

Medium, Message, and Misinformation

IF IN 1990 YOU HAD TOLD US THAT BY 2020 NEARLY HALF OF THE people on the planet would carry a wallet-size machine that could instantly look up any fact in the world—a “smartphone”—we would have predicted an end to bullshit. How could you bullshit someone who could check your claims easily, immediately, and costlessly?

Apparently people have neither the time nor the inclination to use smartphones this way. Instead, smartphones have become just one more vehicle for spreading bullshit. On the positive side, you can have a decent dinner conversation without being fact-checked thirty times. On the negative side, bullshit goes largely unchallenged.

Technology didn’t eliminate our bullshit problem, it made the problem worse. In this chapter, we will explore how that happened. In short, the rise of the Internet changed what kinds of information get produced, how information is shared, and the ways in which we find the information that we want. While much good has come of the Internet revolution, there have been major drawbacks as well. Fluff and glitter have overtaken serious, in-depth, thoughtful content. News coverage has become increasingly partisan. Misinformation, disinformation, and fake news abound. We will consider these issues in turn.

THE BROTHEL OF THE PRINTING PRESS

Pity the soul who hopes to hold back a revolution in information technology. Priest and scribe Filippo de Strata lived through one

such revolution. In 1474, he railed against the damage wrought by the invention of the printing press. Printers, de Strata argued, “shamelessly print, at a negligible price, material which may, alas, inflame impressionable youths, while a true writer dies of hunger. . . .” By massively lowering the cost of manufacturing books, the printing press was bound to reduce the value and authority of text. When every book had to be written by hand, only royalty and the clergy could commission a well-trained scribe like de Strata to produce a copy of a book. The great expense of hiring scribes served as a filter on the kind of information committed to paper. There was little demand for books that served only as trivial entertainment; most new books were copies of the Bible and other documents of huge importance. But the advent of the printing press opened a spillway through which less serious content could flood the market. Publicly, Filippo de Strata fretted that that the “brothel of the printing press” would lead readers to cheap, salacious entertainment—the works of Ovid, even. Privately, he may have been more concerned about his own job security.

Others worried about a proliferation of fluff that would obscure important information. Pioneers in cataloging human knowledge, such as Conrad Gessner in the sixteenth century and Adrien Baillet in the seventeenth century, cautioned that the printing press would bring scholarship to a halt as readers became overwhelmed by the range of options for study. They were wrong. With the advantage of a few centuries’ hindsight, we can see that Gutenberg’s revolution brought vastly more good than harm. The printing press—later coupled with public libraries—democratized the written word. In the year 1500, German writer Sebastian Brant described this change:

In our time . . . books have emerged in lavish numbers. A book that once would’ve belonged only to the rich—nay, to a king—can now be seen under a modest roof. . . . There is nothing nowadays that our children . . . fail to know.

Still, Filippo de Strata was right that when the cost of sharing information drops dramatically, we see changes in both the nature of

the information available and the ways that people interact with that information.*

Roughly five hundred years after Filippo de Strata sounded the alarm about the printing press, sociologist Neil Postman echoed his sentiments:

The invention of new and various kinds of communication has given a voice and an audience to many people whose opinions would otherwise not be solicited, and who, in fact, have little else but verbal excrement to contribute to public issues.

If we wanted to condemn blogs, Internet forums, and social media platforms, we could scarcely say it better. But Postman was not referring to social media or even the Internet. He delivered this line a half century ago. In a 1969 lecture, he lamented the lowbrow programming on television, the vacuous articles in newspapers and magazines, and the general inanity of mass media. This kind of infotainment, he maintained, distracts consumers from the information that does matter—and distraction can itself be a form of disinformation. If religion is the opiate of the masses, *Jersey Shore* and *Temptation Island* are the spray canisters from which the masses huff metallic paint fumes.

Since Postman's lecture, we've undergone another revolution. The Internet has changed the way we produce, share, and consume information. It has altered the way we do research, learn about current events, interact with our peers, entertain ourselves, and even think. But why has the Internet also triggered a bullshit pandemic of unprecedented proportions?

Let's begin by looking at what gets published. Through the 1980s, publishing required money—a lot of it. Typesetting was expensive, printing required substantial overhead, and distribution involved getting physical paper into the hands of readers. Today, anyone with a personal computer and an Internet connection can produce professional-looking documents and distribute them around the world without cost. And they can do so in their pajamas.

