

The Honorable Michael F. Doyle
Chair, Subcommittee on Communications and Technology
Committee on Energy & Commerce, House of Representatives

The Honorable Robert E. Latta
Ranking Member, Subcommittee on Communications and Technology
Committee on Energy & Commerce, House of Representatives

The Honorable Janice D. Schakowsky
Chair, Subcommittee on Consumer Protection and Commerce
Committee on Energy & Commerce, House of Representatives

The Honorable Cathy McMorris Rodgers
Ranking Member, Subcommittee on Consumer Protection and Commerce
Committee on Energy & Commerce, House of Representatives

Dear Mr. Doyle, Mr. Latta, Ms. Schakowsky, and Ms. McMorris Rodgers:

A growing body of academic and journalistic research has provided evidence for the patterns of who produces, shares, and is exposed to false or misleading news. Aided by access to new datasets and analysis tools, recent work has emphasized the sometimes-large levels of engagement on pieces of misinformation. However, **relatively little is known about whether people are susceptible to believing false news when they encounter it online**. While the presence of false news in information ecosystems might be problematic *per se*, knowing whether people are susceptible to believing false information is an important contribution to our understanding of the possible threats posed by our networked digital environment. This knowledge is especially important during the pandemic, when susceptibility to false news has the potential to undermine the public's adoption of evidence-based public health responses and policies.

To this end, we developed a pipeline for collecting real-time evaluations of popular news articles by both professional fact checkers and ordinary respondents within 48 hours of an article's publication. From these data—which includes 5,408 American adults and covers the period from December 2019 through February 2020—we find evidence supporting two key findings:

1. **People aren't good at identifying false news:** 33.7% of the time laypeople saw false news, they incorrectly identified it as true. Only one-third of respondents correctly identified false news articles, with the remaining third unable to make a determination in either direction.
2. **Similar levels of susceptibility were seen on false news related to COVID-19 in the early days of the pandemic:** During this period, our transparent article selection mechanism sourced four fake articles that promoted the unfounded rumor the novel coronavirus was intentionally developed in a laboratory. On average, we saw higher levels (40%) of respondents unable to determine veracity in either direction. Still, roughly 30% of respondents incorrectly labeling these four articles as true.

The four coronavirus-related articles included in our study, covering only a single rumor, are hardly representative of the diversity of false information social media users are seeing online. While we are currently undertaking research including other types of coronavirus-related misinformation, these findings suggest non-trivial numbers of people are susceptible false news when they first encounter it. These findings also suggest that efforts to remove coronavirus-related misinformation will need to be swift — and implemented early in an article’s life cycle — to stop the spread of false news before people have the opportunity to believe it.

Sincerely,

Joshua A Tucker

Faculty Co-Director, NYU Center for Social Media and Politics

Director, Jordan Center for the Advanced Study of Russia

Zeve Sanderson

Executive Director, NYU Center for Social Media and Politics

It's not easy for ordinary citizens to identify fake news

And fake coronavirus news is no exception.

By **Zeve Sanderson, Kevin Aslett, Will Godel, Nathaniel Persily, Jonathan Nagler, Richard Bonneau** and **Joshua Tucker**

April 7, 2020 at 6:00 a.m. EDT

Fake news makes up a relatively small portion of Americans' news consumption, research has shown. Throughout 2016, the vast majority of U.S. Facebook users did not share any fake news articles. In the month leading up to the election, fake news sources only made up roughly 1 percent of the average Twitter user's political news on the platform.

In 2020, however, even small amounts of fake news about coronavirus can have dire consequences. The current public health crisis requires the coordinated actions of individuals — maintaining social distance, buying reasonable quantities of food and supplies, and following the latest medical advice rather than bogus cures.

This means fake covid-19 news has potential to undermine the public's adoption of evidence-based public health responses and policies. And in the area of public health, externalities are quite meaningful: Quack cures can lead to more covid-19 cases. Disinformation can be damaging in other ways — in the United Kingdom, for

instance, fake news blaming 5G wireless for coronavirus apparently led to arson attacks against cellphone towers.

How good are people at sifting out fake news? In a collaboration between the NYU Center for Social Media and Politics and the Stanford Cyber Policy Center (supported by the Hewlett Foundation, we've been investigating whether ordinary individuals in the United States who encounter news when it first appears online — before fact-checkers like Snopes and PolitiFact have an opportunity to issue reports about an article's veracity — are able to identify whether articles contain true or false information.

Unfortunately, it seems quite difficult for people to identify false or misleading news, and the limited number of coronavirus news stories in our collection are no exception.

How we did our research

Over a 13-week period, our study allowed us to capture people's assessments of fresh news articles in real time. Each day of the study, we relied on a fixed, preregistered process to select five popular articles published within the previous 24 hours.

