



Research Article

“Fake news” may have limited effects beyond increasing beliefs in false claims

Since 2016, there has been an explosion of interest in misinformation and its role in elections. Research by news outlets, government agencies, and academics alike has shown that millions of Americans have been exposed to dubious political news online. However, relatively little research has focused on documenting the effects of consuming this content. Our results suggest that many claims about the effects of exposure to false news may be overstated, or, at the very least, misunderstood.

Author: Andrew M. Guess (1), Dominique Lockett (2), Benjamin Lyons (3), Jacob M. Montgomery (4), Brendan Nyhan (5), and Jason Reifler (6)

Affiliation: (1) Department of Politics and Woodrow Wilson School, Princeton University; (2,4) Department of Political Science, Washington University in St. Louis; (3) Department of Communication, University of Utah; (5) Department of Government, Dartmouth College, (6) Department of Politics, University of Exeter

How to cite: Guess, Andrew M.; Lockett, Dominique; Lyons, Benjamin; Montgomery, Jacob M.; Nyhan, Brendan; Reifler, Jason (2020). “Fake news” may have limited effects beyond increasing beliefs in false claims, *The Harvard Kennedy School (HKS) Misinformation Review*, Volume 1, Issue 1

Received: Nov. 11, 2019 Accepted: Dec. 20, 2019 Published: Jan. 14, 2020

Research question

Does news from untrustworthy websites (“fake news”) influence political attitudes and beliefs? What can observational and experimental data tell us about the effects of exposure to this type of content?

Essay summary

- We examine the connection between the websites people visit or are exposed to and their answers to questions in public opinion surveys about politics.
- We analyze these relationships using two types of data: real-world consumption of news from untrustworthy websites and randomized exposure to an article from an untrustworthy website as part of an opinion survey. (We use “consume”/“consumption” and “expose”/“exposure,” respectively, to distinguish these approaches.)
- Our data come from three nationally representative samples of Americans that we refer to as the summer 2018, fall 2018, and winter 2018 data. These studies included survey questions; observational web traffic data from respondents, which was collected in anonymized form with informed consent; and controlled experiments in which some respondents were shown an article from an untrustworthy website.

¹ A publication of the Shorenstein Center for Media, Politics, and Public Policy at Harvard University, John F. Kennedy School of Government.

- We identify real-world consumption of “untrustworthy websites” by comparing their web browsing history in the 1–3 weeks before the survey to a comprehensive list compiled by Grinberg et al. (2019) of sites that frequently publish false or unverified information. This list includes sites flagged by fact-checkers as publishing false content as well additional websites that were labeled as unreliable by that research team.
- We also randomly showed respondents one of three different articles in the summer and fall 2018 studies — a pro-Democratic false story from an untrustworthy website, a pro-Republican false story from an untrustworthy website, or a neutral (non-political) story — and then examined how exposure to that article affected subsequent survey responses.
- Untrustworthy website consumption or exposure seems to increase misperceptions. Our results indicate that untrustworthy website consumption is associated with greater willingness to endorse politically congenial misperceptions among Republicans, the group for whom the content of those sites is most often congenial. Similarly, randomized exposure to a misleading article from an untrustworthy website increases belief in the article’s central claim, especially for Democrats when the article is congenial.
- However, the effects for other variables are inconsistent. People who consume news from untrustworthy websites tend to have more negative views of the media and more polarized feelings toward the political parties, but they do not differ measurably from other Americans in their intentions to participate in politics. By contrast, controlled exposure to a dubious article increased self-reported intentions to vote, but had no measurable effect on feelings toward the media, the political parties, or intended political activity beyond voting.
- These differences may be the result of the different ways people encounter untrustworthy news in our data. Real-world consumption occurs naturally, while controlled exposure was experimental. Alternatively, these differences may reflect the nature of the two approaches; the experiment shows only one story while consumption may be repeated frequently over a long period of time.
- Though further research is needed, these results suggest that exposure to “fake news” may have limited effects aside from increasing beliefs in false claims.

Implications

The combination of observational and experimental approaches that we employ offers important new insights into the effects of “fake news.” Most notably, both types of data indicate how untrustworthy online content can promote inaccurate factual beliefs. Using our real-world consumption data, we see that people who visit untrustworthy websites are more likely to hold a misperception.

Similarly, controlled exposure to a false article from an untrustworthy website increases belief in its central claim. As a result, we are confident that exposure to misinformation increases misperceptions about specific and relevant factual claims.