This is the democratizing promise of the Internet: endless new

* Eörs Szathmáry and John Maynard Smith make a similar point, applied to living systems, in their 1995 book *The Major Transitions in Evolution*. Their ideas about biological transitions in information use have been instrumental to us in developing this section.

voices brought into a worldwide conversation. Members of marginalized groups, who previously might have lacked the financial and social capital to publish and publicize their work, can now make their stories heard. At the same time, the new technology captures the long tail of interests and creates communities around even the rarest of obsessions. Want to build your own air calliope? Explore the Scooby-Doo cartoons from a critical theory perspective? Learn the complex dice games played by the protagonists of *The Canterbury Tales*? The Internet has you covered.

This democratization has a dark side as well. Aided by viral spread across social media, amateur writers can reach audiences as large as those of professional journalists. But the difference in the reporting quality can be immense. A typical Internet user lacks the journalistic training, let alone the incentives to report accurately. We can access more information than ever, but that information is less reliable.

Prior to the Internet, mass media filled our living rooms with voices from afar—but these voices were familiar to us. We listened to Ed Murrow; we read the words of familiar newspaper columnists; we watched Walter Cronkite, “the most trusted man in America”; and we dove into the fictional worlds created by famous authors. In today’s social media world, our friends treat us to saccharine drivel about their latest soul mates, square-framed snapshots of their locally sourced organic brunches, and tiresome boasts about their kids’ athletic, artistic, or academic accomplishments. But our homes are also filled with the voices of strangers—often anonymous strangers—that our friends have seen fit to share. We don’t know these people. What they write is seldom written with the attention to accuracy we would expect from a commercial media outlet. And some of the “authors” are paid human agents or computer programs spreading disinformation on behalf of corporate interests or foreign powers.

Back when news arrived at a trickle, we might have been able to triage this information effectively. But today we are confronted with a deluge. As we are writing this chapter, we both have multiple browser windows open. Each window has roughly ten open tabs, and each of those tabs contains a news story, journal article, blog post, or other information source that we intend to revisit but never will. Additional stories and tidbits are scrolling across our social media feeds faster than we could track, even if we did nothing else. Because there

is so much more volume and so much less filtering, we find ourselves like the Sorcerer's Apprentice: overwhelmed, exhausted, and losing the will to fight a torrent that only flows faster with every passing hour.

THE INADEQUACY OF THE UNVARNISHED TRUTH

Filippo de Strata feared that the works of Ovid might crowd out the Bible. We fear that the mindless lists, quizzes, memes, and celebrity gossip that proliferate on social media might crowd out thoughtful analyses of the sort you see in *The New York Times* or *The Wall Street Journal*. Every generation thinks that its successor's lazy habits of mind will bring on a cultural and intellectual decline. It may be a stodgy lament that has been repeated for thousands of years, but it's our turn now, and we're not going to miss the opportunity to grumble.

Prior to the Internet, newspapers and magazines made money by selling subscriptions.* Subscribing to a periodical, you embarked on a long-term relationship. You cared about the quality of information a source provided, its accuracy, and its relevance to your daily life. To attract subscribers and keep them, publishers provided novel and well-vetted information.

The Internet news economy is driven by clicks. When you click on a link and view a website, your click generates advertising revenue for the site's owner. The Internet site is not necessarily designed to perpetuate a long-term relationship; it is designed to make you click, now. Quality of information and accuracy are no longer as important as *sparkle*. A link needs to catch your eye and pull you in. Internet publishers are not looking for Woodward and Bernstein. Instead, they want "Seven Cats That Look Like Disney Princesses," "Eight Amazing Nutrition Secrets Your Personal Trainer Doesn't Want You to Know," "Nine Never-Before-Published Photos of Elvis Found in Retiree's Attic," and "Ten Ways That Experts Spot Quantitative Bullshit."

Publishers produce this fluff because we click on it. We might aspire to patronize quality news sources that provide nuanced analysis.

* A considerable fraction of the revenue to most newspaper and magazine publishers has always come from advertising, but advertising revenue scaled with the subscription base, so again publishers needed to maximize subscriptions.

But faced with the temptation to click on the informational equivalent of empty calories, the mental junk food usually wins.

You can see this trend toward fluff in headlines. Headlines draw our attention—and in the social media environment, where many people never read any further, they are a significant source of information. A satirical website published a headline proclaiming that “70% of Facebook Users Only Read the Headline of Science Stories before Commenting.” The story began by noting that most people don’t read stories before sharing them on social media either. After a couple of sentences, the text gave way to paragraph after paragraph of the standard “*lorem ipsum dolor . . .*”—random text used as filler for webpage layouts. The post was shared tens of thousands of times on social media, and we don’t know how many of those who did so were in on the joke.

