Student Perspectives on COVID-19 and Digital Technology

Course on Governing Digital Technology

Harvard Law School

Table of Contents

The American Government Needs to Improve Information Sharing with Private Sector Actors to Fight Coronavirus Disinformation

Casey Corcoran

In late April, in a disappointing show of irony, the European Union buckled to Chinese pressure and deleted an accusation in an official report that China was running a global disinformation campaign to deflect blame for the pandemic. The United States government has proved somewhat more resistant to such pressure, warning that Russian and Chinese actors are engaged in an effort to spread virus related disinformation. At the same time, however, the United States has neglected to share the evidence behind their accusation with at risk social media companies, effectively announcing that platforms face a formidable threat and then neglecting to help. This lack of transparency has caused confusion, blunting platform responses as they struggle to identify the coordinated state action. To combat state-run COVID-19 disinformation campaigns, the United States needs to focus both its messaging and its response, condemning state run disinformation campaigns and coordinating with private actors who have the most power to stop them.

State-run disinformation campaigns are often lengthy, meticulous operations designed to evade detection even from sophisticated detection algorithms. While platforms are skilled at identifying large inauthentic networks designed to gather a mass audience before infecting them with disinformation, they struggle to detect operations that gain traction by infiltrating pre-existing social circles and manipulating their influential members as proxies. In these cases, once the hazardous activity rises to a detectable level, the disinformation has already spread and it is hard to undo its effects. The key then is to catch the campaign early, but this relies on intimate knowledge of the state actor's plan in order to identify below-the-radar tactics before otherwise possible and attribute proxy behavior to a foreign source. This type of information is more accessible to a government intelligence officer than a company's information officer. Companies can take down malicious accounts but only if they know what to look for. The government must take the lead on finding and supplying this information.

So far during the current infodemic, the government has not yet stepped into this role, as evidenced by divided messaging among public and private sector actors. A recent Facebook report did not give any indication that it had detected a spike in virus-related state-run, coordinated, inauthentic behavior. And although a Chinese foreign ministry employee speculated in a highly publicized tweet that the United States may have started the pandemic, Twitter also reported that it had not noticed a significant coordinated campaign to spread coronavirus disinformation. At the same time, the U.S. State Department is insisting that such campaigns exists. While several Senators have been quick to criticize Twitter's refusal to take down the Chinese official's tweet, more need to follow Senator Cory Booker's example and press government agencies to increase information flow to private actors. It is more important

to give companies access to the data they need to target state directed operations than to pressure them to take down a single post.

China's reported recent disinformation campaign's heavy reliance on text messages and encrypted messaging platforms provides an example of exactly why public, private cooperation is necessary. China's alleged tactics frustrate traditional content moderation efforts because the encryption and privacy norms do not allow carriers or platforms to analyze messages for false information. Stopping the spread of disinformation on encrypted platforms has traditionally involved modifying features commonly used by coordinated actors, while mobile carriers most analogous challenges has been flagging and blocking spam robo-calls. Both of these defenses rely on intimate knowledge of how an adversary is operating and taking measures closely tailored to mitigate that behavior. The private sector must ultimately find innovative ways to reduce disinformation on mediums where content moderation is impossible, but government coordination to show them exactly how adversaries are operating is critical to their success.

Of course, the government has good reasons for being wary of sharing too much information. The head of the Global Engagement Center, Lea Gabrielle, has argued that too much openness could have the perverse effect of aiding adversaries as publicizing intelligence-based reports runs the risk of compromising sources and methods. Separately, unrestrained cooperation between the national security apparatus and the companies who oversee the nation's communication infrastructure could be abused to degrade civil liberties or chill legitimate political speech. However, the government still must find a way to share sensitive information with key private sector actors. Identifying and stopping state-run COVID-19 disinformation campaigns requires knowing what behavior to look for, and right now the government alone seems to know what that is. Knowledge is not helpful if it cannot be used to empower those who need to act.

Government Coronavirus Information Isn't Just Contradictory: It's Overwhelming, Confusing, and Inaccessible

Nancy Fairbank

In the face of a rising tide of COVID-19 misinformation, the U.S. government has failed to consistently provide clear and convincing information to the public. In part, this lack of clarity has arisen from ongoing public clashes between the President and experts (including other government officials). This lack of unity in messaging leaves civilians in a lurch, especially when it comes to making judgments on the basis of complex scientific and medical information – it's hard enough to pronounce hydroxychloroquine without also having to unpack the nuances of the dispute over the drug's usefulness and safety.

But it's not just that government messaging around the pandemic has been contradictory and volatile. Perhaps even more critical is the fact that much of the authoritative information currently in circulation is presented in ways that are inaccessible to laypeople. It's hard for high-quality information to compete with flashy falsehoods when it's confusing and dense. For example, take this CDC Tweet, which specifically indicates it's for laypeople like me who want to learn how to protect themselves and their families:

Instead of finding a list of helpful safety tips, people who click the Tweet's link are surprised by a page full of paragraphs like this one, filled with cohort percentages that require far more interpretation to be made usable and accessible for regular folks:

> It is not yet known whether the severity or level of control of underlying health conditions affects the risk for severe disease associated with COVID-19. Many of these underlying health conditions are common in the United States: based on self-reported 2018 data, the prevalence of diagnosed diabetes among U.S. adults was 10.1% (*7*), and the U.S. age-adjusted prevalence of all types of heart disease (excluding hypertension without other heart disease) was 10.6% in 2017 (*8*). The age-adjusted prevalence of COPD among U.S. adults is 5.9% (*9*), and in 2018, the U.S. estimated prevalence of current asthma among persons of all ages was 7.9% (*7*). CDC continues to develop and update resources for persons with underlying health conditions to reduce the risk of acquiring COVID-19 (*10*). The estimated higher prevalence of these conditions among those in this early group of U.S. COVID-19 patients and the potentially higher risk for more severe disease from COVID-19 associated with the presence of underlying conditions highlight the importance of COVID-19 prevention in persons with underlying conditions.