By contrast, our findings often conflict when we examine the relationship with other outcomes of interest. For instance, our regression analysis reveals that consuming information from untrustworthy websites is associated with lower media trust, but reading a false article in our experiment did not worsen attitudes towards the media. This result suggests that people may choose to consume factually dubious content *because* they lack trust in the media. We observe a similar discrepancy between our observational and experimental data for political participation. Respondents who visit untrustworthy websites report similar levels of political participation to those who do not visit such websites. Experiments can be a powerful tool to help show causal effects that are unclear from observational data alone. Randomized exposure to a single false article increased respondents’ self-reported likelihood to vote. This

experimental effect is likely short-lived, but underscores the broader point that further research is needed to understand whether and how false content affects political attitudes and behavior other than factual beliefs.

To make further progress, future studies of the effects of “fake news” must confront three additional challenges. First, our measures of untrustworthy website consumption do not capture exposure on social media platforms or mobile devices. As a result, we cannot measure exposure to headlines that people see in a news feed but do not click or website visits that take place on devices such as phones and tablets (the behavioral data we study only includes limited data from mobile devices). Second, it is unclear how much untrustworthy news content should be shown to participants in studies. The effects of exposure to this type of content may only be measurable if we administer a higher dosage than our current studies provide, which in turn raises ethical concerns that researchers should consider carefully. Finally, we should test whether the effects we find in our experiments would be observed among people who would consume untrustworthy websites in the real world. It is possible that Americans who visit these sites do not react as strongly to them as others would.

Until these challenges are addressed, it is difficult to fully assess the effects of “fake news.” Our results suggest, however, that this type of content has limited effects on political attitudes and behavior beyond factual beliefs. The most worrisome aspect of these untrustworthy websites, in other words, is that they often succeed in misleading Americans.

Findings

Each finding below is illustrated using one or more graphs reporting the average difference in the outcome measure between respondents who did not consume information from untrustworthy websites and those who did (for the non-experimental data) and the causal effect of exposure to a pro-Democratic or pro-Republican false story on the outcome measure (for the experimental data). We report differences separately for Democrats and Republicans in some cases.

Finding 1: Consumption of untrustworthy websites is associated with more negative feelings toward the media.

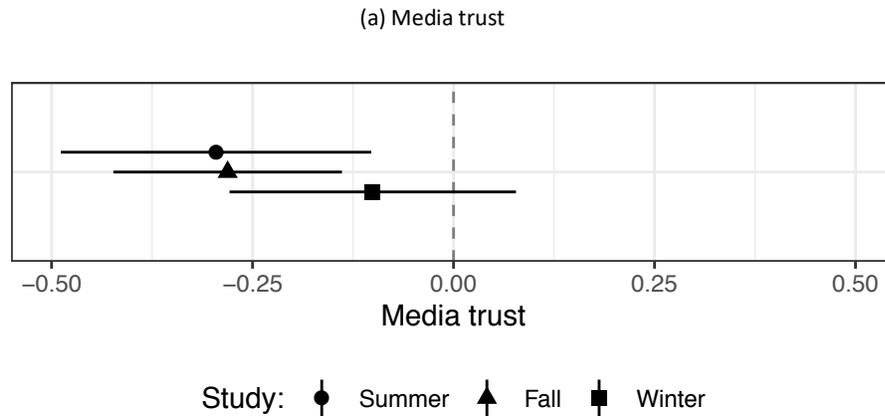
We find that Americans who consume information from untrustworthy websites tend to report less “trust and confidence” in the media and more negative feelings towards it (see Figures 1a and 1b). Previous research has shown mixed evidence about the connection between views toward mainstream news and untrustworthy website consumption. Low media trust can be associated with alternative media consumption (Tsafati and Cappella 2003), though more recent research has called for a re-examination of this relationship (Kalogeropoulos et al. 2019). Similarly, behavioral evidence suggests that consumption of news from untrustworthy websites does not crowd out consumption of mainstream news, which is consistent with evidence that many people who view the media negatively still consume a great deal of mainstream news (Guess et al. 2019 a, b).

Finding 2: Consumption of untrustworthy websites is associated with polarized feelings toward the political parties.