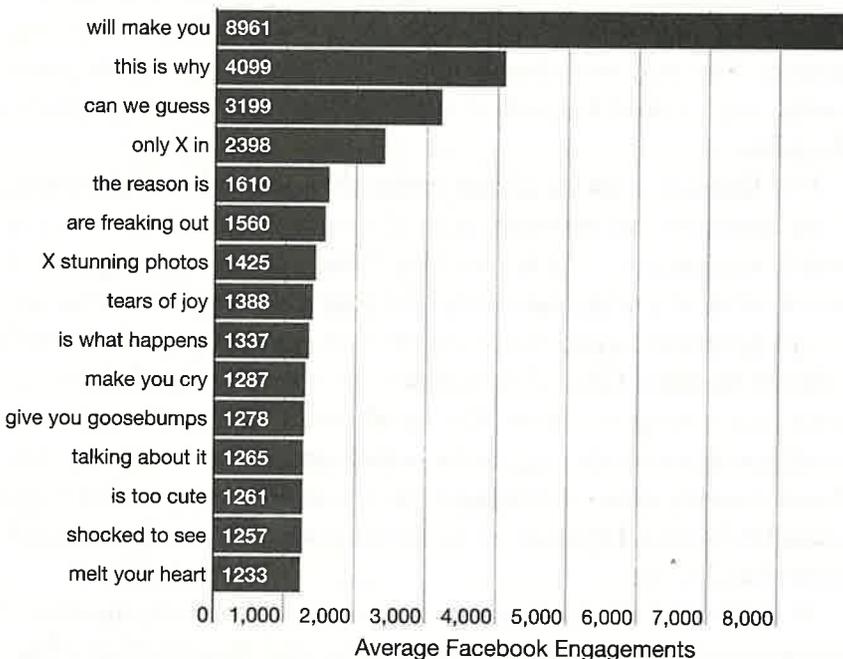
Headlines, written by editors rather than journalists, have always been somewhat inconsistent with the stories they announce. But within a single issue of *The New York Times*, for example, the articles aren’t competing with one another for your attention. The newspaper is trying to create a package of stories that together provide as much value as possible. Click-driven media, on the other hand, drives an arms race among headlines. On social media sites and news feeds, headlines from competing media outlets are presented side by side. Readers rarely read everything—there is simply too much content available. Instead they click on the most tantalizing or titillating headlines that they see.

How do you win an arms race to come up with catchy headlines? Sensationalism works. Tabloids have long used sensational headlines to draw attention at the newsstand, but major subscription papers largely eschewed this practice. But sensationalism is not the only way. Entrepreneur Steve Rayson looked at 100 million articles published in 2017 to determine what phrases were common in the headlines of articles that were widely shared. Their results will make you gasp in surprise—unless you’ve spent a few minutes on the Internet at some point in the past few years.

The study found that the most successful headlines don’t convey facts, *they promise you an emotional experience*. The most common phrase among successful Facebook headlines, by nearly twofold, is “will make you,” as in “will break your heart,” “will make you fall in love,”

“will make you look twice,” or “will make you gasp in surprise” as above. This phrase is also highly successful on Twitter. Other top phrases include “make you cry,” “give you goosebumps,” and “melt your heart.” Intellectual experiences cannot compete. Pause for a moment and think about what a huge shift this represents. Can you imagine *The New York Times* or your local newspaper with headlines that told you how you’d feel about each story, but not what the story actually entailed?

Top Headline Phrases



Headlines once aimed to concisely convey the essence of a story: “Kennedy Is Killed by Sniper as He Rides in Car in Dallas; Johnson Sworn In on Plane.” “Men Walk on Moon. Astronauts Land on Plain; Collect Rocks, Plant Flag.” “East Germany Opens Wall and Borders, Allowing Citizens to Travel Freely to the West.”

With click-driven advertising, if a headline tells too much, there is little incentive to click on the story. Headlines now go through con-tortions *not* to tell you what the story says. While these so-called forward reference headlines are most commonly used by Internet media

companies, traditional media are getting into the game as well. “One-Fifth of This Occupation Has a Serious Drinking Problem,” announces *The Washington Post*. “How to Evade the Leading Cause of Death in the United States,” CNN promises to inform you. “Iceland Used to Be the Hottest Tourism Destination. What Happened?” asks *USA Today*. (So as not to leave you in suspense: lawyers; don’t get in a car accident; and nobody knows.)

Headlines also lure us in by making the story about *us*. In the world of social media, news is a two-way street in which everyone is both a consumer and a producer. As we were writing this section, the following headlines were coming across our social media feeds.

