

The Grim Conclusions of the Largest-Ever Study of Fake News

Falsehoods almost always beat out the truth on Twitter, penetrating further, faster, and deeper into the social network than accurate information.

• [ROBINSON MEYER](#)

THE ATLANTIC

March 8, 2018

“Falsehood flies, and the Truth comes limping after it,” Jonathan Swift once wrote.

It was hyperbole three centuries ago. But it is a factual description of social media, according to [an ambitious and first-of-its-kind study](#) published Thursday in *Science*.

The massive new study analyzes every major contested news story in English across the span of Twitter’s existence—some 126,000 stories, tweeted by 3 million users, over more than 10 years—and finds that the truth simply cannot compete with hoax and rumor. By every common metric, falsehood consistently dominates the truth on Twitter, the study finds: Fake news and false rumors reach more people, penetrate deeper into the social network, and spread much faster than accurate stories.

“It seems to be pretty clear [from our study] that false information outperforms true information,” said Soroush Vosoughi, a data scientist at MIT who has studied fake news since 2013 and who led this study. “And that is not just because of bots. It might have something to do with human nature.”

The study has already prompted alarm from social scientists. “We must redesign our information ecosystem in the 21st century,” write a group of 16 political scientists and legal scholars in an essay also [published](#) Thursday in *Science*. They call for a new drive of interdisciplinary research “to reduce the spread of fake news and to address the underlying pathologies it has revealed.”

“How can we create a news ecosystem ... that values and promotes truth?” they ask.

The new study suggests that it will not be easy. Though Vosoughi and his colleagues only focus on Twitter—the study was conducted using exclusive data that the company made available to MIT—their work has implications for Facebook, YouTube, and every major social network. Any platform that regularly amplifies engaging or provocative content runs the risk of amplifying fake news along with it.

Though the study is written in the clinical language of statistics, it offers a methodical indictment of the accuracy of information that spreads on these platforms. A false story is much more likely to go viral than a real story, the authors find. A false story reaches 1,500 people six times quicker, on average, than a true story does. And while false stories outperform the truth on every subject—including business, terrorism and war, science and technology, and entertainment—fake news about politics regularly does best.

Twitter users seem almost to *prefer* sharing falsehoods. Even when the researchers controlled for every difference between the accounts originating rumors—like whether that person had more followers or was verified—falsehoods were still 70 percent more likely to get retweeted than accurate news.

And blame for this problem cannot be laid with our robotic brethren. From 2006 to 2016, Twitter bots amplified true stories as much as they amplified false ones, the study found. Fake news prospers, the authors write, “because humans, not robots, are more likely to spread it.”

Political scientists and social-media researchers largely praised the study, saying it gave the broadest and most rigorous look so far into the scale of the fake-news problem on social networks, though some disputed its findings about bots and questioned its definition of news.

“This is a really interesting and impressive study, and the results around how demonstrably untrue assertions spread faster and wider than demonstrable true ones do, within the sample, seem very robust, consistent, and well supported,” said [Rasmus Kleis Nielsen](#), a professor of political communication at the University of Oxford, in an email.

“I think it’s very careful, important work,” Brendan Nyhan, a professor of government at Dartmouth College, told me. “It’s excellent research of the sort that we need more of.”

“In short, I don’t think there’s any reason to doubt the study’s results,” said [Rebekah Tromble](#), a professor of political science at Leiden University in the Netherlands, in an email.

What makes this study different? In the past, researchers have looked into the problem of falsehoods spreading online. They’ve often focused on rumors around singular events, like the speculation [that preceded the discovery of the Higgs boson](#) in 2012 or the rumors [that followed the Haiti earthquake](#) in 2010.

This new paper takes a far grander scale, looking at nearly the entire lifespan of Twitter: every piece of controversial news that propagated on the service from September 2006 to December 2016. But to do that, Vosoughi and his colleagues had to answer a more preliminary question first: *What is truth? And how do we know?*

It’s a question that can have life-or-death consequences.