Here's another CDC Tweet, which gives a cute little #DYK? hashtag – along with content that seems both meaningless and anxiety-inducing for most laypeople:

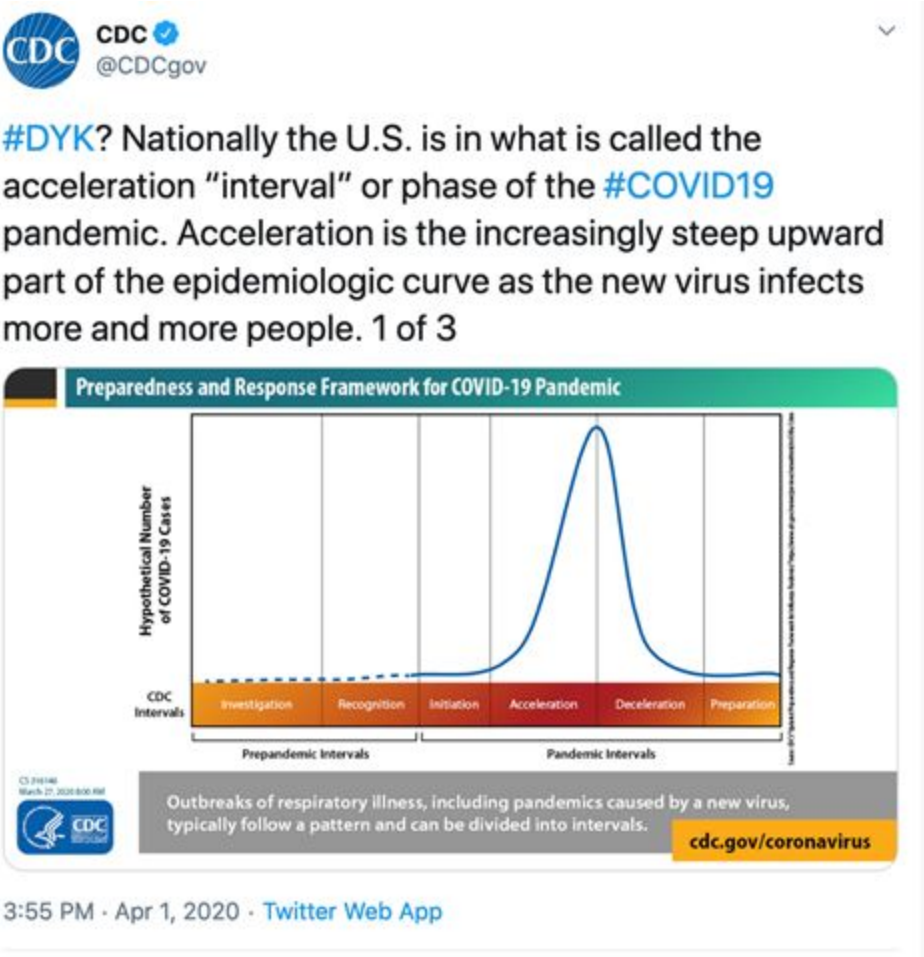

This same worryingly vague and vaguely worrying 'Preparedness and Response Framework for COVID-19 Pandemic' graph was tweeted out again by @CDCgov on April 3rd, and @CDCDirector tweeted it three separate times over a three-day period, on April 1st, 2nd, and 3rd.

States are having similar issues. For example, when I initially visited the Arizona Department of Health webpage, I got this bundle of stats and graphs before I was ever presented any information on what to do if I feel sick, or a button for the 'What Everyone Needs To Do' page. (The page has since been updated.)

