

The Dark Psychology of Social Networks

Why it feels like everything is going haywire



Story by Jonathan Haidt and Tobias Rose-Stockwell

DECEMBER 2019 ISSUE | TECHNOLOGY

SUPPOSE THAT THE biblical story of Creation were true: God created the universe in six days, including all the laws of physics and all the physical constants that apply throughout the universe. Now imagine that one day, in the early 21st century, God became bored and, just for fun, doubled the gravitational constant. What would it be like to live through such a change? We'd all be pulled toward the floor; many buildings would collapse; birds would fall from the sky; the Earth would move closer to the sun, reestablishing orbit in a far hotter zone.

 TheAtlantic - Why It Feels Like Everything Is Going Haywire - The Atlantic - ...

To hear more feature stories, see our full list or get the Audm iPhone app.

Let's rerun this thought experiment in the social and political world, rather than the physical one. The U.S. Constitution was an exercise in intelligent design. The Founding Fathers knew that most previous democracies had been unstable and short-lived. But they were excellent psychologists, and they strove to create institutions and procedures that would work with human nature to resist the forces that had torn apart so many other attempts at self-governance.

For example, in "Federalist No. 10," James Madison wrote about his fear of the power of "faction," by which he meant strong partisanship or group interest that "inflamed [men] with mutual animosity" and made them forget about the common good. He thought that the vastness of the United States might offer some protection from the ravages of factionalism, because it would be hard for anyone to spread outrage over such a large distance. Madison presumed that factious or divisive leaders "may kindle a flame within their particular States, but will be unable to spread a general conflagration through the other States." The Constitution included mechanisms to slow things down, let passions cool, and encourage reflection and deliberation.

Madison's design has proved durable. But what would happen to American democracy if, one day in the early 21st century, a technology appeared that—over the course of a decade—changed several fundamental parameters of social and political life? What if this technology greatly increased the amount of "mutual animosity" and the speed at which outrage spread? Might we witness the political equivalent of buildings collapsing, birds falling from the sky, and the Earth moving closer to the sun?

America may be going through such a time right now.

What Social Media Changed

Facebook's early mission was "to make the world more open and connected"—and in the first days of social media, many people assumed that a huge global increase in connectivity would be good for democracy. As social media has aged,

however, optimism has faded and the list of known or suspected harms has grown: Online political discussions (often among anonymous strangers) are experienced as angrier and less civil than those in real life; networks of partisans co-create worldviews that can become more and more extreme; disinformation campaigns flourish; violent ideologies lure recruits.

The problem may not be connectivity itself but rather the way social media turns so much communication into a public performance. We often think of communication as a two-way street. Intimacy builds as partners take turns, laugh at each other's jokes, and make reciprocal disclosures. What happens, though, when grandstands are erected along both sides of that street and then filled with friends, acquaintances, rivals, and strangers, all passing judgment and offering commentary?

The social psychologist Mark Leary coined the term *sociometer* to describe the inner mental gauge that tells us, moment by moment, how we're doing in the eyes of others. We don't really need self-esteem, Leary argued; rather, the evolutionary imperative is to get *others* to see us as desirable partners for various kinds of relationships. Social media, with its displays of likes, friends, followers, and retweets, has pulled our sociometers out of our private thoughts and posted them for all to see.

Human beings evolved to gossip, preen, manipulate, and ostracize. We are easily lured into this new gladiatorial circus.

If you constantly express anger in your private conversations, your friends will likely find you tiresome, but when there's an audience, the payoffs are different—outrage can boost your status. A 2017 study by William J. Brady and other researchers at NYU measured the reach of half a million tweets and found that each moral or emotional word used in a tweet increased its virality by 20 percent, on average. Another 2017 study, by the Pew Research Center, showed that posts exhibiting "indignant disagreement" received nearly twice as much engagement—including likes and shares—as other types of content on Facebook.

The philosophers Justin Tosi and Brandon Warmke have proposed the useful phrase *moral grandstanding* to describe what happens when people use moral talk to enhance their prestige in a public forum. Like a succession of orators speaking to a skeptical audience, each person strives to outdo previous speakers, leading to some common patterns. Grandstanders tend to “trump up moral charges, pile on in cases of public shaming, announce that anyone who disagrees with them is *obviously* wrong, or exaggerate emotional displays.” Nuance and truth are casualties in this competition to gain the approval of the audience.

Grandstanders scrutinize every word spoken by their opponents—and sometimes even their friends—for the potential to evoke public outrage. Context collapses. The speaker’s intent is ignored.

Human beings evolved to gossip, preen, manipulate, and ostracize. We are easily lured into this new gladiatorial circus, even when we know that it can make us cruel and shallow. As the Yale psychologist Molly Crockett has argued, the normal forces that might stop us from joining an outrage mob—such as time to reflect and cool off, or feelings of empathy for a person being humiliated—are attenuated when we can’t see the person’s face, and when we are asked, many times a day, to take a side by publicly “liking” the condemnation.