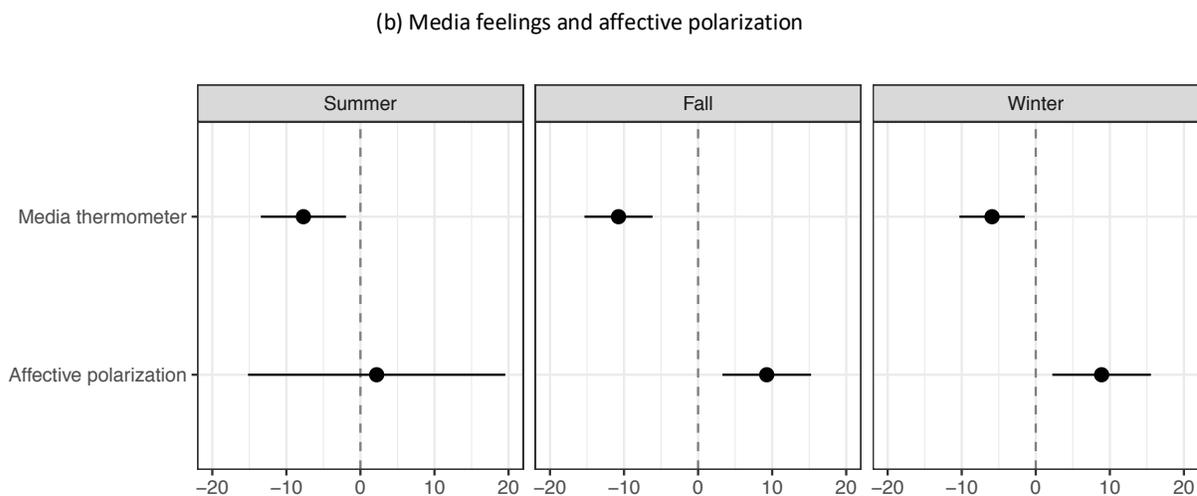
We find evidence that the consumption of untrustworthy websites is associated with more polarized feelings towards the two political parties (see Figure 1b). Previous research examining other forms of media content suggests that exposure to pro-attitudinal information and partisan criticism (the types of content that untrustworthy websites often provide) can inspire more polarized attitudes toward the

parties (Iyengar et al. 2012; Garrett et al. 2014; Suhay et al. 2018). We assess whether the relationship holds for this relatively new type of digital content. In our fall and winter 2018 data, which were collected immediately before and after the midterm elections, we find greater affective polarization (which we measure as the difference in people’s feelings toward their preferred party and the opposition party) among consumers of information from untrustworthy websites.

Figure 1. Association between untrustworthy website consumption and media/party attitudes.



Change in media trust associated with consumption of untrustworthy websites. Point estimates (with 95% confidence intervals) from OLS regressions. Outcome measure is a four-point scale of self-reported trust and confidence in the mass media ranging from “None at all” (1) to “A great deal” (4). Control variables include partisanship, ideology, political knowledge, feelings toward President Trump, age, race, and college education. See Table B1 in Online Appendix B for full results.



Change in media thermometer rating and affective polarization associated with consumption to untrustworthy websites. Point estimates (with 95% confidence intervals) from separate OLS regressions. Outcome measures are respondent feelings toward the media on a 0–100 scale (media thermometer) and the difference in feelings respondents express toward their preferred party and the opposition party (both measured on a 0–100 scale). Control variables include partisanship, ideology, political knowledge, feelings toward President Trump, age, race, and college education. See Tables B1 and B2 in Online Appendix B for full results.