- “People Are Freaking Out About This Photo That Might Show That Amelia Earhart Survived Her Crash” (*BuzzFeed*)
- “‘This is a huge bombshell.’ Twitter reacts to NCAA arrests by FBI” (*IndyStar*)
- “McDonald’s Invented ‘Forks’ and the Internet Can’t Stop Talking About Them” (*HuffPost*)

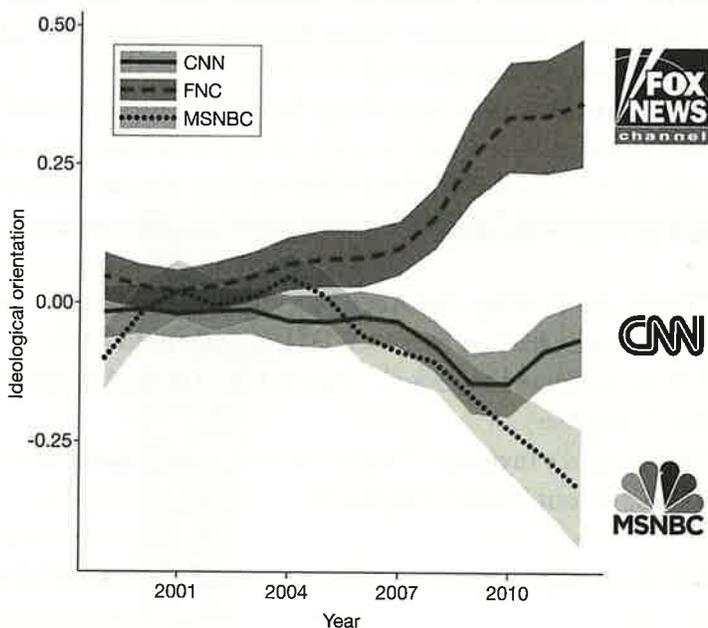
What we are saying becomes more interesting than what is happening.

All of this fluff and glitter does more than just dumb down the national conversation: It opens the door for bullshit. *The unvarnished truth is no longer good enough*. Straight-up information cannot compete in this new marketplace.

PARTISANSHIP, PERSONALIZATION, AND POLARIZATION

Much as the invention of the printing press allowed for a more diverse array of books, the advent of cable television allowed people to select specialized media outlets that closely reflected their views. Prior to 1987, the Fairness Doctrine of the US Federal Communications Commission (FCC) strived to ensure balanced coverage of controversial issues in news programming. But it was repealed under President Ronald Reagan. Hastened by the advent of the 24-hour news cycle, cable news channels proliferated and specialized in delivering specific

political perspectives. In the United States, mainstream news has become increasingly partisan over the past twenty years. The figure below illustrates the diverging ideological positions of three prominent cable news channels, as estimated from broadcast transcripts.



Online, it's the same story, only more so. Even mainstream outlets deliver news with a partisan slant. We find ourselves isolated in separate echo chambers. Publishers such as Breitbart News Network and The Other 98% go one step further, pushing what is known as hyperpartisan news. Their stories may be based in fact, but they are so strongly filtered through an ideological lens that they often include significant elements of untruth.

Publishers churn out partisan and hyperpartisan content because it pays to do so. Social media favors highly partisan content. It is shared more than mainstream news, and once shared, it is more likely to be clicked on. Deepening the ideological divide has become a lucrative business.

MIT professor Judith Donath has observed that even when people appear to be talking about other things, they're often talking about themselves. Suppose I log on to Facebook and share a false—even absurd—story about how airplane contrails are endocrine-disrupting

chemicals being sprayed as part of a liberal plot to lower the testosterone levels of America's youth. I may not be as interested in having you believe my claims about contrails as I am in signaling my own political affiliations. Sharing an article like this signals that I belong to a group of people who believe in conspiracy theories and distrust the "liberal agenda" in America. And if that's my aim, it doesn't matter to me whether the story is true or false. I may not have read it, I may not care if you read it, but I want you to know that I am a fellow tinfoil hatter.

The signal itself becomes the point. If I share a story about how the IRS is investigating Donald Trump's business dealings prior to the 2016 election, my political affiliation is unclear. But if I share a story that says Donald Trump has sold the Washington Monument to a Russian oligarch, it's clear that I hate Trump. And I'm demonstrating a political allegiance so strong that I can suspend disbelief when it comes to stories of Trump's treachery.