[From October 2018: America is living James Madison’s nightmare]

In other words, social media turns many of our most politically engaged citizens into Madison’s nightmare: arsonists who compete to create the most inflammatory posts and images, which they can distribute across the country in an instant while their public sociometer displays how far their creations have traveled.

Upgrading the Outrage Machine

At its inception, social media felt very different than it does today. Friendster, Myspace, and Facebook all appeared between 2002 and 2004, offering tools that helped users connect with friends. The sites encouraged people to post highly curated versions of their lives, but they offered no way to spark contagious outrage. This changed with a series of small steps, designed to improve user experience, that collectively altered the way news and anger spread through

American society. In order to fix social media—and reduce its harm to democracy—we must try to understand this evolution.

When Twitter arrived in 2006, its primary innovation was the timeline: a constant stream of 140-character updates that users could view on their phone. The timeline was a new way of consuming information—an unending stream of content that, to many, felt like drinking from a fire hose.

Later that year, Facebook launched its own version, called the News Feed. In 2009, it added the “Like” button, for the first time creating a public metric for the popularity of content. Then it added another transformative innovation: an algorithm that determined which posts a user would see, based on predicted “engagement”—the likelihood of an individual interacting with a given post, figuring in the user’s previous likes. This innovation tamed the fire hose, turning it into a curated stream.

Mark Pernice

The News Feed’s algorithmic ordering of content flattened the hierarchy of credibility. Any post by any producer could stick to the top of our feeds as long as it generated engagement. “Fake news” would later flourish in this environment, as a personal blog post was given the same look and feel as a story from *The New York Times*. ^[P]_[SEP]

Twitter also made a key change in 2009, adding the “Retweet” button. Until then, users had to copy and paste older tweets into their status updates, a small obstacle that required a few seconds of thought and attention. The Retweet

button essentially enabled the frictionless spread of content. A single click could pass someone else's tweet on to all of your followers—and let you share in the credit for contagious content. In 2012, Facebook offered its own version of the retweet, the “Share” button, to its fastest-growing audience: smartphone users.

Chris Wetherell was one of the engineers who created the Retweet button for Twitter. He admitted to *BuzzFeed* earlier this year that he now regrets it. As Wetherell watched the first Twitter mobs use his new tool, he thought to himself: “We might have just handed a 4-year-old a loaded weapon.”

The coup de grâce came in 2012 and 2013, when Upworthy and other sites began to capitalize on this new feature set, pioneering the art of testing headlines across dozens of variations to find the version that generated the highest click-through rate. This was the beginning of “You won't believe ...” articles and their ilk, paired with images tested and selected to make us click impulsively. These articles were not usually intended to cause outrage (the founders of Upworthy were more interested in uplift). But the strategy's success ensured the spread of headline testing, and with it emotional story-packaging, through new and old media alike; outrageous, morally freighted headlines proliferated in the following years. In *Esquire*, Luke O'Neil reflected on the changes wrought on mainstream media and declared 2013 to be “The Year We Broke the Internet.” The next year, Russia's Internet Research Agency began mobilizing its network of fake accounts, across every major social-media platform—exploiting the new outrage machine in order to inflame partisan divisions and advance Russian goals.

The internet, of course, does not bear sole responsibility for the pitch of political anger today. The media have been fomenting division since Madison's time, and political scientists have traced a portion of today's outrage culture to the rise of cable television and talk radio in the 1980s and '90s. A multiplicity of forces are pushing America toward greater polarization. But social media in the years since 2013 has become a powerful accelerant for anyone who wants to start a fire.

The Decline of Wisdom

Even if social media could be cured of its outrage-enhancing effects, it would still raise problems for the stability of democracy. One such problem is the degree to which the ideas and conflicts of the present moment dominate and displace older

ideas and the lessons of the past. As children grow up in America, rivers of information flow continually into their eyes and ears—a mix of ideas, narratives, songs, images, and more. Suppose we could capture and quantify three streams in particular: information that is new (created within the past month), middle-aged (created 10 to 50 years ago, by the generations that include the child's parents and grandparents), and old (created more than 100 years ago).

Citizens are now more connected to one another, on platforms that have been designed to make outrage contagious.

Whatever the balance of these categories was in the 18th century, the balance in the 20th century surely shifted toward the new as radios and television sets became common in American homes. And that shift almost certainly became still more pronounced, and quickly so, in the 21st century. When the majority of Americans began using social media regularly, around 2012, they hyper-connected themselves to one another in a way that massively increased their consumption of new information—entertainment such as cat videos and celebrity gossip, yes, but also daily or hourly political outrages and hot takes on current events—while reducing the share of older information. What might the effect of that shift be?