“[Fake news] has become a white-hot political and, really, cultural topic, but the trigger for us was personal events that hit Boston five years ago,” said [Deb Roy](#), a media scientist at MIT and one of the authors of the new study.

On April 15, 2013, two bombs exploded near the route of the Boston Marathon, killing three people and injuring hundreds more. Almost immediately, wild conspiracy theories

about the bombings took over Twitter and other social-media platforms. The mess of information only grew more intense on April 19, when the governor of Massachusetts asked millions of people to remain in their homes as police conducted a huge manhunt.

“I was on lockdown with my wife and kids in our house in Belmont for two days, and Soroush was on lockdown in Cambridge,” Roy told me. Stuck inside, Twitter became their lifeline to the outside world. “We heard a lot of things that were not true, and we heard a lot of things that did turn out to be true” using the service, he said.

The ordeal soon ended. But when the two men reunited on campus, they agreed it seemed seemed silly for Vosoughi—then a Ph.D. student focused on social media—to research anything but what they had just lived through. Roy, his adviser, blessed the project.

He made a truth machine: an algorithm that could sort through torrents of tweets and pull out the facts most likely to be accurate from them. It focused on three attributes of a given tweet: the properties of its author (were they verified?), the kind of language it used (was it sophisticated?), and how a given tweet propagated through the network.

“The model that Soroush developed was able to predict accuracy with a far-above-chance performance,” said Roy. He earned his Ph.D. in 2015.

After that, the two men—and Sinan Aral, a professor of management at MIT—turned to examining how falsehoods move across Twitter as a whole. But they were back not only at the “what is truth?” question, but its more pertinent twin: How does *the computer* know what truth is?

They opted to turn to the ultimate arbiter of fact online: the third-party fact-checking sites. By scraping and analyzing six different fact-checking sites—including [Snopes](#), [Politifact](#), and [FactCheck.org](#)—they generated a list of tens of thousands of online rumors that had spread between 2006 and 2016 on Twitter. Then they searched Twitter for these rumors, using a proprietary search engine owned by the social network called [Gnip](#).

Ultimately, they found about 126,000 tweets, which, together, had been retweeted more than 4.5 million times. Some linked to “fake” stories hosted on other websites. Some started rumors themselves, either in the text of a tweet or in an attached image. (The team used [a special program](#) that could search for words contained within static tweet images.) And some contained true information or linked to it elsewhere.

Then they ran a series of analyses, comparing the popularity of the fake rumors with the popularity of the real news. What they found astounded them.

Speaking from MIT this week, Vosoughi gave me an example: There are lots of ways for a tweet to get 10,000 retweets, he said. If a celebrity sends Tweet A, and they have a couple million followers, maybe 10,000 people will see Tweet A in their timeline and decide to retweet it. Tweet A was broadcast, creating a big but shallow pattern.

Meanwhile, someone without many followers sends Tweet B. It goes out to their 20 followers—but one of those people sees it, and retweets it, and then one of *their* followers sees it and retweets it too, on and on until tens of thousands of people have seen and shared Tweet B.

Tweet A and Tweet B both have the same size audience, but Tweet B has more “depth,” to use Vosoughi’s term. It chained together retweets, going viral in a way that Tweet A never did. “It could reach 1,000 retweets, but it has a very different shape,” he said.

Here’s the thing: Fake news dominates *according to both metrics*. It consistently reaches a larger audience, *and* it tunnels much deeper into social networks than real news does. The authors found that accurate news wasn’t able to chain together more than 10 retweets. Fake news could put together a retweet chain 19 links long—and do it 10 times as fast as accurate news put together its measly 10 retweets.

These results proved robust even when they were checked by humans, not bots. Separate from the study, a group of undergraduate students fact-checked a random selection of roughly 13,000 English-language tweets from the same period. They found that false information outperformed true information in ways “nearly identical” to the main data set, according to the study.

What does this look like in real life? Take two examples from the last presidential election. In August 2015, a rumor circulated on social media that Donald Trump had let a sick child use his plane to get urgent medical care. *Snopes* [confirmed almost all of the tale as true](#). But according to the team’s estimates, only about 1,300 people shared or retweeted the story.