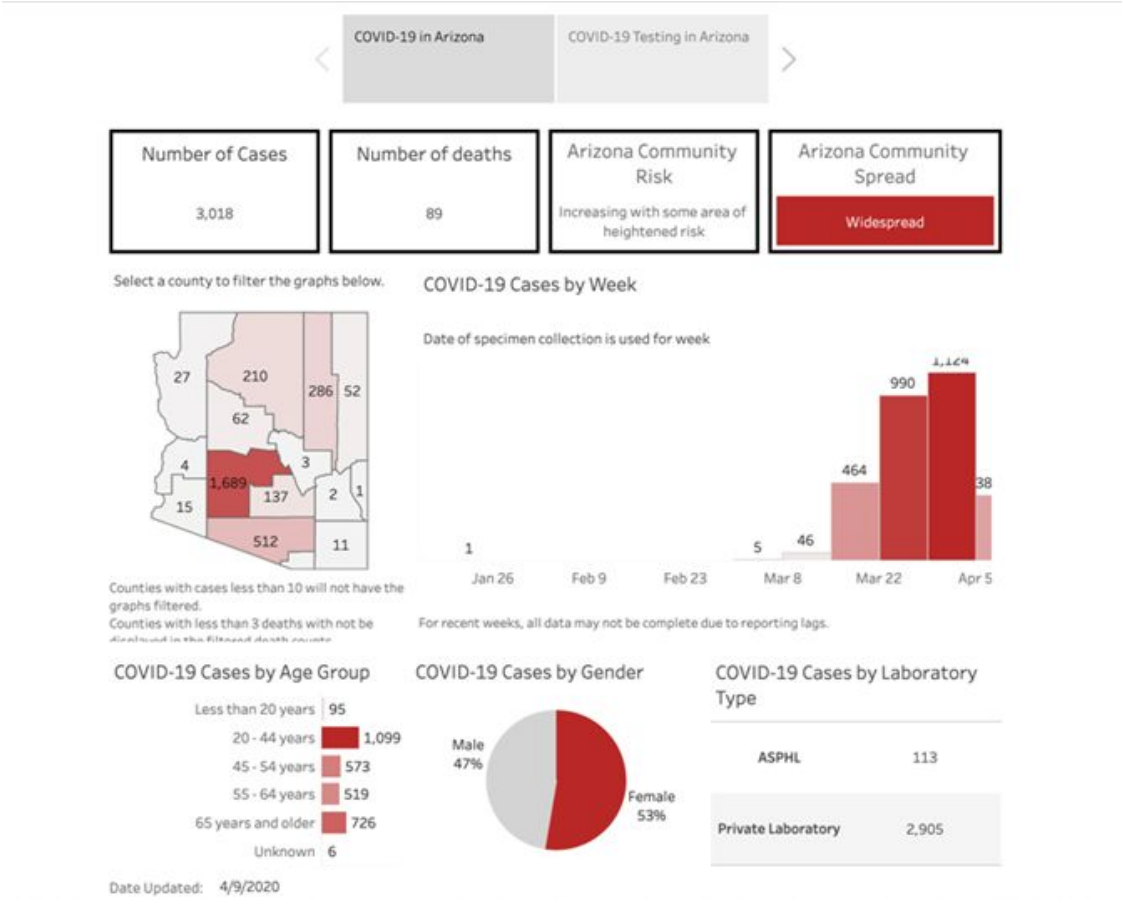

The Nevada Department of Health site also buries general information for the public under comparable stats and graphs:

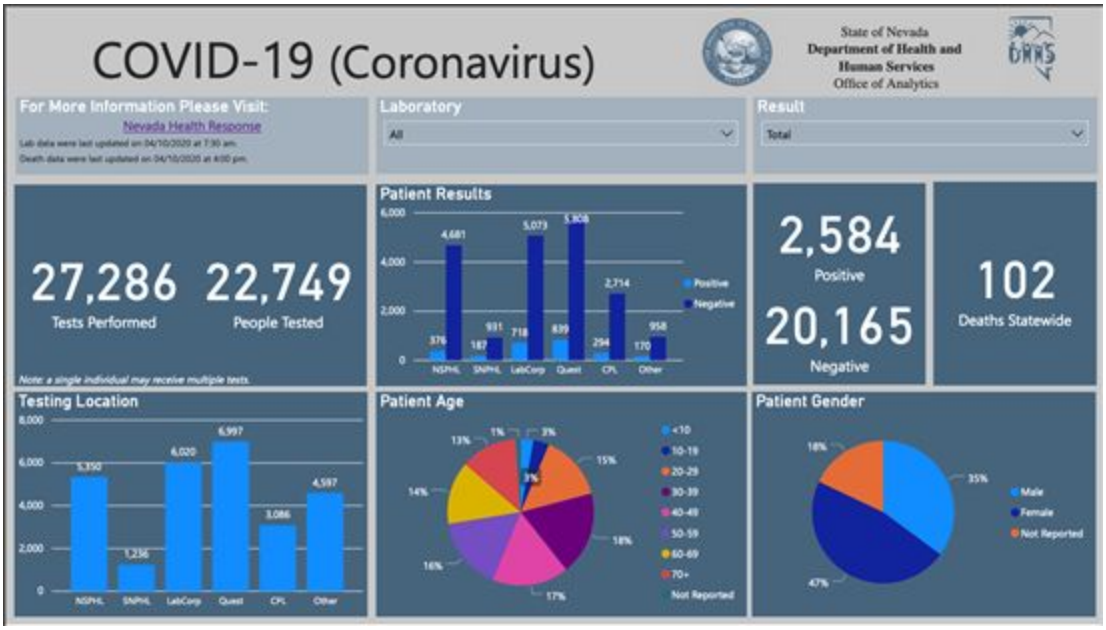

Much of this data is either irrelevant to most laypeople or may actively cause confusion (for example, a layperson may think that a higher case rate in women means coronavirus infects women more easily – something that can't be safely assumed by correlation alone).

And when I clicked on the 'Home Care' link of the Nevada site, I was not met with basic information on how to care for a sick household member. Instead, there was a collection of links, the first of which led to a technical bulletin requiring healthcare providers to report coronavirus deaths – not exactly comforting for someone seeking information on caring for their infected loved one at home.

Nevada Department of Health and Human Services
DIVISION OF PUBLIC AND BEHAVIORAL HEALTH

Technical Bulletin

Date: April 08, 2020
Topic: Mandatory Reporting of All COVID-19 Mortality
Contact: Melissa Peek-Bullock, State Epidemiologist, Office of Public Health Investigations and Epidemiology
To: All Health care Providers and Facilities, Medical Examiners, Coroners and Funeral Homes

Effective immediately, all providers of health care, health care facilities, medical examiners, coroners and funeral homes must report immediately to their local and state health department, each death that is caused by, associated with or related to COVID-19 infection consistent with the CDC criteria outlined below.

In sum, much of the government information currently on offer is either (1) highly technical, (2) vague and inapplicable to civilians, or (3) useful but buried underneath the highly technical and vague/inapplicable. By contrast, many forms of coronavirus misinformation are easy to understand, presented simply and directly, and offer a much greater degree of certainty to fearful citizenry seeking clarity in the midst of a global crisis. Instead of throwing a bunch of content at readers and causing information overload, many of these misinformation messages are just a few sentences long and focus on a single aspect of the pandemic. For example, this coronavirus misinformation meme is just three sentences (including the one telling you to share it with others), with absolutely no jargon or statistics in its 'medical' explanation:

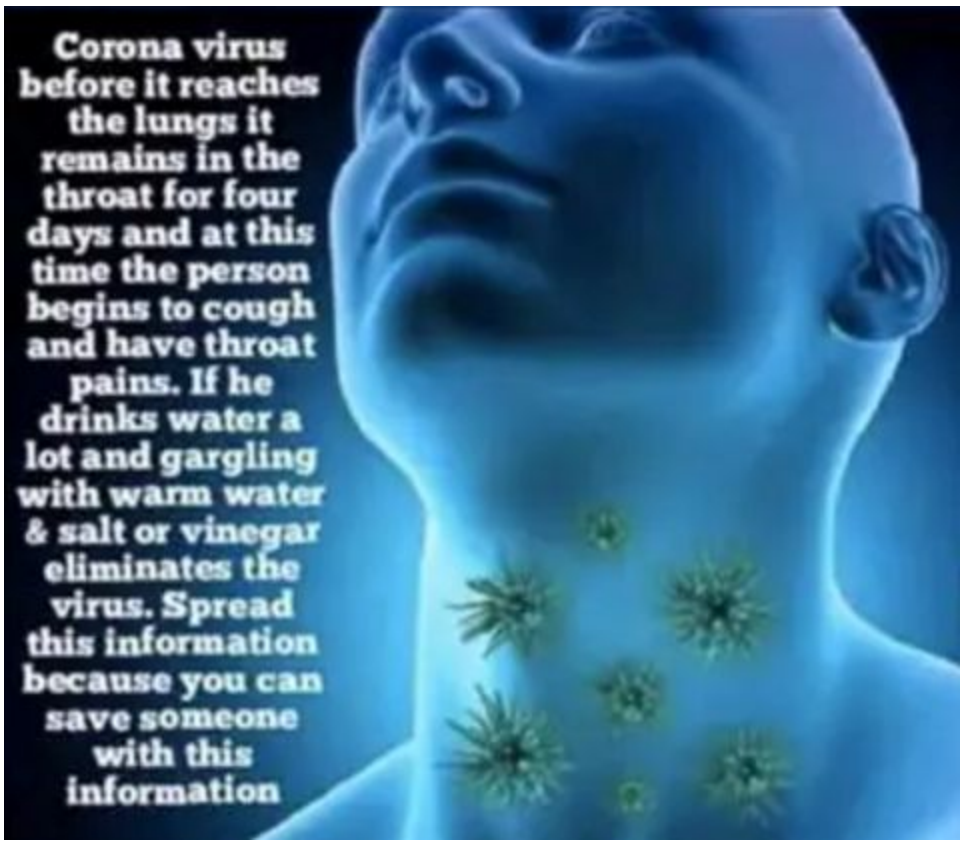

This misinformation is simple, direct, and easy to understand and apply (though, of course, false). While a prevalent response by online platforms combatting the infodemic has been to redirect users interacting with misinformation to more authoritative sources, it's questionable how useful this strategy will be if official sources provide content that's unclear and inaccessible to laypeople.

That's why it's critical that government focuses on presentation and not just content. For example, studies on combatting vaccine misinformation indicate that emotionally-framed information is better transmitted than medically-framed information. Government messaging incorporating emotional appeals rather than just laying out statistics may increase public

cooperation with health regulations. NY's Governor Cuomo has done an excellent job of implementing emotional framing, such as with this Tweet accompanying a video he posted:

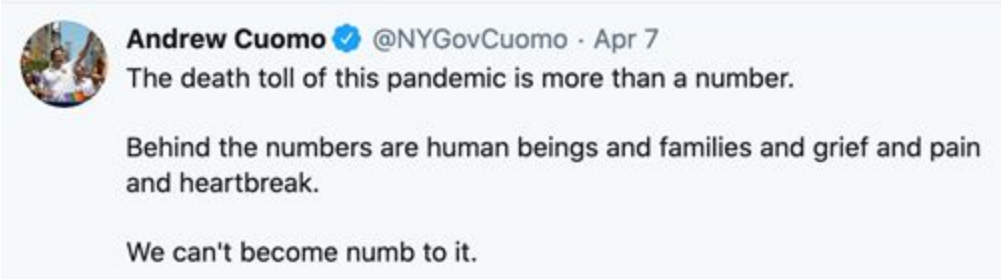

Officials can also utilize the framing effect, which finds that people are more likely to avoid risk when a positive frame (ex. potential lives saved) is presented than when a negative frame (ex. potential lives lost) is used. This concept is demonstrated well in a recent Tweet by the Big Cities Health Coalition:

Finally, agencies should follow the CDC's own guidance on simplifying messages in times of crisis. This means ensuring information targeted toward regular civilians is free of scientific and medical jargon, and is made as short, direct, and actionable as possible. This will likely require creating more clearly separated channels for public-focused and expert-focused communications. At the minimum, basic health and safety information should be at the top of official webpages rather than hidden underneath information intended for experts (experts will go fishing: the public won't). Providing information alone is not enough – it must be provided in a way that's accessible to make the most difference.

To Protect Users from COVID-19 Misinformation, Platforms Need to Focus on Mental Health

Grant Fergusson

If you're like a majority of Americans, the COVID-19 outbreak has been a major source of stress and anxiety. Your stress may come from a variety of factors—the health and wellbeing of your loved ones, social isolation, and job loss among them—but the result is the same: COVID-19 is negatively impacting our mental as well as physical health.

Thus far, digital platforms have tailored their efforts to reducing the risks of physical harm associated with the spread of COVID-19 misinformation. Those risks are real and deserve the heightened attention they're receiving. However, COVID-19 misinformation causes more than physical harm. By spreading counter-narratives and conspiracy theories, COVID-19 misinformation stokes users' confusion and anxieties in ways that increase distrust and paranoia—and may incite violence.

Why do so many people believe COVID-19 misinformation? It's not just because it causes physical harm. Like other forms of misinformation, COVID-19 misinformation entices users with convenient explanations for the pandemic and provides them with a sense of security and control that factual information can't. And unlike other forms of misinformation, COVID-19 misinformation has a uniquely susceptible audience. High levels of stress and anxiety appear to predict one's propensity to believe conspiracy theories, and the COVID-19 outbreak has seen stress levels skyrocket across populations. This is especially true for people who spend more time on social media—in part because of what platforms have done to increase user engagement. As people increasingly turn to these platforms for their information and social interactions during the pandemic, platforms can no longer ignore the mental health impacts of their products and policies. The pandemic is a mental health crisis as much as a physical one. To protect users from the harms of COVID-19 misinformation, platforms will need to consider more active steps to improve and maintain the mental health of their users.