In 1790, the Anglo-Irish philosopher and statesman Edmund Burke wrote, “We are afraid to put men to live and trade each on his own private stock of reason; because we suspect that this stock in each man is small, and that the individuals would do better to avail themselves of the general bank and capital of nations and of ages.” Thanks to social media, we are embarking on a global experiment that will test whether Burke's fear is valid. Social media pushes people of all ages toward a focus on the scandal, joke, or conflict of the day, but the effect may be particularly profound for younger generations, who have had less opportunity to acquire older ideas and information before plugging themselves into the social-media stream.

Our cultural ancestors were probably no wiser than us, on average, but the ideas we inherit from them have undergone a filtration process. We mostly learn of

ideas that a succession of generations thought were worth passing on. That doesn't mean these ideas are always right, but it does mean that they are more likely to be valuable, in the long run, than most content generated within the past month. Even though they have unprecedented access to all that has ever been written and digitized, members of Gen Z (those born after 1995 or so) may find themselves less familiar with the accumulated wisdom of humanity than any recent generation, and therefore more prone to embrace ideas that bring social prestige within their immediate network yet are ultimately misguided.

For example, a few right-wing social-media platforms have enabled the most reviled ideology of the 20th century to draw in young men hungry for a sense of meaning and belonging and willing to give Nazism a second chance. Left-leaning young adults, in contrast, seem to be embracing socialism and even, in some cases, communism with an enthusiasm that at times seems detached from the history of the 20th century. And polling suggests that young people across the political spectrum are losing faith in democracy.

Is There Any Way Back?

Social media has changed the lives of millions of Americans with a suddenness and force that few expected. The question is whether those changes might invalidate assumptions made by Madison and the other Founders as they designed a system of self-governance. Compared with Americans in the 18th century—and even the late 20th century—citizens are now more connected to one another, in ways that increase public performance and foster moral grandstanding, on platforms that have been designed to make outrage contagious, all while focusing people's minds on immediate conflicts and untested ideas, untethered from traditions, knowledge, and values that previously exerted a stabilizing effect. This, we believe, is why many Americans—and citizens of many other countries, too—experience democracy as a place where everything is going haywire.

[There was a time in American public life when atonement was seen as a form of strength—a way not only to own up to one's missteps but also to control the narrative. That time is over, Megan Garber writes.]

It doesn't have to be this way. Social media is not intrinsically bad, and has the power to do good—as when it brings to light previously hidden harms and gives voice to previously powerless communities. Every new communication technology brings a range of constructive and destructive effects, and over time, ways are found to improve the balance. Many researchers, legislators, charitable foundations, and tech-industry insiders are now working together in search of such improvements. We suggest three types of reform that might help:

(1) Reduce the frequency and intensity of public performance. If social media creates incentives for moral grandstanding rather than authentic communication, then we should look for ways to reduce those incentives. One such approach already being evaluated by some platforms is “demetrication,” the process of obscuring like and share counts so that individual pieces of content can be evaluated on their own merit, and so that social-media users are not subject to continual, public popularity contests.

(2) Reduce the reach of unverified accounts. Bad actors—trolls, foreign agents, and domestic provocateurs—benefit the most from the current system, where anyone can create hundreds of fake accounts and use them to manipulate millions of people. Social media would immediately become far less toxic, and democracies less hackable, if the major platforms required basic identity verification before anyone could open an account—or at least an account type that allowed the owner to reach large audiences. (Posting itself could remain anonymous, and registration would need to be done in a way that protected the information of users who live in countries where the government might punish dissent. For example, verification could be done in collaboration with an independent nonprofit organization.)

(3) Reduce the contagiousness of low-quality information. Social media has become more toxic as friction has been removed. Adding some friction back in has been shown to improve the quality of content. For example, just after a user submits a comment, AI can identify text that's similar to comments previously flagged as toxic and ask, “Are you sure you want to post this?” This extra step has been shown to help Instagram users rethink hurtful messages. The quality of information that is spread by recommendation algorithms could likewise be

improved by giving groups of experts the ability to audit the algorithms for harms and biases.

MANY AMERICANS MAY THINK that the chaos of our time has been caused by the current occupant of the White House, and that things will return to normal whenever he leaves. But if our analysis is correct, this will not happen. Too many fundamental parameters of social life have changed. The effects of these changes were apparent by 2014, and these changes themselves facilitated the election of Donald Trump.

If we want our democracy to succeed—indeed, if we want the *idea* of democracy to regain respect in an age when dissatisfaction with democracies is rising—we’ll need to understand the many ways in which today’s social-media platforms create conditions that may be hostile to democracy’s success. And then we’ll have to take decisive action to improve social media.

This article appears in the December 2019 print edition with the headline “Why It Feels Like Everything Is Going Haywire.”