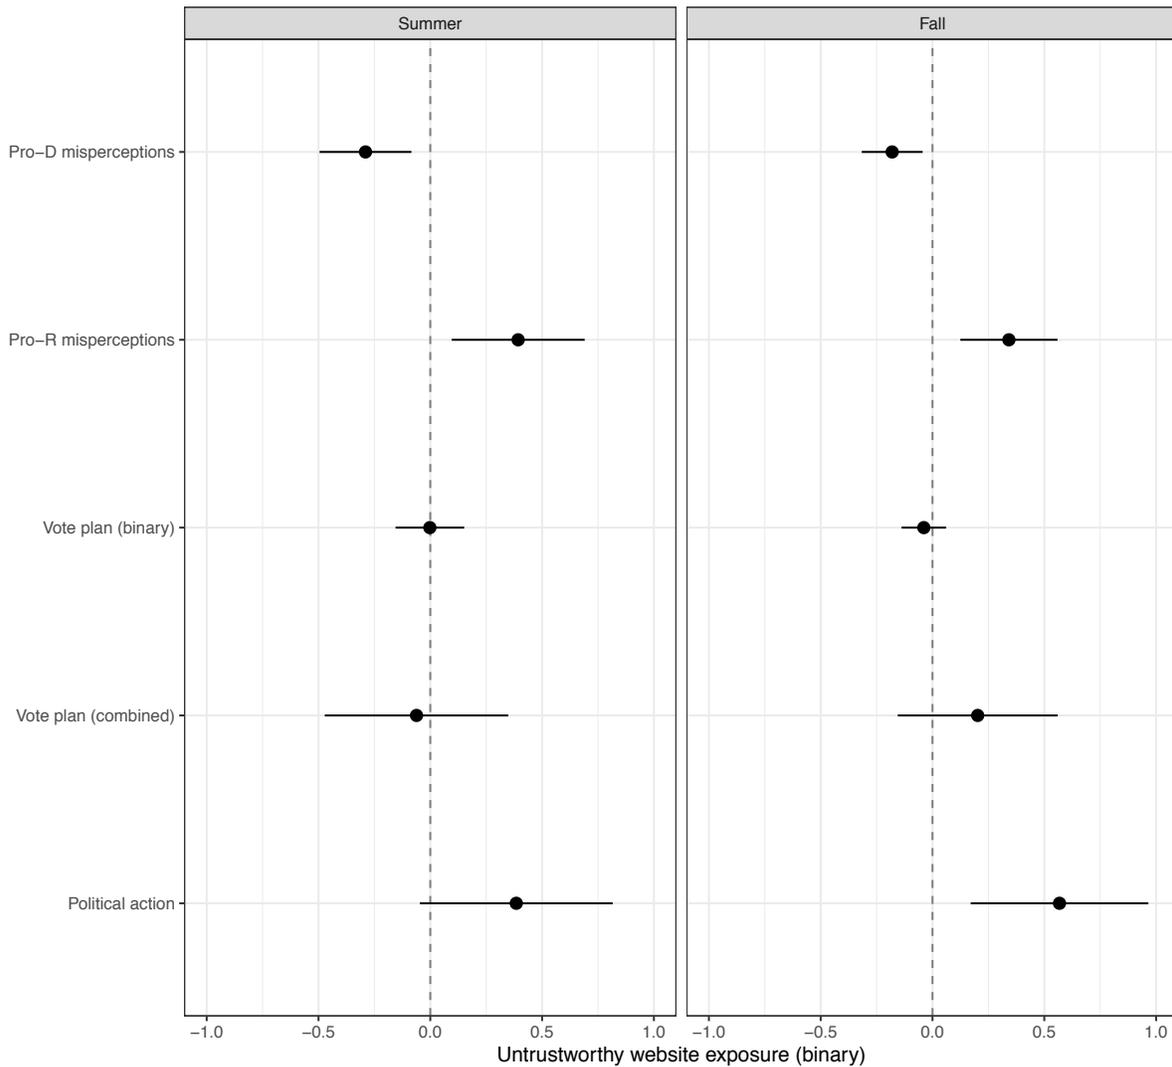
Finding 3: Untrustworthy website consumption is associated with greater belief in pro-GOP misperceptions and reduced belief in pro-Democratic ones.

People who consume information from untrustworthy websites report greater belief in pro-Republican misperceptions and less in pro-Democratic misperceptions even after we account for partisanship, ideology, and feelings toward President Trump (see Figure 2). Previous research found that the untrustworthy websites which attracted the most attention and traffic before the 2016 election had a conservative or pro-Republican orientation (e.g., Benkler et al. 2018; Grinberg et al. 2019; Guess et al. 2019b). We infer that this pattern of slant explains the positive (negative) relationship between visiting untrustworthy websites and pro-Republican (pro-Democratic) misperceptions, though we cannot determine whether this result reflects self-selection or exposure to misinformation from these sites.

Finding 4: The consumption of untrustworthy websites is generally not associated with political participation intentions.

Our data indicate that Americans who consume news from untrustworthy websites generally had similar intentions to participate in politics to those who do not visit such websites (see Figure 2). Given that anger helps to mobilize voters (e.g., Valentino et al. 2011) and that articles on untrustworthy websites are often irate and inflammatory, a frequent question is whether this content increases intent to engage in political activities. However, visiting untrustworthy websites was not associated with intent to vote in the next election. In fall 2018, respondents who visited untrustworthy websites reported significantly greater intent to engage in non-voting forms of political action, but this relationship was not observed in summer 2018. Subsequent research should examine if this the pattern holds in future elections.

Figure 2. Association between untrustworthy website consumption and beliefs/behavior.