What does this mean, exactly? First, platforms need to consider how COVID-19 misinformation harms their users' mental health. Misinformation—particularly misinformation involving conspiracy theories—effectively stokes paranoia and distrust among users, driving them toward more extreme and isolated social niches. Theories based on misinformation also tend to be self-reinforcing and resistant to challenges; those who hold beliefs based in misinformation tend to feel distress when faced with factual information that challenges their beliefs. And even when conspiracy theories don't risk imminent physical harm, they can nonetheless incite believers to act violently. While some platforms have already taken steps to combat incitement caused by COVID-19 misinformation—Facebook and YouTube have taken down posts related to conspiracy theories linking 5G networks to COVID-19—many platforms have yet to focus directly on the mental health harms of COVID-19 misinformation. To protect users from dangerous misinformation in the short-term, platforms should expand the kinds of harms they consider

when determining whether to flag and remove COVID-19 misinformation to include the risks of mental health harms to users—either by exacerbating existing levels of anxiety and depression or by stoking paranoia and cognitive dissonance.

Second, platforms need to broaden the scope of their COVID-19 responses to include preventative—not just proactive—measures. While many platforms have implemented a variety of machine-learning algorithms to identify, flag, and remove harmful COVID-19 misinformation before it spreads, platforms have done little to reduce user susceptibility to misinformation. Because users with higher levels of stress and anxiety tend to accept misinformation more readily than their less stressed and anxious counterparts, strategies designed to improve user mental health and well-being appear to be effective ways to build user resistance to COVID-19 misinformation before they see it. Some platforms have already taken steps in this direction: in March, for example, Snapchat introduced a new, in-app mental health tool, "Here For You," designed to reduce anxiety related to the coronavirus. Still, platforms can do more to improve their users' mental health in both the short- and long-term. Like TikTok, platforms could implement default measures meant to remind users to take breaks and reduce their screen time. (Facebook and Instagram have introduced similar, optional features already.) Platforms could also incorporate new algorithms meant to direct users toward positive and health-reinforcing content—a capability for which Facebook previously received backlash, and one similar to that being implemented by YouTube as a part of its broader battle against conspiracy theories. Lastly, platforms could partner with mental health experts, state actors, and telehealth service providers to increase access to mental health resources and counseling on their platforms, whether through official hotlines and services or by connecting people to online peer support groups.

As we all settle into our new, homebound routines, the risk of COVID-19 related physical harm has declined, but for many, the risk of worsening mental health has only grown. While some platforms have openly discussed their attempts to reduce the mental health impacts of stay-at-home orders or seen their users provide mental health support services to each other, many have yet to consider an official mental health response to COVID-19—even as their users' mental health declines—and no platforms have published policies to address how COVID-19 misinformation may exacerbate existing mental health problems among their users. Whether caused by COVID-19 misinformation directly or by our homebound environments, we are facing a mental health crisis, and so we need to ask: are digital platforms doing enough to protect their users' mental health?

My answer is no.

A Field Guide for the COVID-19 Infodemic: Where to Find Good Data and Guidance

Sasha Peters

It's no revelation that the COVID-19 pandemic has led to much uncertainty as to how many in the U.S. have been affected. In an ideal world, there would be an accessible, centralized, thorough, and accurate source of data covering test counts, deaths, probable cases, disparate impacts on racial and ethnic groups, state policies, testing centers, and predictions of how the virus will spread, among other things. Of course, no such source exists: while the CDC website should ostensibly be the greatest resource for information about the pandemic, the CDC has been slow to gather and publish this data.

Luckily, non-governmental parties have stepped in to give the most accurate information possible on their websites, providing users with the information they need to make informed decisions about where they can go, what they can do, and why they should do it.

That said, the lack of government leadership in providing resources carries serious downsides. First, these third parties cannot compel states to provide consistent and thorough information across the board, and since we rely heavily on state reporting, data on these sites are likely poorly contextualized underestimations. Second, the lack of a centralized hub for this data inevitably means that accurate data is scattered around the internet amongst mis- and disinformation, and it can be difficult to tell which is which. The former problem almost certainly requires government intervention, but the latter problem may be addressed in part by gathering and bolstering the visibility of trustworthy sources of information. Below are some of the most accurate and up-to-date websites providing COVID-19 data and resources.

DATA

TRACKING CORONAVIRUS CASES

Johns Hopkins University: COVID-19 United States Cases by County

This interactive map from Johns Hopkins is the most reliable tracker, and the one that most news organizations use. The U.S. map includes test counts and results, confirmed cases, and deaths down to the county. There's also an international version, subdivided by country.

Data sources: "The data sources [for both the U.S. and world maps] include the World Health Organization, the U.S. Centers for Disease Control and Prevention, the European Center for Disease Prevention and Control, the National Health Commission of the People's Republic of China, 1point3acres, Worldometers.info, BNO, state and national government health

departments, local media reports, and the DXY, one of the world's largest online communities for physicians, health care professionals, pharmacies and facilities."

How often it's updated: The U.S. map is updated "once per day after 8 p.m. Eastern to allow the system to pull county-level data." The world map is updated in real time.

The COVID Tracking Project

The COVID Tracking Project, created by *The Atlantic*, tracks COVID-19 testing and outcomes in the U.S based on state and other local reporting. The COVID Tracking Project has also assigned a letter grade to each state and territory showing how reliable state and local information is, based on sixteen different factors.

Data sources: "All our information comes from state/district/territory public health authorities—or, occasionally, from trusted news reporting, official press conferences, or (very occasionally) tweets or Facebook updates from state public health authorities or governors."

How often it's updated: "We update the full dataset each day between 4pm and 5pm EDT, with limited updates as new information arrives."

IMPACT ACROSS RACIAL AND ETHNIC GROUPS

While we know that COVID-19 has disproportionately affected people of color, especially African Americans, the CDC's website did not start showing data by race until April 8, in response to political pressure to do so. However, that data is incomplete, in part because reporting of racial data varies across states; the latest CDC data shows that racial data is missing in at least 55.4% of cases.

