


# How Your Brain Tricks You Into Believing Fake News



By [Katy Steinmetz](#)  
August 9, 2018

Sitting in front of a computer not long ago, a tenured history professor faced a challenge that billions of us do every day: deciding whether to believe something on the Internet.

On his screen was an article published by a group called the American College of Pediatricians that discussed how to handle bullying in schools. Among the advice it offered: schools shouldn't highlight particular groups targeted by bullying because doing so might call attention to "temporarily confused adolescents."

Scanning the site, the professor took note of the ".org" web address and a list of academic-looking citations. The site's sober design, devoid of flashy, autoplaying videos, lent it credibility, he thought. After five minutes, he had found little reason to doubt the article. "I'm clearly looking at an official site," he said.

What the professor never realized as he focused on the page's superficial features is that the group in question is a socially conservative splinter faction that broke in 2002 from the mainstream American Academy of Pediatrics over the issue of adoption by same-sex couples. It has been accused of promoting antigay policies, and the Southern Poverty Law Center designates it as a hate group.



Trust was the issue at hand. The bookish professor had been asked to assess the article as part of an experiment run by Stanford University psychologist Sam Wineburg. His team, known as the Stanford History Education Group, has given scores of subjects such tasks in hopes of answering two of the most vexing questions of the Internet age: Why are even the smartest among us so bad at making judgments about what to trust on the web? And how can we get better?

Wineburg's team has found that Americans of all ages, from digitally savvy tweens to high-IQ academics, fail to ask important questions about content they encounter on a browser, adding to research on our online gullibility. Other studies have shown that people retweet links without clicking on them and rely too much on search engines. A 2016 Pew poll found that nearly a quarter of Americans said they had shared a made-up news story. In his experiments, MIT cognitive scientist David Rand has found that, on average, people are inclined to believe false news at least 20% of the time. "We are all driving cars, but none of us have licenses," Wineburg says of consuming information online.

Our inability to parse truth from fiction on the Internet is, of course, more than an academic matter. The scourge of "fake news" and its many cousins—from clickbait to "deep fakes" (realistic-looking videos showing events that never happened)—have experts fearful for the future of democracy. Politicians and technologists have warned that meddlers are trying to manipulate elections around the globe by spreading disinformation. That's what Russian agents did in 2016, according to U.S. intelligence agencies. And on July 31, [Facebook](#) revealed that it had found evidence of a political-influence campaign on the platform ahead of the 2018 midterm elections. The authors of one now defunct page got thousands of people to express interest in attending a made-up protest that apparently aimed to put white nationalists and left-wingers on the same streets.

But the stakes are even bigger than elections. Our ability to vet information matters every time a mother asks [Google](#) whether her child should be vaccinated and every time a kid encounters a Holocaust denial on Twitter. In India, false rumors about child kidnappings that spread on WhatsApp have prompted mobs to beat innocent people to death. "It's the equivalent of a public-health crisis," says Alan Miller, founder of the nonpartisan News Literacy Project.

There is no quick fix, though tech companies are under increasing pressure to come up with solutions. Facebook lost more than \$120 billion in stock value in a single day in July as the company dealt with a range of issues limiting its growth, including criticism about how conspiracy theories spread on the platform. But engineers can't teach machines to decide what is true or false in a world where humans often don't agree.

In a country founded on free speech, debates over who adjudicates truth and lies online are contentious. Many welcomed the decision by major tech companies in early August to remove content from florid conspiracy theorist Alex Jones, who has alleged that passenger-jet contrails are damaging people's brains and spread claims that families of Sandy Hook massacre victims are actors in an elaborate hoax. But others cried censorship. And even if law enforcement and intelligence agencies could ferret out every bad actor with a keyboard, it seems unwise to put the government in charge of scrubbing the Internet of misleading statements.

What is clear, however, is that there is another responsible party. The problem is not just malicious bots or chaos-loving trolls or Macedonian teenagers pushing phony stories for profit. The problem is also us, the susceptible readers. And experts like Wineburg believe that the better we understand the way we think in the digital world, the better chance we have to be part of the solution.

**We don't fall** for false news just because we're dumb. Often it's a matter of letting the wrong impulses take over. In an era when the average American spends 24 hours each week online—when we're always juggling inboxes and feeds and alerts—it's easy to feel like we don't have time to read anything but headlines. We are social animals, and the desire for likes can supersede a latent feeling that a story seems dicey. Political convictions lead us to lazy thinking. But there's an even more fundamental impulse at play: our innate desire for an easy answer.

Humans like to think of themselves as rational creatures, but much of the time we are guided by emotional and irrational thinking. Psychologists have shown this through the study of cognitive shortcuts known as heuristics. It's hard to imagine getting through so much as a trip to the grocery store without these helpful time-savers. "You don't and can't take the time and energy to examine and compare every brand of yogurt," says Wray Herbert, author of *On Second Thought: Outsmarting Your Mind's Hard-Wired Habits*. So we might instead rely on what is known as the familiarity heuristic, our tendency to assume that if something is familiar, it must be good and safe.

