



6

Truth, Information & Democracy

What you are will show in what you do.

—Thomas Edison

The reverberations of Silicon Valley’s ascendency can be felt in nearly every corner of our lives. Countless technologies allow us to seamlessly connect with our loved ones around the world, work remotely, access the world’s information quickly and easily, and enjoy the immense scientific and medical advancements that technology affords us. In some cases, tech is also working to address what should be humanity’s primary concern—making the world more sustainable and environmentally sound—though these companies are still in the minority.

Whether you’re a fan of Isaac Newton, Eastern Mysticism, or the Hamilton soundtrack, you’ll know that every action has an equal, opposite reaction; forces come in pairs, and rarely is something purely a force for good. Which is why it shouldn’t surprise us that for all the benefits technology provides us, there are an equal number of drawbacks and challenges that arise when the world changes at the rate it has without attendant oversight or accountability. The scope of these side effects is

vast and, in many cases, hugely complex. This section will illustrate some of the most pervasive problems and challenges that have resulted from technology and the ways in which Silicon Valley's values, behaviors, and psychology have contributed to them.

The products that emerge from Silicon Valley impact us both on a macro, social level, and in more individual and personal ways. This section will examine both, starting with an overview of the industry's more global effects, including democracy, misinformation, economic inequality, and job displacement, and then looking at more individual impacts, such as health and mental health, relationships, and cognition. While these phenomena are not the intended effects of the technology that brought them about, but rather side effects of other motivations, they are, nonetheless, socially destructive, urgent problems that require our immediate attention. Facebook and Twitter did not set out to break democracies and incite hatred; YouTube did not plan to drive extremism; Instagram didn't intend to increase anxiety and depression in young people. Nor did the tech industry as a whole plan to drive inequality and employment instability, demolish individual privacy, create a two-class job market, spread misinformation, upend human connection, or negatively affect our cognition. Each of these is a side effect of other aims and decisions made in the service of certain motivations. The following pages will detail each of these and explore how the psychology of the industry has contributed to the unintended but profound consequences we are now enduring as a result of the technology we have embraced, beginning with an exploration of the crisis of truth, information, and democracy, and the business model that underpins it all.

The Dark Arts

Before we attempt to understand the ways in which technology platforms have undermined social institutions and driven social harm, it's useful to grasp the method by which many of the companies complicit in these problems make the majority of their money. The business model is, in some ways, shockingly simple. In 2018, during Facebook's first

congressional hearing, Senator Orrin Hatch asked Mark Zuckerberg how his company sustained a business model in which users didn't pay for the service. Zuckerberg succinctly and honestly replied, "Senator, we run ads." Indeed they do. Advertising accounted for 99 percent of Facebook's 2019 Q1 revenue. The exchange between Hatch and Zuckerberg was mocked for weeks and the hearing largely considered a failure all around; an unblinking Zuck was caricatured as an automaton and Congress as a bunch of out-of-touch fuddy-duddies. The implications of Facebook's business model were never fully fleshed out that day, thanks to everyone's love of the sound of their own voice. It was, however, the most important question and the most lucid answer of the entire hearing.

In 1998, Sergey Brin and Larry Page, then PhD students at Stanford, released a paper about their new project, Google. Google was a search engine prototype designed to organize academic search results. In the paper, the pair acknowledged the increasing commercialization of the internet, and warned against both the "black box" effect of algorithmic search engines and a business model where search could be commoditized and controlled by advertisers.

Aside from tremendous growth, the Web has also become increasingly commercial over time. In 1993, 1.5% of Web servers were on .com domains. This number grew to over 60% in 1997. At the same time, search engines have migrated from the academic domain to the commercial. Up until now most search engine development has gone on at companies with little publication of technical details. This causes search engine technology to remain largely a black art and to be advertising oriented (see Appendix A in the full version). With Google, we have a strong goal to push more development and understanding into the academic realm.¹

In October 2000, Google began selling advertising on its platform, embracing the very business model that, less than two years previously, Brin and Page had warned against. As Google's user base grew and the technology driving it advanced, the company invented and pushed targeted advertising, as described by former CEO and Executive Chairman Eric Schmidt:

You have to have both a technological idea, but you also have to have a significant change in the way the revenue will come in. In our case, we invented targeted advertising, which is really much better than untargeted advertising. And that's what happened, and we rode that really, really hard. That gave us this engine.²

The departure from Brin and Page's original intentions illustrates a significant shift in motivation. Noam Cohen notes that the original motivations of Google's founders were predominantly academic and prosocial in nature and that Brin and Page repeatedly "stressed the social benefits of their new search engine," which they promised "would be open to the scrutiny of other researchers and wouldn't be advertising-driven."

The public needed to be assured that searches were uncorrupted, that no one had put his finger on the scale for business reasons. To illustrate their point, Mr. Brin and Mr. Page boasted of the purity of their search engine's results for the query 'cellular phone'; near the top was a study explaining the danger of driving while on the phone. The Google prototype was still ad-free, but what about the others, which took ads? Mr. Brin and Mr. Page had their doubts: 'We expect that advertising-funded search engines will be inherently biased towards the advertisers and away from the needs of the consumers.'³

The advertising-driven business model that Brin and Page initially denounced is the same model that in 2018 made Google \$116,320,000,000, or 71% of their total revenue.^{4,5} It is the same model enthusiastically acknowledged by Zuckerberg during Facebook's Congressional and Senate hearings. It is also the same model that is responsible for the majority of the social upheaval that has been experienced over the past decade.

