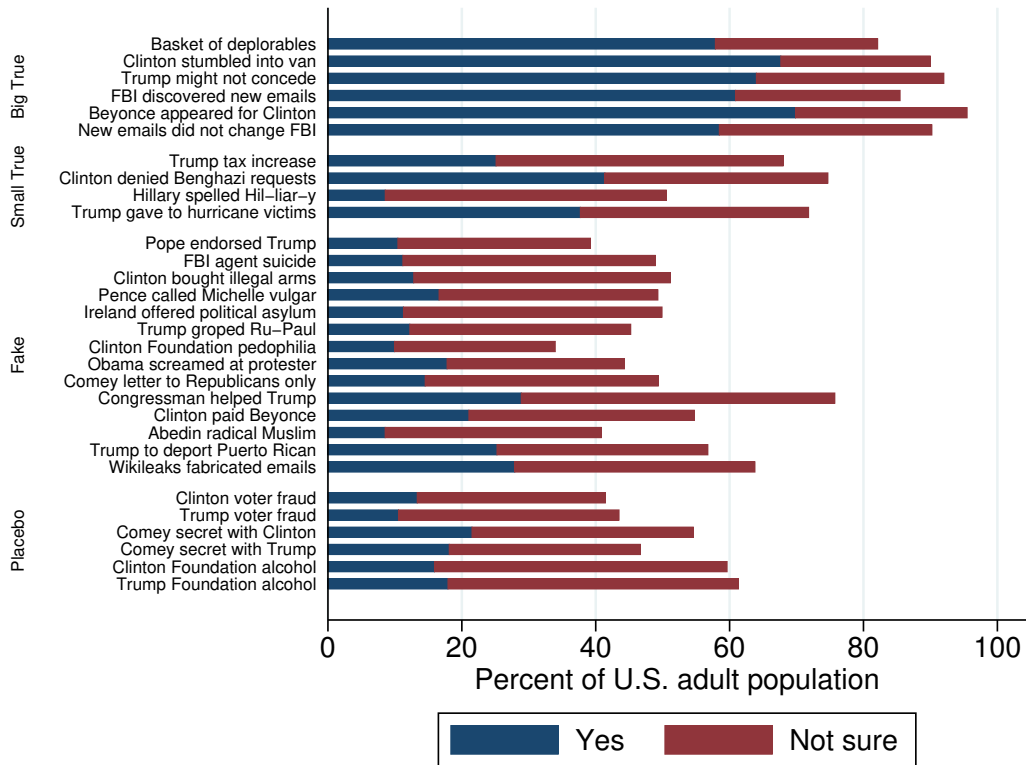
A Appendix

Appendix Figure 1: Percent of U.S. adult population that recalled seeing election news, by article