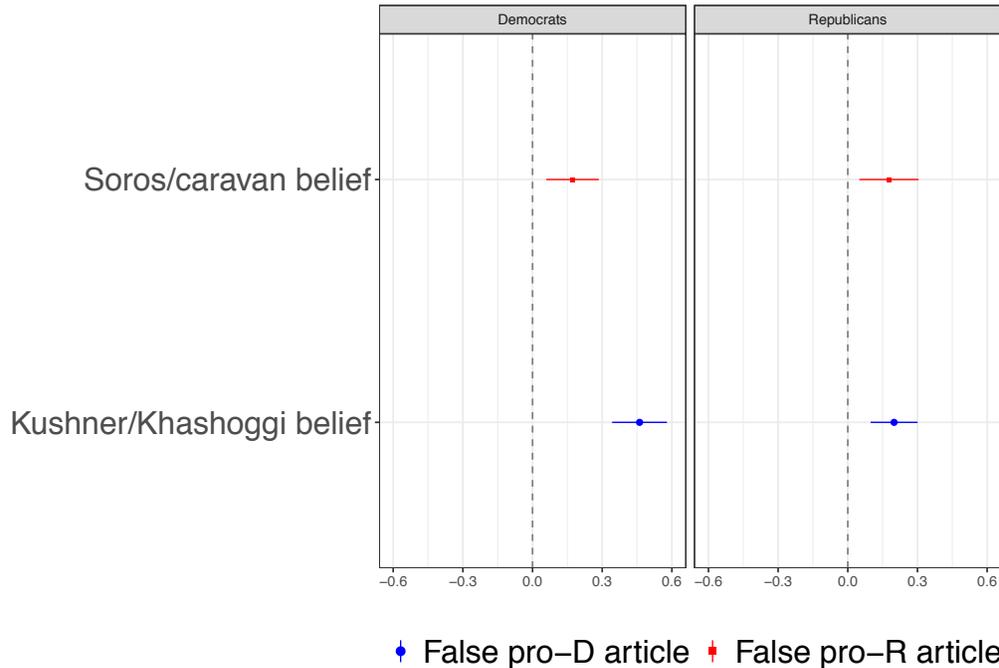
Change in misperceptions, vote intentions, and willingness to take political action associated with consumption of news from untrustworthy websites. Point estimates (with 95% confidence intervals) from separate OLS regressions. Outcome measures are average belief in pro-Democratic or pro-Republican misperceptions on a four-point accuracy scale ranging from “Not at all accurate” (1) to “Very accurate” (4); whether respondents say they are “absolutely certain” to vote or report already having voted (0/1); a five-point measure of vote intention ranging from answering no when asked if they plan to vote (1) to saying they are “absolutely certain” to vote or report already having voted (5); and respondents’ average response when asked if they would be willing to take a series of political actions on a scale from “Disagree strongly” (1) to “Agree strongly” (5). Control variables include partisanship, ideology, political knowledge, feelings toward President Trump, age, race, and college education. See Tables B3 and B4 for full results. Note: Voting and participation regressions are estimated among respondents in the control condition for the false article experiment.

Finding 5: Randomized exposure to a dubious article increases belief in the article’s central claim, especially for Democrats when the claim was congenial.

Respondents who were randomly exposed to an inflammatory but misleading article from an untrustworthy website rated the article’s central claim as more accurate than those who did not read the

article in question (see Figure 3). This effect was magnified for Democrats exposed to a pro-Democratic false story. We ran this experiment as part of the fall 2018 survey, which randomly assigned respondents to one of three articles: a politically neutral story on food, a false story about how George Soros was sponsoring a migrant caravan from Central America (a prominent conspiracy theory at the time), and a false story suggesting that Jared Kushner, Donald Trump’s son-in-law, played a role in the death of Jamal Khashoggi at the hands of Saudi Arabian intelligence agents (a more obscure conspiracy theory). We found that exposure to these stories, which are provided in Online Appendix A, increased respondents’ beliefs in the underlying claims. However, the effect was not symmetric by party. Exposure to the Kushner/Kashoggi article increased belief in its central claim among identifiers with both major parties, but the effect was differentially larger among Democrats compared to Republicans. By contrast, the increase in belief in the Soros/caravan conspiracy claim resulting from exposure to the article in question were similar between parties. These results may reflect differences in pre-treatment awareness of these stories since the Soros/caravan rumor was covered more widely in the period before the study was fielded; prior encounters with information can reduce or change its impact (Druckman and Leeper 2012).²

² We omit the effect of exposure to one false article on respondent belief in the claim expressed in the other false article from the figure. In one of four cases, exposure to a false article decreased belief in the other claim (ratings of the Soros/caravan claim among Republicans who saw the Kushner/Khashoggi article; see Table B5). Given that we only observe one effect of this type, we are not certain if these results are meaningful.