Professor Donath's insight springs from a broader tradition in the field known as *communication theory*. We often think of communication solely as the transmission of information from sender to receiver. But this ignores a second, broader social aspect of communication, one that is revealed by its origins in the Latin verb *communicare*, "to make shared or common."

Communication is how we establish, reinforce, and celebrate a shared framework for thinking about the world. Think about a religious mass, or even the scripted, ordered regularity of the nightly news. Communication over social media does the same thing: It creates and structures social communities. When we send out a tweet or Facebook post or Instagram image, we are affirming our commitment to the values and beliefs of our particular online community. As the community responds, these common values are reaffirmed through likes, shares, comments, or retweets.

Blindfolded and submerged in a pool, I shout "Marco!" If I do so correctly, my network of acquaintances sends back an encouraging chorus. "Polo! Polo! Polo!" Participating on social media is only secondarily about sharing new information; it is primarily about maintaining and reinforcing common bonds. The danger is that, in the process, what was once a nationwide conversation fragments beyond repair. People begin to embrace *tribal epistemologies* in which the truth

itself has less to do with facts and empirical observation than with who is speaking and the degree to which their message aligns with their community's worldview.

Algorithms only make matters worse. Facebook, Twitter, and other social media platforms use algorithms to find "relevant" posts and stories for you, personalizing your feed. These algorithms are not designed to keep you informed; they are designed to keep you active on the platform. The aim is to feed you content that is sufficiently compelling to prevent you from wandering elsewhere on the Web or, Lord forbid, going to bed at a decent hour. The problem is that the algorithms create a vicious cycle, giving you more of what they think you want to hear, and fewer opportunities to read divergent points of view. The actual details of these algorithms are hidden from view; but what you like and what you read, whom your friends are, your geolocation, and your political affiliations all influence what you see next. The algorithms amplify content that aligns with their guesses about your sociopolitical orientation, and they suppress alternative viewpoints.

On the Web, we are all test subjects. Commercial websites are continually running large-scale experiments to see what keeps us online and engaged. Online media companies experiment with different variations on a headline, different accompanying images, even different fonts or "click to continue" buttons. At the same time, Facebook and other platforms offer advertisers—including political advertisers—the ability to target specific consumers with messages designed to cater to their interests. These messages may not even be clearly identified as advertisements.

Think about what YouTube can learn as they experiment by recommending different videos and observing what users select to watch. With billions of videos viewed every day and vast computational resources, they can learn more about human psychology in a day than an academic researcher could learn in a lifetime. The problem is, their computer algorithms have learned that one way to retain viewers is to recommend increasingly extreme content over time. Users who watch left-leaning videos are quickly directed to extreme-left conspiracy theories; users who enjoy right-leaning material soon get recommendations for videos from white supremacists or Holocaust deniers. We've seen this ourselves. As Jevin and his six-year-old son watched

real-time video from the International Space Station 254 miles above spherical Earth, YouTube filled the sidebar of their screen with videos claiming that the earth is actually flat.

Riffing on Allen Ginsberg, tech entrepreneur Jeff Hammerbacher complained in 2011 that “the best minds of my generation are thinking about how to make people click ads. That sucks.” The problem is not merely that these “best minds” could have been devoted to the artistic and scientific progress of humankind. The problem is that all of this intellectual firepower is devoted to hijacking our precious attention and wasting *our* minds as well. The Internet, social media, smartphones—we are exposed to increasingly sophisticated ways of diverting our attention. We become addicted to connectivity, to meaningless checking, to a life of fragmented attention across myriad streams of digital information. In short, the algorithms driving social media content are bullshitters. They don’t care about the messages they carry. They just want our attention and will tell us whatever works to capture it.

MISINFORMATION AND DISINFORMATION

Social media facilitates the spread of *misinformation*—claims that are false but not deliberately designed to deceive. On social media platforms, the first outlet to break a story receives the bulk of the traffic. In the race to be first, publishers often cut any fact-checking out of the publication process. You can’t beat your competitors to press if you pause to rigorously fact-check a story. Being careful is admirable, but it doesn’t sell ads.

Social media is also fertile ground for *disinformation*, falsehoods that are spread deliberately.

A study in 2018 found that about 2.6 percent of US news articles were false. This might not seem like a big percentage, but if every American read one article per day, it would mean that nearly eight million people a day were reading a false story.

Sometimes false information is just a nuisance. One satirical news site claimed that musician Taylor Swift was dating notorious anti-communist senator Joseph McCarthy—who died forty-two years before she was born. Predictably, some fans were unable to realize the absurdity of the story and reacted with disgust. But no businesses

were undermined, no lives were at stake, and even Ms. Swift is unlikely to suffer any serious harm to her reputation.