These habits of mind surely helped our ancestors survive. The problem is that relying on them too much can also lead people astray, particularly in an online environment. In one of his experiments, MIT's Rand illustrated the dark side of the fluency heuristic, our tendency to believe things we've been exposed to in the past. The study presented subjects with headlines—some false, some true—in a format identical to what users see on Facebook. Rand found that simply being exposed to fake news (like an article that claimed President Trump was going to brink back the draft) made people more likely to rate those stories as accurate later on in the experiment. If you've seen something before, "your brain subconsciously uses that as an indication that it's true," Rand says.

This is a tendency that propagandists have been aware of forever. The difference is that it has never been easier to get eyeballs on the message, nor to get enemies of the message to help spread it. The researchers who conducted the Pew poll noted that one reason people knowingly share made-up news is to "call out" the stories as fake. That might make a post popular among like-minded peers on social media, but it can also help false claims sink into the collective consciousness.

Academics are only beginning to grasp all the ways our brains are shaped by the Internet, a key reason that stopping the spread of misinformation is so tricky. One attempt by Facebook shows how introducing new signals into this busy domain can backfire. With hopes of curtailing junk news, the company started attaching warnings to posts that contained claims that fact-checkers had rated as false. But a study found that this can make users more likely to believe any unflagged post. Tessa Lyons-Laing, a product manager who works on Facebook's News Feed, says the company toyed with the idea of alerting users to hoaxes that were traveling around the web each day before realizing that an "immunization approach" might be counterproductive. "We're really trying to understand the problem and to be thoughtful about the research and therefore, in some cases, to move slower," she says.



Part of the issue is that people are still relying on outdated shortcuts, the kind we were taught to use in a library. Take the professor in Wineburg's study. A list of citations means one thing when it appears in a book that has been vetted by a publisher, a fact-checker and a librarian. It means quite another on the Internet, where everyone has access to a personal printing press. Newspapers used to physically separate hard news and commentary, so our minds could easily grasp what was what. But today two-thirds of Americans get news from social media, where posts from publishers get the same packaging as birthday greetings and rants. Content that warrants an emotional response is mixed with things that require deeper consideration. "It all looks identical," says Harvard researcher Claire Wardle, "so our brain has to work harder to make sense of those different types of information."

Instead of working harder, we often try to outsource the job. Studies have shown that people assume that the higher something appears in Google search results, the more reliable it is. But Google's algorithms are surfacing content based on keywords, not truth. If you ask about using apricot seeds to cure cancer, the tool will dutifully find pages asserting that they work. "A search engine is a search engine," says Richard Gingras, vice president of news at Google. "I don't think anyone really wants Google to be the arbiter of what is or is not acceptable expression."

That's just one example of how we need to retrain our brains. We're also inclined to trust visuals, says Wardle. But some photos are doctored, and other legitimate ones are put in false contexts. On Twitter, people use the size of others' followings as a proxy for reliability, yet millions of followers have been paid for (and an estimated 10% of "users" may be bots). In his studies, Wineburg found that people of all ages were inclined to evaluate sources based on features like the site's URL and graphic design, things that are easy to manipulate.

It makes sense that humans would glom on to just about anything when they're so worn out by the news. But when we resist snap judgments, we are harder to fool. "You just have to stop and think," Rand says of the experiments he has run on the subject. "All of the data we have collected suggests that's the real problem. It's not that people are being super-biased and using their reasoning ability to trick themselves into believing crazy stuff. It's just that people aren't stopping. They're rolling on."

**That is, of course,** the way social-media platforms have been designed. The endless feeds and intermittent rewards are engineered to keep you reading. And there are other environmental factors at play, like people's ability to easily seek out information that confirms their beliefs. But Rand is not the only academic who believes that we can take a big bite out of errors if we slow down.

Wineburg, an 18-year veteran of Stanford, works out of a small office in the center of the palm-lined campus. His group's specialty is developing curricula that teachers across the nation use to train kids in critical thinking. Now they're trying to update those lessons for life in a digital age. With the help of funding from Google, which has devoted \$3 million to the digital-literacy project they are part of, the researchers hope to deploy new rules of the road by next year, outlining techniques that anyone can use to draw better conclusions on the web.

His group doesn't just come up with smart ideas; it tests them. But as they set out to develop these lessons, they struggled to find research about best practices. "Where are the studies about what superstars do, so that we might learn from them?" Wineburg recalls thinking, sitting in the team's office beneath a print of the Tabula Rogeriana, a medieval map that pictures the world in a way we now see as upside-down. Eventually, a cold email to an office in New York revealed a promising model: professional fact-checkers.