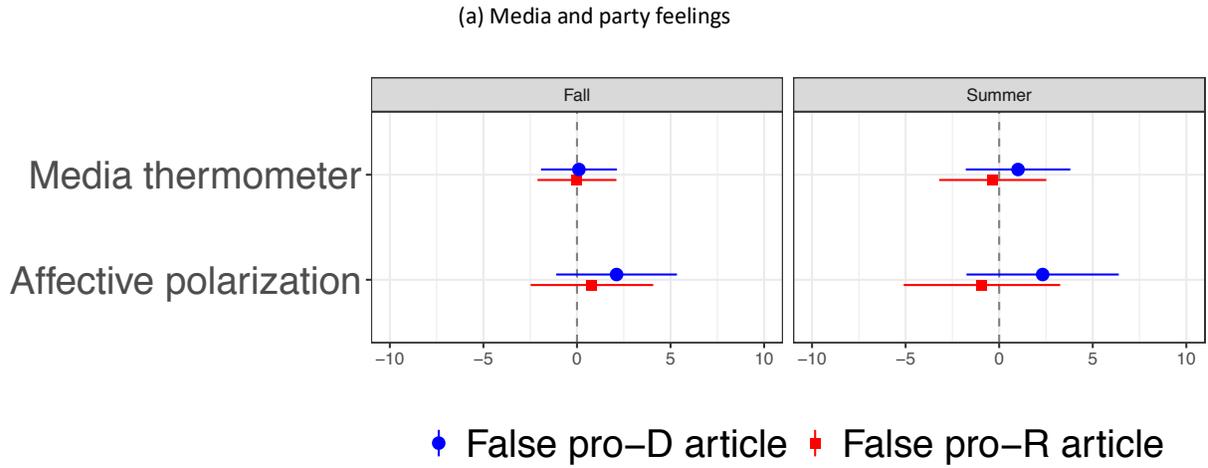
Figure 3. Effect of false article exposure on claim belief.

Effect of exposure to a false article from an untrustworthy website on belief in the claim expressed in the article (fall 2018). Point estimates (with 95% confidence intervals) from separate OLS regressions. Outcome measure is a four-point accuracy scale ranging from “Not at all accurate” (1) to “Very accurate” (4). Control variables include partisanship, ideology, political knowledge, feelings toward President Trump, age, race, and college education. See Table B5 for full results.

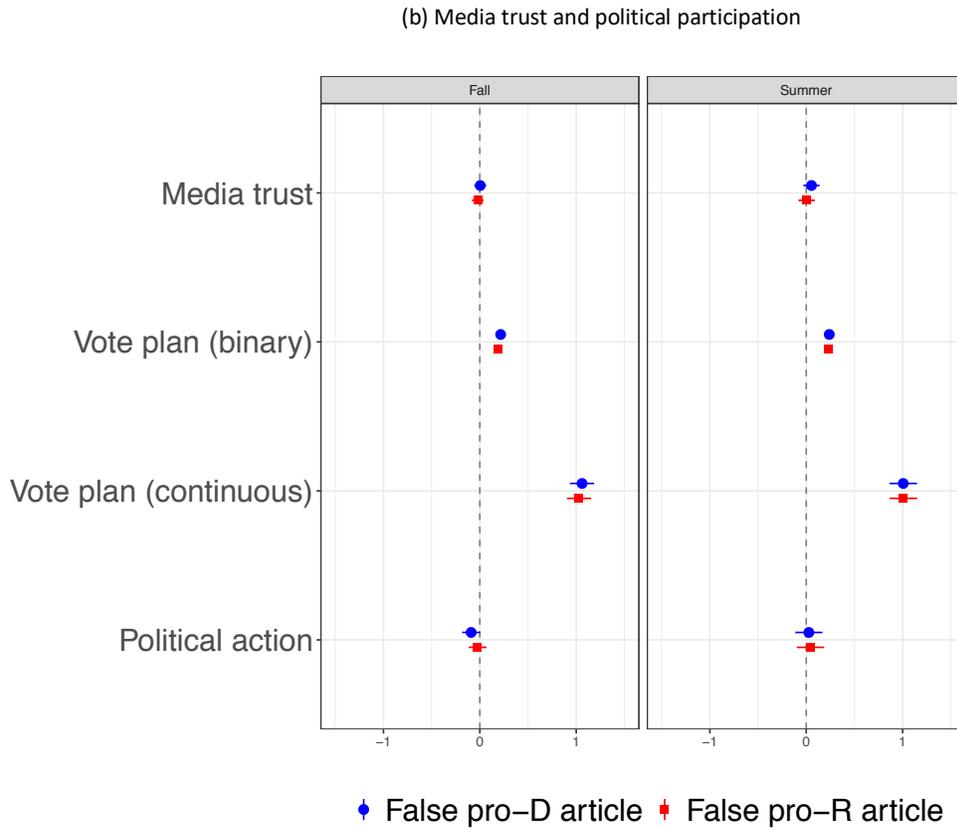
Finding 6: Randomized exposure to a dubious article increases intent to vote. However, article exposure has no measurable effect on other forms of participation, views toward the media, or on affective polarization.