The five articles were balanced between conservative, liberal and nonpartisan sources, as well as from mainstream news websites and from websites known to produce fake news. In total, we sent 150 total articles to 90 survey respondents each, which meant over the course of the study we received over 13,500 ratings of news stories from ordinary citizens, whom we reached via Qualtrics. To our knowledge, no individual took the study twice; over the course of the study, we ended up with 5,408 unique respondents, who were nationally representative.

We also sent these articles separately to six independent fact-checkers, and treated their most common response — true, false/misleading, or cannot determine — for each article as the “correct” answer for that article.

People aren’t good at identifying false news

When shown an article that was rated “true” by the professional fact-checkers, respondents correctly identified the article as true 62 percent of the time. When the source of the true news story was a mainstream news source, respondents correctly identified the article as true 73 percent of the time.

However, for each article the professional fact-checkers rated “false/misleading,” the study participants were as likely to say it was true as they were to say it was false or misleading. Roughly one-third of the time they told us they were unable to determine the veracity of the article. In other words, people on the whole were unable to correctly classify false or misleading news.

How did people gauge coronavirus-related misinformation?

Between Jan. 22 and Feb. 6 — this was the first two weeks after confirmed cases of the coronavirus appeared outside of China, but before any social distancing orders were in place in the United States — four of the articles in our study that fact-checkers rated as false or misleading were related to the coronavirus.

All four articles promoted the unfounded rumor that the virus was intentionally developed in a laboratory. Although accidental releases of pathogens from labs have previously caused significant morbidity and mortality, in the current pandemic multiple pieces of evidence suggest this virus is of natural origin. There’s little evidence the virus was manufactured or altered.

When we asked people to rate the veracity of these four articles, the results mirrored those from the full study: Only 30 percent of participants correctly classified them

as false or misleading.

Additionally, on average, respondents seemed to have more trouble deciding what to think about false covid-19 stories, leading to a higher proportion of “could not determine” responses than we saw for the stories on other topics our professional fact-checkers rated as “false/misleading.” This finding suggests it may be particularly difficult to identify misinformation in newly emerging topics (although further research would be needed to confirm this).

Education matters, a little

Study participants with higher levels of education did better on identifying both fake news overall and coronavirus-related fake news — but were far from being able to correctly weed out misinformation all of the time. In fact, no group, regardless of education level, was able to correctly identify the stories that the professional fact-checkers had labeled as false or misleading more than 40 percent of the time.

What do these findings reveal?

Taken together, our findings suggest there is widespread potential for vulnerability to misinformation when it first appears online. This is especially worrying during the current pandemic, when governments, public health officials and the media are constantly updating information. Nothing we saw in our data suggested the potential for misinformation on covid-19 was any lower than the overall findings — though it's important to note that our study did not include misinformation related to cures or treatments.

How individuals respond to the pandemic will depend, in large part, on the quality of information to which they are exposed — and the stories and the reports they find credible. In the current environment, misinformation has the potential to undermine social distancing efforts, to lead people to hoard supplies, or to promote the adoption of potentially dangerous fake cures.

Since the outbreak of covid-19, platforms such as Facebook and Twitter have been inundated with a diversity of false narratives, ranging from accusations of partisan hoaxes to medical misinformation propagating unverified treatments. This suggests the rapidly changing information environment surrounding covid-19 has the potential to magnify the vulnerabilities revealed by previous studies of online misinformation.

The four news articles included in our study, covering only a single rumor, are hardly representative of what social media users are seeing online. While more research is needed about other types of coronavirus-related misinformation, our findings suggest non-trivial numbers of people will believe false information to be true when they first encounter it. And it suggests that efforts to remove coronavirus-related misinformation will need to be swift — and implemented early in an article’s life cycle — to stop the spread of something else that’s dangerous: misinformation.

Zeve Sanderson is the executive director of the NYU Center for Social Media and Politics

Kevin Aslett is a postdoctoral researcher at the NYU Center for Social Media and Politics.

Will Godel is a PhD student in NYU’s Department of Politics and a graduate research affiliate of the NYU Center for Social Media and Politics.

Nathaniel Persily is the James B. McClatchy Professor of Law at Stanford University and co-director of Stanford’s Cyber Policy Center.

Jonathan Nagler is a professor of politics at NYU and co-director of the NYU Center for Social Media and Politics.

Richard Bonneau is a professor of biology and computer science at NYU and co-director of the NYU Center for Social Media and Politics.

Joshua Tucker

Joshua A. Tucker is Professor of Politics at New York University, Director of NYU's Jordan Center for Advanced Study of Russia, and a co-Director of the NYU Center for Social Media and Politics.

Follow