But misinformation and disinformation can be far more serious. As more of the world comes online, the problem multiplies. For example, nearly half a billion Indian citizens received access to the Internet for the first time between 2010 and 2020. Overall, the rapid expansion of connectivity benefits both those coming online and those already there. Unfortunately, new Internet users tend to be more susceptible.

In addition to basic social messaging functions, WhatsApp serves as a news source to its 1.5 billion users worldwide. It is also a potent vector for spreading misinformation. In early 2018, Indian users widely shared a series of fake videos that purported to illustrate children being kidnapped by organized gangs. A fear of strangers spread, with disastrous consequences. Visiting a temple in Tamil Nadu, one family stopped to ask for directions. Locals became suspicious that they might be the type of kidnappers seen in the WhatsApp videos. A crowd gathered; the family was pulled from their car; the mob stripped them naked and beat them brutally with metal bars and sticks. One was killed and the others permanently disfigured. Motivated by the same false story, mobs attacked dozens of other innocent people, beating and often killing them.

Police tried to counter the online misinformation and put a stop to the killings. But the rumors traveled too fast. In some regions, authorities had to shut down the Internet completely to slow them down. WhatsApp tried its own interventions, changing how many times a message could be shared. Previously a message could be forwarded 250 times. They dropped this to five. Still the mob attacks continued.

It is not only new Internet users who are fooled. In December 2016, a website called AWD News published a frightening headline: "Israeli Defense Minister: If Pakistan Send Ground Troops to Syria on Any Pretext, We Will Destroy This Country with a Nuclear Attack."

The story contained several cues that should have tipped off a careful reader. The headline contained grammatical errors ("send" instead of "sends"). The story named the wrong person as the Israeli defense

minister.* The article sat next to other implausible headlines such as “Clinton Staging Military Coup against Trump.” But it still fooled the person you would least want it to: Khawaja Muhammad Asif, defense minister of Pakistan. Asif responded with a threat of his own via Twitter: “Israeli def min threatens nuclear retaliation presuming pak role in Syria against Daesh . . . Israel forgets Pakistan is a Nuclear state too.”

A single fake news piece led one major power to threaten another with a nuclear attack. It is one thing to mislead the infinitely gullible about Taylor Swift’s latest romance. It is another to tip the globe toward nuclear war.

And then there is political propaganda. Social media is a more effective medium for spreading propaganda than leaflets dropped from airplanes or high-powered radio transmissions directed into enemy territory. Social media posts are unconstrained by most borders. And they are shared organically. When social media users share propaganda they have encountered, they are using their own social capital to back someone else’s disinformation. If I come across a political leaflet or poster on the street, I am immediately skeptical. If my dear uncle forwards me a story on Facebook that he “heard from a friend of a friend,” my guard drops. Disinformation flows through a network of trusted contacts instead of being injected from outside into a skeptical society.

In 2017, Facebook admitted that over the past two years, 126 million US users—half of the adult population and about three-quarters of its US user base—had been exposed to Russian propaganda on the site. More than one hundred thirty thousand messages from these accounts were designed to deepen preexisting ideological divides within the US, and to seed mistrust between neighbors. They focused on emotionally charged issues such as race relations, gun rights, border security, welfare, and abortion, and played both sides of each topic to reach the largest possible audience. The goal was to amplify the loudest and most extreme voices in each political camp, while drowning out the more reasonable and productive discussants. When the divide

* The defense minister at the time was Avigdor Lieberman, but the story attributed the quotation to Moshe Yaalon, who had previously served in that role.

between political factions grows deep enough, conversations between them stop. Our trust in people and institutions erodes. We lose faith in our ability to make collective decisions. Ultimately we start to question the democratic process itself.

While we think of propaganda as designed to convince people of specific untruths, much modern propaganda has a different aim. The “firehose strategy” is designed to leave the audience disoriented and despairing of ever being able to separate truth from falsehood. Social media makes it easy to broadcast large volumes of disinformation at high rates across multiple channels. This is part of the firehose strategy. The other part is to deliberately eschew consistency. Rather than being careful to convey only a single cohesive story, the aim is to confuse readers with a large number of mutually contradictory stories. In 2016, chess grand master Garry Kasparov summarized this approach in a post on Twitter: “The point of modern propaganda isn’t only to misinform or push an agenda. It is to exhaust your critical thinking, to annihilate truth.”

