


IDEAS

Virus Experts Aren't Getting the Message Out

If the authorities can't satisfy the public's desire to know more, others will fill the void with misinformation.

By Renée DiResta



The Atlantic

MAY 6, 2020

IN A 10-WEEK SPAN at the end of 2019, 83 people—most under the age of 5—died from a disease outbreak in the South Pacific island nation of Samoa. The government undertook drastic measures to stop the highly contagious disease that had infected thousands, resulting in the hospitalization of 33 percent of those who contracted it. Schools were shut down and children were banned from public gatherings.

In that case, the disease was measles, which is preventable. The measles, mumps, and rubella (MMR) vaccination rate had already begun to fall in

Samoa when, in 2018, two children died after receiving improperly prepared injections—after which the vaccination rate dropped to just 34 percent. When the 2019 measles outbreak struck, the disease spread quickly. As the Samoan government declared a state of emergency, anti-vaccine activists emerged to do what they always do during epidemics: push conspiracy theories, hawk specious cures, and spam the social-media pages of the government and health authorities trying to get accurate information to the public. The anti-vaccine groups—some with more than 100,000 members—worked together to try to undermine the vaccination campaign the government implemented, directing members to comment on specific posts and vote in online public-opinion polls. To any observer who didn't know better, it looked like a mass public uprising—one with its own folk hero.

[Read: The small, small world of Facebook's anti-vaxxers](#)


A man named Edwin Tamasese, the manager of a coconut-farming collective and a self-proclaimed holistic healer, warned people to avoid the MMR vaccine in favor of papaya-leaf extract and vitamin C. Anti-vaccine activists in the United States supported Tamasese via social media; they started drives to collect vitamin C to mail to him, bombarded the official Facebook page of the Samoan government with one-star reviews, and then started a GoFundMe campaign to “free Edwin” after he was eventually arrested for incitement against a government order.

Health officials—and not just in Samoa—need better ways of countering misinformation online. The posts that reach people on Facebook, YouTube, and other platforms aren't those with the most reliable information; they're the ones that have the most compelling memes, get the most likes, or are shared by influencers with large audiences. Elevating popularity over facts is dangerous during a disease outbreak. That's why, even before the outbreak in

Samoa, social-media companies had begun to take steps to tackle health misinformation differently from other types of misleading posts, which included making a policy decision to elevate content from the World Health Organization and the U.S. Centers for Disease Control and Prevention.

The paradox, however, is that the WHO, the CDC, and other leading health institutions—experts in real-world virality—have failed to adapt to the way information now circulates. Agencies accustomed to writing press releases and fact sheets for consumption by professional reporters are unequipped to produce the style and speed of information that the social platforms have made routine, and that the public has come to expect.

All too often, the people responsible for protecting the public do not appear to understand how information moves in the internet era. Meanwhile, people who best understand what content is likely to go viral are using that knowledge to mislead.



WE LIVE IN AN INFORMATION glut. The volume of content produced each minute exceeds the limits of human time and attention. Commanding a share of that attention has become a power struggle for states, media, and aspiring populist influencers alike. In the early days of social networks, users mostly saw the status updates, pictures, and other content that their friends posted—and in rough chronological order. But as the social sphere (formerly conducted offline) and the information sphere (formerly dominated by recognized institutions and news outlets) merged onto one internet infrastructure, they came to be governed by the same arbiter: curation algorithms. The creators who figured out how these algorithms worked became consistent winners. Older, somewhat sclerotic institutions paid little attention to the new dynamics, under which the ability to gain attention and shape perception depended on what we might call the “consensus of the most liked”: If enough people clicked on something, social-media outlets found it worthy of being pushed out, unsolicited, to still more people.

This was a problem with regard to health information. The algorithms were

bad at judging accuracy or authoritativeness. But Google, Facebook, and other platforms did not like the political, technical, and sociological complexities of building a notion of “authoritative information” into their ranking algorithms. For many years, the propaganda that anti-vaccine fever swamps pushed out at the first sign of an outbreak was, in Silicon Valley’s eyes, a matter of free expression. And if the anti-vaxxers’ content was more popular than the WHO’s, then so be it.