Johns Hopkins Coronavirus Resource Center: Racial Data Transparency

Johns Hopkins has released a visualization showing state-level data by race in three categories: confirmed cases, deaths, and testing. Not all states are collecting or providing racial data.

Data sources: State reports.

How often it's updated: As state reporting policies change.

COVID Racial Data Tracker

The Atlantic's COVID Tracking Project has collaborated with the Antiracist Research & Policy Center to create the COVID Racial Data Tracker, which compiles testing and death data by race and ethnicity across states. The COVID Racial Data Tracker also offers a state-by-state report card which outlines the quality of racial data provided by state officials.

Data sources: Typically state reports, although "data sources remain in flux."

How often it's updated: Frequently—the time of last update appears on the data spreadsheet.

PREDICTING CORONAVIRUS

R_t Covid-19

Started by two Instagram founders, this site shows state-by-state data for R_t, which measures how fast the virus is growing. If R_t is above 1, the virus is still spreading; if it is below 1, the virus is slowing down. Keep in mind that the data assumes that infectiousness begins with symptoms; as such, the site suggests that "a simple heuristic is to shift all values of Rt 5 days into the past."

Data sources: The site gets its case count information from the COVID Tracking Project. You may download a Jupyter notebook showing how R_t is calculated here.

How often it's updated: Not stated, but it appears to be updated as the COVID Tracking Project is updated.

EpiForecasts.io: National and Subnational Estimates for the United States of America

This site identifies "changes in the reproduction number, rate of spread, and doubling time during the course of the COVID-19 outbreak whilst accounting for potential biases due to delays in case reporting both nationally and subnationally in the United States of America."

Data sources**:** "For sub-national analyses, the source of the data is reported on each page, the data are fetched from government departments or from individuals who maintain a data source if no official data are available. Regions within countries are only reported if at least 40 cases have been reported in a single day." The data is also adjusted for reporting delays.

How often it's updated: "[R]egularly as new data becomes available."

STATE SHUT DOWN/REOPENING STATUS

The New York Times: See Which States Are Reopening and Which Are Still Shut Down

The New York Times regularly updates this interactive map, which shows the policies in place in each state, the expiration dates of those policies, and the types of businesses that are open or slated to reopen.

Data sources: "State-level case data is from a New York Times database of reports from state and local health agencies and hospitals."

How often it's updated: "Regularly."

RESOURCES

INFORMATION AND MISINFORMATION

First Draft: Find information about coronavirus (Covid-19)

First Draft, a non-profit aimed at providing best practices for accurate and ethical journalism, has created a database of coronavirus information and misinformation across more than twenty languages. A search for "ibuprofen," for example, yields ten articles in four different languages debunking the myth that ibuprofen is harmful to those who have coronavirus. Users may also submit information they have received regarding coronavirus for First Draft to investigate.

Data sources: Information is drawn from the World Health Organization (WHO), Google's Fact Check Explorer, Poynter's International Fact-Checking Network (IFCN) and First Draft's trusted partners.

How often it's updated: Not stated.

WHERE TO GET TESTED

Location of Drive-Thru COVID19 Testing Sites across America

Jim Kyung-Soo Liew, an expert on big data and artificial intelligence and a professor at Johns Hopkins Carey Business School, created this project with help from his software company and student volunteers from Carey.

Data sources: "[V]olunteers search for drive-through kiosks across the U.S. by perusing online local news sites and input that data into a form. The form populates a database, and the data

renders up on our site." Visitors to the site may also report new testing locations through a form on the site.

How often it's updated: Not stated.

WHERE TO VOLUNTEER

VolunteerMatch

This site offers both virtual volunteer opportunities and in-person opportunities that may be narrowed by location.

Data sources: Submitted by volunteer organizations.

How often it's updated: Not stated; likely as admissions come in.

Mutual Aid Hub

This map gives locations and contact data for mutual aid networks around the U.S. and offers resources for starting a mutual aid network. The networks on this site are tailored specifically to providing local assistance during the pandemic.

Data sources: Submitted by volunteer organizations.

How often it's updated: Not stated; likely as admissions come in.

Coronavirus: The Geopolitical Blame Game

Katherine Shen

The "fake news" and propaganda campaigns that plagued the 2016 US elections catalyzed much of the existing research on misinformation, disinformation, and mal-information. At first glance, studies of the 2016 presidential election seem to have little relevance to the 2020 coronavirus infodemic: politically motivated disinformation is primarily the product of Russian interference, whereas COVID-19 misinformation seems to be a natural response to the rapid spread of a newly-discovered, infectious disease. "Miracle cures", fake tests and lock-down rumors do not seem to have a distinctly political lean because they spread via pandemic-related fear and not through partisanship. However, the boundary between propaganda and truth has become increasingly blurred as governments, news outlets, and politicians manipulate COVID-related news to further self-interested political narratives. Disentangling politics from coronavirus disinformation is absolutely critical to controlling the outbreak and moving into the post-pandemic world.

Regardless of its content and source, ideologically-driven coronavirus disinformation tends to be structured around a fairly consistent set of goals: to sow panic and widen domestic political divisions in another country, incite suspicion of foreign governments, and reinforce narratives about "good" vs. "bad" actors. In light of these shared objectives, it may be more useful to distinguish between different types of coronavirus disinformation by considering whether a story is intended to affect domestic politics (inward-looking) or international relations (outward-looking). Coronavirus disinformation has mapped onto the existing political divisions within the U.S. primarily in the form of diverging attitudes towards lockdown policies and the administration's management of the pandemic. COVID conspiracies are new in content but familiar in form: coronavirus disinformation is spread through existing media structures and exacerbated by the political polarization that has existed long before the pandemic.