Meanwhile, authoritarian governments have embraced social media. They were originally fearful of this medium and prone to censor its use. But more recently, governments such as those of China, Iran, and Russia have discovered that social media offers an ideal platform on which to monitor public sentiment, track dissent, and surreptitiously manipulate popular opinion.

Still, fake news is not *primarily* a propaganda tool. Most fake and hyperpartisan news is created for a different reason: to generate advertising revenue. Anyone, anywhere can get in on the action. In the final days of the 2016 US election, Barack Obama talked extensively about the fake news factories in Macedonia. The people running these factories—often teenagers—created at least 140 popular fake news websites during the election. When a story went viral, it generated huge advertising revenues for the site owners. Some of them were making in excess of \$5,000 per month, compared with the average Macedonian monthly salary of \$371. The teens writing these stories didn’t care whether Trump or Clinton won; they cared only about clicks. The most shared fake news story in that whole election proclaimed that “Pope Francis Shocks the World, Endorses Donald Trump for President.” This story was created by a group of teenagers

in Macedonia, under the aegis of WT05 News, and received nearly a million engagements on Facebook. To put that in perspective, the top *New York Times* article during this same period received 370,000 engagements.

Pope Francis was not happy about this story. He issued the following statement about the salacious reporting in general and about fake news in particular. “I think the media have to be very clear, very transparent, and not fall into—no offence intended—the sickness of coprophilia, that is, always wanting to cover scandals, covering nasty things, even if they are true,” he said. “And since people have a tendency towards the sickness of coprophagia, a lot of damage can be done.”

Fake news purveyors, take note. When the pope himself says you eat shit, it is time to reevaluate your life choices.

THE NEW COUNTERFEITERS

Since the advent of money, governments have had to deal with counterfeiting. Precious-metal coins came into use in the Mediterranean world in the sixth century B.C.E.; soon after counterfeiters started to produce facsimiles with gold or silver plating over a cheaper base metal. They’ve been at it ever since. Conducted on a sufficiently large scale, counterfeiting can undermine public trust in a currency, devalue the currency, and drive runaway inflation. Counterfeiting has often been used this way during wartime—by the British in the American Revolutionary War, by the Union in the American Civil War, and by the Nazis in the Second World War, to list just a few examples.

In an Internet-connected world, governments have to worry about a new kind of counterfeiting—not of money, but of *people*. Researchers estimate that about half of the traffic on the Internet is due not to humans, but rather “bots,” automated computer programs designed to simulate humans. The scale of the problem is staggering. By 2018, Facebook had over two billion legitimate users—but in the same year deleted even more fake accounts: nearly three billion. Some bots act as information providers, pushing out their messages, usually for advertising purposes but occasionally in service of propaganda aims. Others

emulate information consumers. “Click farms” use massive banks of cell phones to generate webpage views or YouTube video views for a fee.

For a representative democracy to function, constituents must be able to share their views with their elected officials. The ability to counterfeit people threatens to swamp real voices with fake ones. In mid-2017, the FCC solicited public comments on a proposal to eliminate “net neutrality,” the requirement that Internet service providers treat all information that they transmit the same, regardless of source or content. The FCC received a staggering 21.7 million citizen comments in response—but a large fraction of these appeared to be fraudulent. Over half came from throwaway email addresses or from addresses used to send multiple comments. A strong signature of bot activity is the simultaneous transmission of huge numbers of messages. At 2:57:15 P.M. EDT on July 19, 2017, half a million similar comments were sent at the exact same second. One bot submitted over a million different anti-net neutrality comments following the same basic structure, using a MadLibs approach of plugging in synonyms for various words. Half a million comments on net neutrality came directly from Russian email addresses. The New York State attorney general estimated that nearly 10 million of the comments were sent using stolen identities of people who had no idea their names were being used. In the end, the vast majority of the twenty-one million comments the FCC received were in favor of abandoning net neutrality—even though there is reason to believe public opinion ran strongly in the opposite direction.

The most influential fake accounts are run not by bots but by real people pretending to be someone else. All-American girl Jenna Abrams was an Internet celebrity who spent a lot of time commenting on pop culture, but also broadcast provocative right-wing views to her seventy thousand Twitter followers. The problem is, Jenna Abrams wasn’t real. She was a creation of a Moscow propaganda outfit known as the Internet Research Agency. Still, she was very effective. Her tweets were retweeted extensively and covered by top news outlets including *The New York Times* and *The Washington Post*. She fooled the magazine *Variety* into posting a fake news story about CNN airing pornography. Once exposed, she turned around and mocked the media for covering her: “Don’t blame the media for including my

tweets in their articles. If you pull the right strings a puppet will dance any way you desire.”