Read: Why the coronavirus is so confusing

But last year, things began to change. Six months or so before the crisis in Samoa began, Brooklyn reached the peak of a measles outbreak in which more than 650 people were infected and dozens were hospitalized. In response, the Senate held hearings about whether the growing frequency of preventable disease outbreaks had anything to do with the health misinformation proliferating on social networks. By this point, stories about misinformation and disinformation—about Macedonian spammers and Russian bots—had become a staple of media coverage.

Health misinformation was one of the easier problems for platforms to take a stand on. It had demonstrably negative downstream effects. In deciding how to surface health information, product managers and content moderators could look to scientific consensus—an option not available when judging political information. In mid-2019, the platforms rolled out new policies to mitigate the flood of anti-vaccine nonsense, downranking the grifters and conspiracists and upranking authoritative information. At least on health matters, the consensus of the most liked got a hard override: Queries about disease-related topics returned links to legitimate medical agencies, not to whichever anti-vax group had the most followers. Pinterest, Twitter, and Facebook began to direct users to the WHO and the CDC, both for queries related to specific disease outbreaks and for vaccine information in general.

Not long after the resumption of normal activity in Samoa, the coronavirus pandemic began. In late February, major social-media platforms’ policies against health misinformation were extended to content related to the coronavirus. But COVID-19, as it turned out, was a tougher problem.

Measles is a well-understood disease, with an effective vaccine, established treatment protocols, and overwhelming scientific consensus around a body of reputable, established facts; COVID-19 is not. This time, authorities and institutions themselves were learning on the fly. Early information coming out of China was unreliable—which didn’t prevent WHO officials from repeating it—and political considerations distorted the early public stances of global-health leaders and American politicians alike. But as it became clear that *this* disease was not going to remain confined overseas, global public demand for

information skyrocketed.

Peter Beinart: What the measles outbreak really says about America

The benefit of the information glut is that something novel always appears at the top of our Facebook and Twitter feeds. However, this engineered serendipity has come at a cost: We've become conditioned to expect engaging, up-to-date content *when we want it*. When a major crisis breaks—a mass shooting, a natural disaster, an emerging political scandal—we refresh and refresh, and usually see new information every time. But the COVID-19 pandemic is not like those crises. Under normal circumstances, rigorous scientific research about a new disease takes months or years. Reliable new information is, simply, slow to appear. Yet people keep searching—for information about symptoms, spread, treatments, fatality rates. In response, the algorithm returns *something*.

The feed abhors a vacuum. But in many cases, algorithms have little or no authoritative content to push to users—because experts haven't bothered to produce any, or because what they have produced simply isn't compelling to the average social-media user. Their work is locked in journals, while bloggers produce search-engine-optimized, Pinterest-ready posts offering up their personal viewpoint as medical fact. And with COVID-19, as in past outbreaks, anti-vaxxers and related influencers with a tenuous hold on reality jumped on the emerging topic early, posting repeatedly about synthetic-virus and mass-vaccination plots ostensibly hatched by Bill Gates and Big Pharma.

But around the same time, in late January, exceptionally prescient voices were also tweeting with increasing alarm, about very real risks. Reputable figures such as Scott Gottlieb, Donald Trump's former FDA commissioner, and Carl Bergstrom, a biologist at the University of Washington, used Twitter to walk the public through emerging research and information that seemed to indicate that, despite the reassurances of world leaders and health ministers, events were headed down a path that could be very bad indeed. Notably, these people presented evidence while also acknowledging its limitations. As accurate, up-to-date information began to seem like a matter of life and death, an increasing percentage of the public began to wonder if what elected officials, institutions, and the media were telling them was in fact correct.

DETERMINING WHO IS AN authoritative figure worth amplifying is more challenging than ever. Curated, personalized feeds enable bespoke realities. Trump supporters trust Fox News or One America News Network, while liberals follow a very different set of trusted sources. The legitimacy of media outlets is constantly questioned. Internet users have made collages of statements from mainstream publications that did not age well—for instance, early headlines and chyrons that could be interpreted as downplaying the threat from the coronavirus—and tweeted them out to dismiss the competence and quality of *all* mainstream media. Meanwhile, self-published Medium posts and tweetstorms by people with varying degrees of expertise—including none at all—regularly go viral. Some are highly accurate and well researched, deserve attention, and merit discussion; others are garbage pushed by grifters. The algorithm is responsible for deciding what, out of all this, to surface.