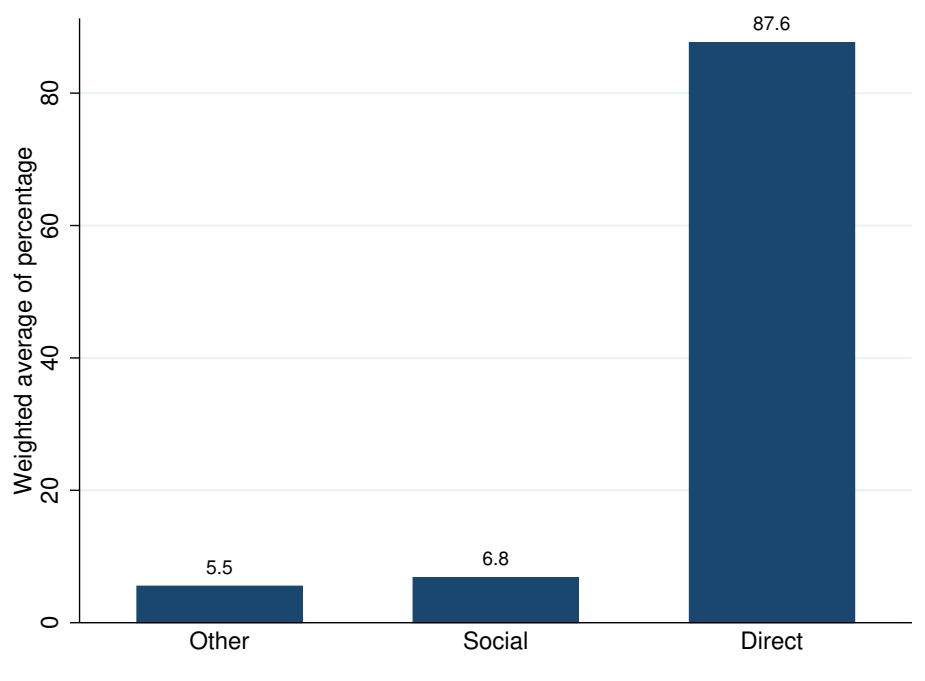
Notes: This figure presents the share of respondents that responded “Yes” and “Not sure” to the question, “Do you recall seeing this reported or discussed before the election,” for each of the 30 headlines listed in table 1. The headline categories written vertically are as defined in Section 2.3. Observations are weighted for national representativeness.

Appendix Figure 2: Percent of U.S. adult population that believed election news, by article



Notes: This figure presents the share of respondents that responded “Yes” and “Not sure” to the question, “At the time of the election, would your best guess have been that this statement was true?” for each of the 30 headlines listed in table 1. The headline categories written vertically are as defined in Section 2.3. Observations are weighted for national representativeness.

Appendix Figure 3: Share of visits to U.S. news websites by source, using comScore data



Notes: This figure presents the share of traffic from different sources for the top 50 U.S. news websites. Sites are weighted by number of monthly visits. Data are from comScore.

Appendix Table 1: Ideological alignment and perceptions of fake news among people with ideologically segregated social networks

	(1)	(2)	(3)	(4)	(5)
		Recall seeing		Believed	
	Fake	Placebo	Fake-Placebo	Fake	Placebo
Pro-Trump × Republican	0.298*** (0.058)	0.290*** (0.079)	0.007 (0.056)	0.364*** (0.069)	0.276*** (0.076)
Pro-Trump × Independent	0.142*** (0.026)	0.062** (0.026)	0.080** (0.035)	0.087*** (0.023)	0.081*** (0.029)
Pro-Trump × Democrat	0.154*** (0.026)	0.102*** (0.027)	0.052 (0.038)	0.075*** (0.020)	0.054*** (0.017)
Pro-Clinton × Republican	0.216*** (0.067)	0.148*** (0.057)	0.068* (0.039)	0.170*** (0.065)	0.121** (0.053)
Pro-Clinton × Independent	0.131*** (0.025)	0.123*** (0.033)	0.008 (0.032)	0.139*** (0.026)	0.173*** (0.035)
Pro-Clinton × Democrat	0.180*** (0.028)	0.189*** (0.035)	-0.008 (0.043)	0.300*** (0.030)	0.251*** (0.040)
N	2,548	1,092	3,640	2,548	1,092

Notes: This table presents linear probability model regressions of indicator variables for recall (columns 1-3) or belief (columns 4 and 5) of news headlines on the interaction of political affiliation indicators and pro-Clinton or pro-Trump headline indicators. There is no omitted category, so the coefficients represent shares of the population. This table parallels table 5, except limiting the sample to people who report having “segregated” networks, i.e. that 75 percent or more of their friends on social media preferred the same presidential candidate. Columns 1 and 4 include only Fake headlines, while columns 2 and 5 include only Placebo headlines. Column 3 presents differences between the previous two columns. Observations are weighted for national representativeness. Standard errors are robust and clustered by survey respondent. *, **, ***: statistically significant from zero with 90, 95, and 99 percent confidence, respectively.