The effects of exposure to dubious articles are largely limited to belief in the false claims they promote. Among all the other outcome measures we tested, the only variable affected by exposure to a dubious article is vote intention (see Figure 4). People who receive a false pro-Democratic or pro-Republican article reported a greater likelihood of voting in the upcoming election. Importantly, these effects did not differ significantly by congeniality of article content (see Tables B8–B10 in Online Appendix B). Respondents were more likely to say they would vote regardless of whether the article favored their preferred party; exposure to the article did not differentially mobilize supporting partisans or demobilize opponents. Moreover, these effects did not extend to other important attitudes and behavioral intentions. Specifically, exposure to a false article did not make respondents feel more negatively toward the media or the other party — we find no measurable effect on attitudes toward the media or affective polarization.

Figure 4. Effect of false article exposure on attitudes and participation intentions.



Effect of exposure to a false article from an untrustworthy website on media thermometer rating and affective polarization. Point estimates (with 95% confidence intervals) from separate OLS regressions. Outcome measures are respondent feelings toward the media on a 0–100 scale (media thermometer) and the difference in feelings respondents express toward their preferred party and the opposition party (both measured on a 0–100 scale). Control variables include partisanship, ideology, political knowledge, feelings toward President Trump, age, race, and college education. See Table B6 for full results.



Effect of exposure to a false article from an untrustworthy website on media trust, vote intentions, and willingness to take political action. Point estimates (with 95% confidence intervals) from separate OLS regressions. Outcome measures are a four-

point scale of self-reported trust and confidence in the mass media ranging from “None at all” (1) to “A great deal” (4); whether respondents say they are “absolutely certain” to vote or report already having voted (0/1); a five-point measure of vote intention ranging from answering no when asked if they plan to vote (1) to saying they are “absolutely certain” to vote or report already having voted (5); and respondents’ average response when asked if they would be willing to take a series of political actions on a scale from “Disagree strongly” (1) to “Agree strongly” (5). Control variables include partisanship, ideology, political knowledge, feelings toward President Trump, age, race, and college education. See Table B7 for full results.

Methods

To test our hypotheses, we draw on data from three two-wave survey panels conducted by the survey company YouGov during and immediately following the 2018 U.S. midterm elections:

- A two-wave panel study fielded June 25–July 3, 2018 (wave 1; N = 1,718) and July 9–17, 2018 (wave 2; N = 1,499), which we refer to as the summer 2018 data;
- A two-wave panel study fielded October 19–26 (wave 1; N=3,378) and October 30–November 6, 2018 (wave 2; N = 2,948), which we refer to as the fall 2018 data;
- A two-wave panel study fielded November 20–December 27, 2018 (wave 1; N = 4,907) and December 14, 2018–January 3, 2019 (wave 2; N = 4,283), which we refer to as the winter 2018 data.

Respondents were selected by YouGov’s matching and weighting algorithm to approximate the demographic and political attributes of the U.S. population. Participants were ineligible to take part in more than one study.

These surveys were designed to include a separate experiment evaluating whether a media literacy intervention affects how people evaluate both mainstream and false news headlines. The results of this experiment are examined in a separate study. Due to a programming error, all respondents in the summer and fall surveys were exposed to the intervention in the first wave of the two-wave survey. The experiment reported in this article was then conducted in the second wave in each of the summer and fall surveys. Given that all respondents also received the media literacy intervention, the effect sizes of our false and inflammatory news exposure treatments may underestimate the true effects of these articles. Future studies that we intend to publish will include additional analyses that were preregistered prior to conducting each of these studies.

Information consumption in this study is measured using behavioral data on respondents’ web visits collected unobtrusively in the 1-3 week period before the initial survey wave. Data is available from users’ laptop or desktop computers as well as their mobile devices and tablets (though coverage of the latter is limited)³. For each survey, we created a binary measure indicating whether we observe a visit to an untrustworthy website in the 1-3 week period for which we have data prior to the respondent completing the survey. This approach allows us to measure consumption of dubious online information directly rather than using survey questions and hoping that respondents both remember and faithfully report all the websites they visited.

We use regression to first examine associations between prior consumption of news from untrustworthy websites and feelings toward the media and the political parties, belief in topical misperceptions, and intent to vote or engage in other forms of political participation. Misperception questions focused on topics such as Trump’s immigration policy, the Special Counsel investigation of Russian interference in the 2016 United States elections, and Brett Kavanaugh’s Supreme Court

³ Web visits are collected anonymously with users’ permission through a mix of browser plug-ins, proxies, and VPNs. The providers of this data were the firms Wakoopa (summer 2018) and Reality Mine (fall and winter 2018).

nomination. These findings are correlational and cannot establish causality. In addition, because intention to vote or engage in other forms of political misinformation were only measured after the experiment, we consider only individuals assigned to the control condition in observational analyses of these outcome measures.