Until recently, creating fake personalities on the Internet was tricky because in a social media world, we want to see the pictures of the people we follow. Without a picture, a user account on the Internet could be anyone. The computer security professional posting to a tech forum could be a kid in his mom’s basement. A fourteen-year-old girl in a chat room could be an undercover cop. The oil heiress in your inbox is undoubtedly a scam artist. But if we can see a picture, we tend to be less suspicious. Fake accounts sometimes use stock photos or images scraped from the Internet—but these were easily tracked down by savvy users using tools such as Google’s reverse image search.

No longer. A new class of algorithms, collectively known as adversarial machine learning, can fashion photorealistic faces of nonexistent people out of whole cloth. The fabricated images are stunningly good. This is a dangerous period for a technology: It is widely available but few people know it’s being used. To raise public awareness, we developed a website called WhichFaceIsReal.com. A real photograph of a real person is paired with a computer-generated image of someone who does not exist. Your aim is to guess which is which. More than a million people have played the game on our website and the results show just how good the fakes are. People do not do much better than chance when they start playing, and even with lots of practice people still are fooled one time in five.

Similar machine learning algorithms are able to “voiceshop,” generating fake audio and video that are nearly indistinguishable from the real thing. By synthesizing audio from previous recordings and grafting expressions and facial movements from a person acting as model onto the visage of a target, these so-called deepfake videos can make it look like anyone is doing or saying anything.

Director and comedian Jordan Peele created a public service announcement about fake news using this technology. Peele’s video depicts Barack Obama addressing the American people about fake news, misinformation, and the need for trusted news sources. Midway through the video, however, the face of Jordan Peele appears next to Obama, speaking the same words in perfect time, clearly the model from which Obama’s facial movements and expressions have been derived. Obama concludes, in his own voice but with Peele’s words:

“How we move forward in the age of information is going to be the difference between whether we survive or whether we become some kind of fucked-up dystopia.”

Confronted with all of these technologies for fabricating reality, one might lose hope of getting to the truth about anything. We are not so pessimistic. Our society adjusted to the anonymity afforded by an Internet on which “nobody knows you’re a dog.” And we adjusted to a Photoshop world in which pictures do lie. How? In a word, we triangulate. We no longer trust a single message, a single image, a single claim. We look for independent witnesses who can confirm testimony. We seek multiple images from multiple vantage points. Society will adjust similarly to a world of deepfakes and whatever reality-bending technologies follow.

There are three basic approaches for protecting ourselves against misinformation and disinformation online. The first is technology. Tech companies might be able to use machine learning to detect online misinformation and disinformation. While this is a hot area for research and development, we are not optimistic. Tech companies have been trying to do this for years, but the problem shows no signs of abating. Microsoft, Facebook, and others have recently started to release large data sets to academic researchers working on this problem; that suggests to us that the tech companies know that they need help. And economically, it’s not clear that Internet companies have sufficient incentives. After all, extreme content is highly effective at drawing an audience and keeping users on a platform. Technologically, the same artificial intelligence techniques used to detect fake news can be used to get around detectors, leading to an arms race of production and detection that the detectors are unlikely to win.

A second approach is governmental regulation. Some countries have already passed laws against creating or spreading fake news, but we worry about this approach for two reasons. First, it runs afoul of the First Amendment to the US Constitution, which guarantees freedom of speech. Second, who gets to determine what is fake news? If a leader doesn’t like a story, he or she could declare it fake news and pursue criminal charges against the perpetrators. That has happened already in some parts of the world. A lighter regulatory touch might help. We have long advocated a legislative ban on targeted political advertising online, and are heartened by Twitter’s self-imposed mora-

torium. We would like to see users control the information that comes across their social media feeds, rather than being forced to rely on a hopelessly opaque algorithm. The justification for the FCC's Fairness Doctrine was one of democratic necessity. For a democracy to function properly, a country needs an informed populace with access to reliable information. Similar arguments could justify governmental regulation of social media.

A third and most powerful approach is education. If we do a good job of educating people in media literacy and critical thinking, the problem of misinformation and disinformation can be solved from the bottom up. That is our focus in this book, and in much of our professional lives.

Every generation has looked back on the past with nostalgia for a simpler and more honest time. We may have greater than usual cause for nostalgia. The Internet has brought about a sea change in the way that information is created, sorted, discovered, spread, and consumed. While this has had a decentralizing effect on the dissemination of information, it has come with a cost. Bullshit spreads more easily in a massively networked, click-driven social media world than in any previous social environment. We have to be alert for bullshit in everything we read.