In February 2016, a rumor developed that Trump’s elderly cousin had recently died and that he had opposed the magnate’s presidential bid in his obituary. “As a proud bearer of the Trump name, I implore you all, please don’t let that walking mucus bag become president,” the obituary reportedly said. But *Snopes* could not find evidence of the cousin, or his obituary, and [rejected the story as false](#).

Nonetheless, roughly 38,000 Twitter users shared the story. And it put together a retweet chain three times as long as the sick-child story managed.

A false story [alleging the boxer Floyd Mayweather had worn a Muslim head scarf](#) to a Trump rally also reached an audience more than 10 times the size of the sick-child story.

Why does falsehood do so well? The MIT team settled on two hypotheses.

First, fake news seems to be more “novel” than real news. Falsehoods are often notably different from the all the tweets that have appeared in a user’s timeline 60 days prior to their retweeting them, the team found.

Second, fake news evoked much more emotion than the average tweet. The researchers created a database of the words that Twitter users used to reply to the 126,000 contested

tweets, then analyzed it [with a state-of-the-art sentiment-analysis tool](#). Fake tweets tended to elicit words associated with surprise and disgust, while accurate tweets summoned words associated with sadness and trust, they found.

The team wanted to answer one more question: Were Twitter bots helping to spread misinformation?

After using two different bot-detection algorithms on their sample of 3 million Twitter users, they found that they automated bots were spreading false news—but they were retweeting it at the same rate that they retweeted accurate information.

“The massive differences in how true and false news spreads on Twitter cannot be explained by the presence of bots,” Aral told me.

But some political scientists cautioned that this should not be used to disprove the role of Russian bots in seeding disinformation recently. An “army” of Russian-associated bots helped amplify divisive rhetoric after the school shooting in Parkland, Florida, [The New York Times has reported](#).

“It can both be the case that (1) over the whole 10-year data set, bots don’t favor false propaganda and (2) in a recent subset of cases, botnets have been strategically deployed to spread the reach of false propaganda claims,” said Dave Karpf, a political scientist at George Washington University, in an email.

“My guess is that the paper is going to get picked up as ‘scientific proof that bots don’t really matter!’ And this paper does indeed show that, if we’re looking at the full life span of Twitter. But the real bots debate assumes that their usage has recently escalated because strategic actors have poured resources into their use. This paper doesn’t refute that assumption,” he said.

Vosoughi agrees that his paper cannot determine whether the use of botnets changed at the end of the sample period. “We did not study the change in the role of bots across time,” he told me in an email. “This is an interesting question and one that we will probably look at in future work.”

Some political scientists also questioned the study’s definition of “news.” By turning to the fact-checking sites, the study blurs together a wide range of false information: outright lies, urban legends, hoaxes, spoofs, falsehoods, *and* “fake news.” It does not just look at fake news by itself—that is, articles or videos that look like news content, and which appear to have gone through a journalistic process, but which are actually made up.

Therefore, the study may undercount “non-contested news”: accurate news that is widely understood to be true. For many years, the most retweeted post in Twitter’s history [celebrated Obama’s re-election as president](#). But as his victory was not a widely disputed fact, *Snopes* and other fact-checking sites never confirmed it.

The study also elides *content* and *news*. “All our audience research suggests a vast majority of users see news as clearly distinct from content more broadly,” Nielsen, the Oxford professor, said in an email. “Saying that untrue content, including rumors, spread faster than true statements on Twitter is a bit different from saying false news and true news spread at different rates.”

But many researchers told me that simply understanding *why* false rumors travel so far, so fast, was as important as knowing that they do so in the first place.

“The key takeaway is really that content that *arouses strong emotions* spreads further, faster, more deeply, and more broadly on Twitter,” said Tromble, the political scientist, in an email. “This particular finding is consistent with research in a number of different areas, including psychology and communication studies. It’s also relatively intuitive.”