[Read: The billion-dollar disinformation campaign to reelect the president](#)

In the case of the coronavirus, the worst-case predictions of some of those early prescient voices on social media were borne out, even as some city leaders were telling the public—in March—to keep going out to theaters and the CDC was still insisting that only a very narrow group of Americans needed to be tested. Frontline doctors and scientists emerged in droves as the pandemic spread, posting on Twitter, Medium, and Reddit to tell the public about the number and severity of cases they were seeing in their hospitals; their stories further contradicted earlier reassurance from institutions that the flu posed a more serious risk to the United States. And so a meta-debate began: Why were social-media companies elevating the WHO and the CDC when some of their information turned out to be incorrect? And if agencies like these were wrong about COVID-19, what else were the so-called experts wrong about?

Populist Twitter decries any misstep by authority as confirmation of wholesale ineptitude or corruption—as if a mistake anywhere casts doubt on expertise everywhere. But these institutions did make a costly error with long-term ramifications for public trust: Rather than communicating transparently, frequently, and directly, explaining the distribution of probabilities and potential outcomes informing their guidance, they were reticent. Institutional medical authorities are bound by an ethical obligation to speak precisely and to hew to the facts—a constraint not shared by the Twitter and Medium commentariat. But when they finally do achieve a sufficient standard of confidence to make a statement, the pronouncement is often something that some faction on the internet has been insisting is true for weeks, so the authorities appear to be leading from behind. The CDC, which during the pandemic has largely operated in the background, is not structurally suited for

the communication environment in which it must operate.

One example has been the guidance on masks. In late January, the CDC said that, because of a lack of evidence of community spread, it was not recommending that Americans cover their face. By the time the agency changed its position in early April, #MaskUp had been trending for days. When people searched for information on the subject, every conceivable entity except the institutional authorities had produced content to fill the void.

[Read: The real reason to wear a mask](#)

Until the expert institutions adapt to modern means of communication—which they must do, and quickly, if they are to regain public confidence—a platform such as Facebook will have nothing compelling from them to show users. And yet simply deferring to the consensus of the most liked is also not a viable solution. Internet users need curation methods that ensure the visibility of *authoritative* voices even when that is not synonymous with *institutional* voices.

Some of the best frameworks for curating good information today remain those that involve a hybrid of humans and artificial intelligence: On Wikipedia, an army of volunteer human editors methodically records the facts while using bots to point out suspicious activity, and an arbitration committee—ArbCom—handles users who repeatedly make edits in bad faith. On Reddit, highly qualified moderators are curating coronavirus subreddits that offer substantive discussions about emerging research, while low-quality, misinformation-heavy subreddits have a warning label on them. Twitter has begun verifying the accounts of doctors and other science communicators, recognizing that channels beyond the official CDC and WHO accounts are providing highly useful, up-to-date information. These processes are difficult to scale because they involve human review, but they also recognize the value of factoring authoritativeness—not just pure popularity—into the way information is curated.

In the end, Samoa's measles outbreak was mitigated not by Facebook influencers, but by trusted public-health officials and experts who worked with the Samoan government to convey to the people directly affected why the vaccination programs mattered. During the 2014 Ebola outbreak in Nigeria, misinformation was rampant, but local public-health organizations used social media and mainstream media, and set up an SMS service, to debunk popular misconceptions. A year later, with Zika, Brazilian physicians and public-health officials got involved in WhatsApp channels to communicate with people on their preferred social platform.

The world is on the cusp of another high-stakes information battle: the one that will take shape surrounding the drug treatments and vaccines developed for COVID-19 over the next year. The consensus of the most liked would have us believe that Bill Gates and Anthony Fauci are preparing to track us all by microchip; countering those narratives as they continue to take hold, mutating slightly to appeal to specific online subcultures, will not be easy. Facebook, Twitter, and YouTube bear significant responsibility for the information environment for which they are hosts, curators, and amplifiers. But they can only do so much. If institutions and authority figures don't adapt to the content and conversation dynamics of the day, other things will fill the void. The time for institutions and authorities to begin communicating transparently is before wild speculation goes viral. Preventing epidemics of misinformation from spreading is easier than curing them once they've taken hold.

ABOUT THE AUTHOR

Renée DiResta

Renée DiResta, an associate research professor at Georgetown University, is the author of *Invisible Rulers: The People Who Turn Lies Into Reality*.

Explore More Topics

[Epidemics And Pandemics](#)