Coronavirus disinformation not only has a serious impact on domestic politics, but has also amplified existing geopolitical tensions and possibly reshaped the global order. Conspiracy theories about the origins of the coronavirus are used to frame foreign states (or other "outsiders") as the villainous progenitors of the virus and enemies of a pseudo-war. For instance, the Trump administration has claimed that the coronavirus escaped from a Chinese lab while Chinese officials have promoted a rumor the US army brought the virus to Wuhan. These conspiracy theories tend to encourage nationalist sentiments, which by extension legitimize both defensive and offensive governmental actions. Rhetoric that stresses the dangers posed by "enemies abroad" galvanizes public support for the adoption of protectionist policies and restrictions on immigration. Stories about government mismanagement of the outbreak serve a dual function: first, they highlight the incompetence of foreign government(s), thereby implying that domestic officials are handling the outbreak better than their counterparts abroad; second, accusing a foreign state of inaction, intentional or otherwise, shifts the blame to external actors. Occasionally, stories about mismanagement of the outbreak go further than

accusations of incompetence and negligence and suggest that the foreign country allowed the virus to spread deliberately. By implying that a foreign State purposely allowed coronavirus to spread, (mis)management narratives dove-tail with conspiracy theories about the virus' origins, seamlessly blending mal-information and disinformation.

Nationalistic disinformation campaigns are a cause for concern because they legitimize government policies and actions that ought to be subject to greater public scrutiny. Coronavirus provides the opportunity for political actors to stretch the Overton Window; for instance, the state-of-emergency has empowered Congress to push through drastic, unprecedented policies. Media and news outlets usually ensure that policies implemented during a crisis remain within acceptable democratic bounds by encouraging public oversight of government actions. However, disinformation campaigns erode the mechanisms for checking State power because they undermine norms-constrained media's check on the dissemination and validation of extremism, thereby affirming the post-crisis Overton Window. The 2020 infodemic has shown that politicians play a major role in spreading disinformation but more alarmingly, the COVID crisis shows that politicians are ready and willing to disseminate political narratives that bolsters their chances of re-election and facilitates the consolidation of power. The spread of propaganda and "fake news" in the 2016 election has led to calls for a heightened bar of professional responsibility for journalists. Perhaps, politicians and government officials should be held to a similar, if not greater, standard for transparency and accountability when they are sharing information and news.

Finally, it makes sense to pause and consider the human impacts of coronavirus disinformation. There is a surprising connection between international narratives and local violence – stories about China's mismanagement of the initial outbreak creates animosity towards the country but also leads to an increase in hate crimes and racism against Asian Americans and Asians residing in the U.S. More often than not, disinformation harms minorities, immigrants and refugees because their socio-political status renders them more vulnerable to violence and discrimination. Governments that are intentionally leveraging coronavirus myths and conspiracies to further their foreign policy goals are also hurting their own citizens, residents, and displaced peoples. While it may be tempting to see the coronavirus as yet another piece in the grand game of geopolitics, the costs of the pandemic are and will continue to be deeply human.

Information Rights in Global Crisis

Parvuna Sulaiman

In response to the COVID-19 pandemic, governments around the world are taking unprecedented measures to control the spread of misinformation surrounding the disease. As nations like Russia and China enact laws to silence members of their own citizenry who repeat popular rumors about potential cures and methods of transmission, they are also fielding accusations of disinformation-spreading from across the international milieu. These trends are not novel in themselves. Misinformation laws are an established part of state toolkits, and countries like the US and Russia have already engaged in the fierce exchange of disinformation accusations with respect to other topics. But in the heat of the pandemic response, we are seeing how false narratives can lead directly with clear attribution to needless deaths. People, particularly vulnerable populations like the elderly, are so afraid of contracting the disease that their usual skepticism about rumors might be temporarily disabled. With these challenges mounting, the American government has faced calls to adopt at least some of the stringent methods used by China and similar nations to control the flow of information and to actively silence those who spread information inconsistent with such ground truth.

Though our collectively rising anxiety might make such options seem very attractive, even necessary, we have to recognize that taking a turn towards pervasive censorship could pose a serious threat to individual rights. If we were to adopt such measures, we would have to understand what they would mean, not just in a time of crisis, but for our everyday lives as people under one of the most powerful states in the world. In order to decide whether or not we want to give the state this set of powers, we first have to decide two fundamental questions about information rights:

> *1) Does the individual have a right to be actively supplied with true information?*
>
> *2) Is there a right against being deceived?*

In the face of what the World Health Organization has called an "infodemic" surrounding COVID-19, a right to be actively supplied with true information sounds particularly suitable and attractive. Assuming that the public trusts the state, a system like this would allow each person the ability to trace information back to a recognizable and trusted source before acting on it. But as a rights theory, the right to true information is nebulous. The right doesn't seem to come with any obvious qualifications beyond impossibility and privacy invasion. Do I have a right to a daily personalized report from the state about any non-sensitive information, like the movement of the stock market or every licensed airplane? Think about the endless resources that would consume. If we were to entertain this kind of information right, we would have to qualify it with so many limitations that we might not recognize it. I think that there is a more elegant framing.

Instead of starting with a right to information and trying to chop it down to what we can manage, we could instead take as our starting point the fundamental bundle of rights that involve the state's responsibility to protect your life and body. In the unique situation of a pandemic, the state can't use its traditional tools for preventing injury or death to its public. Information is now absolutely necessary for any viable defense of the individual's life. We all need to know how the disease spreads and what mitigation efforts work in order to survive this. The circumstances activate the government's duty to provide reliable, up-to-date information in a manner accessible enough to meaningfully afford protection to the individual through knowledge. On this model, the state isn't afforded control over the evolving scientific narrative of the virus; it is charged with the responsibility to consult the scientific community as it performs the same core functions it has always performed.

This still leaves us with the problem of misinformation, and for that we should consider a right against being deceived. Misinformation laws generally pose an interesting question for the right to freedom of speech, a core concern in American society. I would never contest the right to free speech, but I would qualify it without hesitation. My freedom to speak as I choose stops short of deceiving you in a way that manipulates you against your own interest. In order for a society of rational beings to live together in peace and prosper as individuals, we have to create legal systems that protect each person's rights to live free from non-retaliatory coercion. When I deceive you by selling you a fraudulent cure for a deadly virus, for example, I am using a form of coercion against you. I'm using you for my own ends and to your detriment without your informed consent, even if you don't realize I've done it. The state already has a responsibility to stop the commission of fraud, as well as other crimes carried out through dishonest manipulation. The right against being deceived is a familiar concept, and it only needs minimal tuning to adapt for the COVID-19 pandemic.