“False information online is often really novel and frequently negative,” said Nyhan, the Dartmouth professor. “We know those are two features of information generally that grab our attention as human beings and that cause us to want to share that information with others—we’re attentive to novel threats and especially attentive to negative threats.”

“It’s all too easy to create both when you’re not bound by the limitations of reality. So people can exploit the interaction of human psychology and the design of these networks in powerful ways,” he added.

He lauded Twitter for making its data available to researchers and called on other major platforms, like Facebook, to do the same. “In terms of research, the platforms are the whole ballgame. We have so much to learn but we’re so constrained in what we can study without platform partnership and collaboration,” he said.

“These companies now exercise a great deal of power and influence over the news that people get in our democracy. The amount of power that platforms now hold means they have to face a great deal of scrutiny and transparency,” he said. “We can study Twitter all day, but only about 12 percent of Americans are on it. It’s important for journalists and academics, but it’s not how most people get their news.”

In a statement, Twitter said that it was hoping to expand its work with outside experts. In a series of tweets last week, Jack Dorsey, the company’s CEO, [said the company](#) hoped to “increase the collective health, openness, and civility of public conversation, and to hold ourselves publicly accountable toward progress.”

Facebook did not respond to a request for comment.

But Tromble, the political-science professor, said that the findings would likely apply to Facebook, too. “Earlier this year, Facebook announced that it would restructure its News Feed to favor ‘meaningful interaction,’” she told me.

“It became clear that they would gauge ‘meaningful interaction’ based on the number of comments and replies to comments a post receives. But, as this study shows, that only

further incentivizes creating posts full of disinformation and other content likely to garner strong emotional reactions,” she added.

“Putting my conservative scientist hat on, I’m not comfortable saying how this applies to other social networks. We only studied Twitter here,” said Aral, one of the researchers. “But my intuition is that these findings are broadly applicable to social-media platforms in general. You could run this exact same study if you worked with Facebook’s data.”

Yet these do not encompass the most depressing finding of the study. When they began their research, the MIT team expected that users who shared the most fake news would basically be crowd-pleasers. They assumed they would find a group of people who obsessively use Twitter in a partisan or sensationalist way, accumulating more fans and followers than their more fact-based peers.

In fact, the team found that the opposite is true. Users who share accurate information have more followers, and send more tweets, than fake-news sharers. These fact-guided users have also been on Twitter for longer, and they are more likely to be verified. In short, the most trustworthy users can boast every obvious structural advantage that Twitter, either as a company or a community, can bestow on its best users.

The truth has a running start, in other words—but inaccuracies, somehow, still win the race. “Falsehood diffused further and faster than the truth *despite* these differences [between accounts], not because of them,” write the authors.

This finding should dispirit every user who turns to social media to find or distribute accurate information. It suggests that no matter how adroitly people *plan* to use Twitter—no matter how meticulously they curate their feed or follow reliable sources—they can still get snookered by a falsehood in the heat of the moment.

It suggests—to me, at least, a Twitter user since 2007, and someone [who got his start in journalism](#) because of the social network—that social-media platforms do not encourage the kind of behavior that anchors a democratic government. On platforms where every user is at once a reader, a writer, and a publisher, falsehoods are too seductive not to succeed: The thrill of novelty is too alluring, the titillation of disgust too difficult to transcend. After a long and aggravating day, even the most staid user might find themselves lunging for the politically advantageous rumor. Amid an anxious election season, even the most public-minded user might subvert their higher interest to win an argument.

It is unclear which interventions, if any, could reverse this tendency toward falsehood. “We don’t know enough to say what works and what doesn’t,” Aral told me. There is little evidence that people change their opinion because they see a fact-checking site reject one of their beliefs, for instance. Labeling fake news as such, on a social network or search engine, may do little to deter it as well.

In short, social media seems to systematically amplify falsehood at the expense of the truth, and no one—neither experts nor politicians nor tech companies—knows how to reverse that trend. It is a dangerous moment for any system of government premised on a common public reality.