Fact-checkers, they found, didn't fall prey to the same missteps as other groups. When presented with the American College of Pediatricians task, for example, they almost immediately left the site and started opening new tabs to see what the wider web had to say about the organization. Wineburg has dubbed this lateral reading: if a person never leaves a site—as the professor failed to do—they are essentially wearing blinders. Fact-checkers not only zipped to additional sources, but also laid their references side by side, to better keep their bearings.

In another test, the researchers asked subjects to assess the website MinimumWage.com. In a few minutes' time, 100% of fact-checkers figured out that the site is backed by a PR firm that also represents the restaurant industry, a sector that generally opposes raising hourly pay. Only 60% of historians and 40% of Stanford students made the same discovery, often requiring a second prompt to find out who was behind the site.

Another tactic fact-checkers used that others didn't is what Wineburg calls "click restraint." They would scan a whole page of search results—maybe even two—before choosing a path forward. "It's the ability to stand back and get a sense of the overall territory in which you've landed," he says, "rather than promiscuously clicking on the first thing." This is important, because people or organizations with an agenda can game search results by packing their sites with keywords, so that those sites rise to the top and more objective assessments get buried.

The lessons they've developed include such techniques and teach kids to always start with the same question: Who is behind the information? Although it is still experimenting, a pilot that Wineburg's team conducted at a college in California this past spring showed that such tiny behavioral changes can yield significant results. Another technique he champions is simpler still: just read it.

One study found that 6 in 10 links get retweeted without users' reading anything besides someone else's summation of it. Another found that false stories travel six times as fast as true ones on Twitter, apparently because lies do a better job of stimulating feelings of surprise and disgust. But taking a beat can help us avoid knee-jerk reactions, so that we don't blindly add garbage to the vast flotillas already clogging up the web. "What makes the false or hyperpartisan claims do really well is they're a bit outlandish," Rand says. "That same thing that makes them successful in spreading online is the same thing that, on reflection, would make you realize it wasn't true."

**Tech companies** have a big role to play in stemming the tide of misinformation, and they're working on it. But they have also realized that what Harvard's Wardle calls our "information disorder" cannot be solved by engineers alone. Algorithms are good at things like identifying fake accounts, and platforms are flagging millions of them every week. Yet machines could only take Facebook so far in identifying the most recent influence campaign.

One inauthentic page, titled "Resisters," ginned up a counterprotest to a "white civil rights" rally planned for August in Washington, D.C., and got legitimate organizations to help promote it. More than 2,600 people expressed interest in going before Facebook revealed that the page was part of a coordinated operation, disabled the event and alerted users. The company has hired thousands of content reviewers that have the sophistication to weed through tricky mixes of truth and lies. But Facebook can't employ enough humans to manually review the billions of posts that are put up each day, across myriad countries and languages.

Many misleading posts don't violate tech companies' terms of service. Facebook, one of the firms that removed content from Jones, said the decision did not relate to "false news" but prohibitions against rhetoric such as "dehumanizing language." [Apple](#) and Spotify cited rules against hate speech, which is

8/12/2018

## How Your Brain Tricks You Into Believing Fake News | Time

generally protected by the First Amendment. “With free expression, you get the good and the bad, and you have to accept both,” says Google’s Gingras. “And hopefully you have a society that can distinguish between the two.”

You also need a society that cares about that distinction. Schools make sense as an answer, but it will take money and political will to get new curricula into classrooms. Teachers must master new material and train students to be skeptical without making them cynical. “Once you start getting kids to question information,” says Stanford’s Sarah McGrew, “they can fall into this attitude where nothing is reliable anymore.” Advocates want to teach kids other defensive skills, like how to reverse-search an image (to make sure a photo is really portraying what someone says it is) and how to type a neutral query into the search bar. But even if the perfect lessons are dispersed for free online, anyone who has already graduated will need to opt in. They will have to take initiative and also be willing to question their prejudices, to second-guess information they might like to believe. And relying on open-mindedness to defeat tribal tendencies has not proved a winning formula in past searches for truth.

That is why many advocates are suggesting that we reach for another powerful tool: shame. Wardle says we need to make sharing misinformation as shameful as drunk driving. Wineburg invokes the environmental movement, saying we need to cultivate an awareness of “digital pollution” on the Internet. “We have to get people to think that they are littering,” Wineburg says, “by forwarding stuff that isn’t true.” The idea is to make people see the aggregate effect of little actions, that one by one, ill-advised clicks contribute to the web’s being a toxic place. Having a well-informed citizenry may be, in the big picture, as important to survival as having clean air and water. “If we can’t come together as a society around this issue,” Wineburg says, “it is our doom.”

*This appears in the August 20, 2018 issue of TIME.*

### **SPONSORED FINANCIAL CONTENT**