We also use experiments to assess the effects of exposure to a false and inflammatory article from an untrustworthy website. In the second wave of the summer and fall surveys, respondents were randomly assigned to read a false and inflammatory pro-Democratic article with 1/3 probability, a false and inflammatory pro-Republican article with 1/3 probability, or an unrelated placebo article with 1/3 probability (These treatments and relevant question wordings are provided in Online Appendix A.) This approach allows us to measure the causal effect of exposure to an article from an untrustworthy website with greater confidence than we can using non-experimental data.

Because we frequently expect different results by party, we report results for Democrats and Republicans (including those who lean toward one party) only; independents who do not lean toward other party are excluded. This exclusion simplifies interpretation of our results.

Bibliography

- Benkler, Y., Faris, R., and Roberts, H. (2018). *Network Propaganda: Manipulation, Disinformation, and Radicalization in American Politics*. Oxford University Press.
- Druckman, J. N. and Leeper, T. J. (2012). Learning more from political communication experiments: Pretreatment and its effects. *American Journal of Political Science*, 56(4):875–896.
- Garrett, R. K., Gvirsman, S. D., Johnson, B. K., Tsifti, Y., Neo, R., and Dal, A. (2014). Implications of pro-and counter-attitudinal information exposure for affective polarization. *Human Communication Research*, 40(3):309–332.
- Grinberg, N., Joseph, K., Friedland, L., Swire-Thompson, B., and Lazer, D. (2019). Fake news on Twitter during the 2016 U.S. presidential election. *Science*, 363(6425):374–378.
- Guess, A., Nyhan, B., and Reifler, J. (2019a). National news, local lens? findings from the 2019 Poynter media trust survey. October 21, 2019. Downloaded November 14, 2019 from <http://www.dartmouth.edu/~nyhan/media-trust-report-2019.pdf>.
- Guess, A., Nyhan, B., and Reifler, J. (2019b). Selective exposure to misinformation: Evidence from the consumption of fake news during the 2016 U.S. presidential campaign. Unpublished manuscript. Downloaded March 19, 2019 from <http://www-personal.umich.edu/~bnyhan/fake-news-2016.pdf>.
- Iyengar, S., Sood, G., and Lelkes, Y. (2012). Affect, not ideology: A social identity perspective on polarization. *Public Opinion Quarterly*, 76(3):405–431.
- Kalogeropoulos, A., Suiter, J., Udris, L., and Eisenegger, M. (2019). News media trust and news consumption: Factors related to trust in news in 35 countries. *International Journal of Communication*, 13:22.
- Suhay, E., Bello-Pardo, E., and Maurer, B. (2018). The polarizing effects of online partisan criticism: Evidence from two experiments. *International Journal of Press/Politics*, 23(1):95–115.
- Tsifti, Y. and Cappella, J. N. (2003). Do people watch what they do not trust? exploring the association between news media skepticism and exposure. *Communication Research*, 30(5):504–529.
- Valentino, N. A., Brader, T., Groenendyk, E. W., Gregorowicz, K., and Hutchings, V. L. (2011). Election night's alright for fighting: The role of emotions in political participation. *Journal of Politics*, 73(1):156–170.

Funding

Funding support for this research was provided by Democracy Fund, the European Research

Council (ERC) under the European Union’s Horizon 2020 research and innovation program (grant agreement No. 682758), the Nelson A. Rockefeller Center at Dartmouth College, and the Weidenbaum Center on the Economy, Government, and Public Policy at Washington University in St. Louis.

Competing interests

Guess and Nyhan are the principal investigators on separate research projects supported by funds provided by Facebook. The other authors are co-authors on a paper reporting the results of the study for which Nyhan was the principal investigator.

Ethics

All human subjects provided informed consent for the studies reported in this article, which were approved by institutional review boards at Dartmouth College, the University of Michigan, Washington University in St. Louis, the University of Exeter, and Princeton University. All ethnicity/gender categories were determined by respondents from options provided by the researchers or the survey company YouGov. These data were collected to ensure that participants approximated a nationally representative sample.

Copyright

This is an open access article distributed under the terms of the Creative Commons Attribution License, which permits unrestricted use, distribution, and reproduction in any medium, provided that the original author and source are properly credited.

Data availability

The data and code necessary to replicate all the findings in this article will be made available on Dataverse upon publication of this article. Per our human subjects protocols, we will protect respondent privacy by only including individual-level summary data of respondents’ web consumption (e.g., number of untrustworthy websites visited) in the replication data.