We can separate misinformation into two categories: intentional and unintentional. Intentionally spreading misinformation violates the audience's right against being deceived, and therefore the speaker can and ought to be silenced by the state. Unintentional misinformation is trickier. So many of us have heard inaccurate but well-meaning advice from relatives who genuinely believe that eating garlic is a real cure. These people aren't actively deceiving anyone. They are in most cases suffering under the weight of this pandemic, and the application of state force to their cases in the manner adopted by countries like Russia feels hollow and brutal. I would suggest that we treat this kind of case as beyond the scope of one's right against deception, and leave it to the realm of our pandemic-specific right to be supplied with protective truths. Rumor debunking on government websites is one gentle, uniform, and accessible means of deflating the fear that occasions the spread of these rumors.

Given the terrifying spread and devastation of COVID-19, we might be tempted to temporarily opt into the stringent state controls on information that are gaining popularity in other parts of the world, and deal with the consequences once the public health is secured. But invasive government powers like these don't readily fit back in the box after we take them out, and there is a way to control misinformation without compromising our safeguards for individual rights. I

suggest that we combine a right to be supplied with protective truths and a right against being deceived as a two-pronged strategy for controlling the COVID-19 narrative in a just way. The nuances here might seem like an unnecessary consideration in terrifying times. But when we carefully assign the powerful government certain affordances to address a crisis, we can ensure that the law doesn't later morph into an instrument for manipulating the public information exchange or trampling individual speech. As much as we may need order to maintain our societies against pandemics and infodemics alike, we must work creatively within our existing precommitment strategies for human rights protection if we want those societies to be worth maintaining.

A Plea for Platform Transparency

Adam Toobin

To preserve a diverse democratic discourse, major platforms need to clarify how they are regulating speech online.

The clearest challenge that platforms face in adopting proactive "public health" policies to stem the spread of misinformation comes from the inevitable confrontation they will face with those who view their unilateral approach to speech regulation as coercive and partisan. The highly segmented media environment in the United States combined with the politically charged nature of speech regulation of pandemic misinformation has already brought platforms into conflict with conservatives who view these decisions as partisan. As the United States enters the fourth month of a pandemic crackdown that is not expected to end entirely for many months, these tensions are likely only to escalate further. Without a meaningful shift in the degree to which platforms are transparent about how they regulate speech online, the political consensus that supports the status quo may yet fracture, leaving decisions about private speech regulation to the political process.

The controversy over the role the platforms are playing in speech regulation comes at a time when platforms are adopting a new, more comprehensive approach to policing misinformation. Facebook is not only fact-checking more content but sending corrections to users who interact with misinformation. WhatsApp has limited the ability to forward messages to more than a single other user. YouTube is introducing fact-checkers to the United States for the first time. And Twitter announced new labels that it would apply to content it deemed to contain false or misleading information. The Twitter decision is notable, because it almost immediately resulted in a label being applied to a tweet by President Donald Trump that involved a separate issue.

On an industry-wide basis, there are also collective trends that suggest a broader shift in Silicon Valley's way of thinking about its role in monitoring speech. All the major platforms are collaborating in a variety of ways to identify misinformation. They are also providing "verified" information to their users. And they are relying more heavily on the use of artificial intelligence to make individual content decisions. As a general matter, these extensive responses to the challenge of pandemic misinformation — whether they have been adequate or not — suggest the emergence of a new paradigm in favor of speech regulation that is already on a collision course with those opponents who view it as partisan.

These interventions have led to new confrontations with conservatives and other members of the right who have interpreted new efforts to combat misinformation as further veiled attempts to censor ideas that run against prevailing sentiments. Notably, Facebook's decision to take down event pages promoting protests against social distancing policies was interpreted by

many as politically motivated. Moreover, the fact that platforms are empowered to make important decisions about speech regulation unilaterally and without clear public oversight does raise real concerns, particularly about whether platform decisions could influence a close election. As platforms continue to adopt a more active role in regulating speech online, these questions are becoming more urgent, and more closely tied to debates about whether and how to change the permissive legal regime fostered by the Communications Decency Act Section 230.

Under CDA 230, platforms have been able to develop their own regulatory mechanisms without direct external accountability. For some, particularly on the left, CDA 230 has made it almost impossible to punish platforms that do allow harmful or abusive content to spread on their sites — the effects of which tend to fall disproportionately on underrepresented communities. For critics on the right, the law has enabled a private system of speech regulation that has empowered Silicon Valley magnates to dictate controls that suit their political agendas. Reconciling these perspectives may not be entirely possible, but interventions that advance transparency at least advance the underlying objectives on both sides of the debate.

Increasing transparency around speech regulation online — simply explaining publicly and on an on-going basis how and why platforms decide which content users see — would help respond to the concerns from both the left and the right over how platforms are moderating their sites. Where platforms can and will engage in sophisticated speech regulation to serve a variety of goals,[20] there is legitimate concern that platforms may intentionally or inadvertently harm one political group. Moreover, a clearer understanding of how platforms do regulate speech would make it easier for interest groups to collaborate to address concerns over the disparate impacts of harmful speech. As the pandemic forces platforms to expand the mechanisms they employ to regulate speech online, transparency becomes even more important, as their actions — broad, sweeping changes made over the course of a few months — are those that are most likely to have disproportionate effects on certain communities that the platforms themselves may not initially understand.

More transparency will not solve the underlying tension between the right and left over how to regulate speech online, but it would help sustain a consensus that platforms, not the government, are better suited to regulate their own spaces. Were that consensus to splinter and result in changes to CDA 230, there is little assurance that any subsequent policy framework would result in a better balance of free speech and regulation than platforms are capable of producing on their own.