Google, Facebook, and Twitter remain free at the point of use because of their reliance on advertising revenue. The easiest way to explain the relationship between tech companies and advertising, according to Zeynep Tufekci, a Harvard researcher and author of *Twitter and Tear Gas: The Power and Fragility of Networked Protest*, is to think of the former as advertising brokers.

These companies—which love to hold themselves up as monuments of free expression—have attained a scale unlike anything the world has ever seen; they’ve come to dominate media distribution, and they increasingly stand in for the public sphere itself. But at their core, their business is mundane: They’re ad brokers. To virtually anyone who wants to pay them, they sell the capacity to precisely target our eyeballs. They use massive surveillance of our behavior, online and off, to generate increasingly accurate, automated predictions of what advertisements we are most susceptible to and what content will keep us clicking, tapping, and scrolling down a bottomless feed.⁶

The dynamic Tufekci describes can only be accomplished because of the reams of data collected by Google and Facebook, which include not only the basics, like your name, age, gender, location, interests, socioeconomic level, income level, ethnicity, and occupation, but also more private and inferred information, including your political affiliation, search and browser history, purchases, likes, private messages and emails, internet activity, and your location, every second of every day.

Silicon Valley is an extractive industry. Its resource isn’t oil or copper, but data. Companies harvest this data by observing as much of our online activity as they can. This activity might take the form of a Facebook like, a Google search, or even how long your mouse hovers in a particular part of your screen. Alone, these traces may not be particularly meaningful. By pairing them with those of millions of others, however, companies can discover patterns that help determine what kind of person you are—and what kind of things you might buy.⁷

When someone wants to advertise their service, product, or message, Facebook or Google are the most obvious platforms to utilize, as they are able to group their users into highly defined audiences, based on millions of data points they have stored over time.

The difference between selling user data and allowing access to user data for a fee is a significant distinction, which has, until recently, allowed companies to hide behind how they define user privacy. When Zuckerberg vehemently claims, as he did in a 2018 interview with Kara Swisher, that Facebook does not “sell” user data, he is technically being accurate.⁸ What

free platforms such as Facebook, Google, and Twitter do not shout about, however, is why they hoard and protect your data so aggressively—not because selling it would be morally reprehensible, but rather, as New Shift Company explains, because it wouldn't be nearly as profitable. “Facebook does not sell your data. It protects your data like Gollum holding the ring. Selling your data would not be nearly as profitable as leasing access to you, via advertising—over and over again.”⁹

What Facebook and Google realized long before the rest of us, and proceeded to build billion dollar businesses off the back of, is the supreme value of data in the information age. There is simply nothing that has lit a fire under capitalism more in the past century than the ability of tech giants to target consumers, using the very data their users agree to give away for free. As Eric Schmidt bragged, Google was the first to capitalize on this realization, which was soon emulated by Twitter and Facebook. Collectively, these companies perfected the art of personalized, targeted advertising.

Facebook was a social network where legions of users voluntarily offered personally identifying information in exchange for the right to poke each other, like each other, and share their baby pictures with each other. Facebook's founders knew their future lay in connecting that trove of user data to a massive ad platform. In 2008, they hired Sheryl Sandberg, who ran Google's advertising operation, and within a few years, Facebook had built the foundation of what is now the most ruthlessly precise targeting engine on the planet.¹⁰

Given the now inextricable relationship between advertising and technology, it is perhaps unsurprising that the ethos of the former would bleed into that of the latter. Professor William Irwin, who teaches philosophy at King's College in Pennsylvania, explained to Time magazine reporter Coeli Carr that the advertising industry as a whole “has historically wrestled with questionable ethics and a lack of self-awareness,”¹¹ the cultural effects of which are now beginning to materialize.

The purpose of advertising, for those of us who didn't major in marketing or watch *Mad Men*, is to change behavior. Historically, marketing campaigns in non-digital spaces focused on changing our behavior in

order to make us buy stuff: McDonalds wanted you to buy burgers, Estee Lauder wanted you to buy makeup, and the Marlboro Man wanted to get you hooked on cigarettes. Through the whole of the twentieth century, such advertising campaigns were largely unable to reach their target audience in very meaningful ways. Digital advertising, however, brought sellers closer to their consumers than ever before. Instead of marketing diapers to the general public and hoping expectant couples were on the receiving end of the campaign, diaper companies could now target potential parents based on their past purchases, age, browser history, or any number of demographic factors. In the last few years, a handful of brilliant but morally reprehensible assholes realized targeted advertising could also be used, and indeed was especially effective, at changing public opinion and manipulating emotions—particularly in the political sphere.

Facebook's business lies in influencing people. That's what the service (sic) it sells to its customers—advertisers, including political advertisers. As such, Facebook has built a fine-tuned algorithmic engine that does just that. This engine isn't merely capable of influencing your view of a brand or your next smart-speaker purchase. It can influence your mood, tuning the content it feeds you in order to make you angry or happy, at will. It may even be able to swing elections.¹²

As its audience was lulled into watching cat videos, posting baby photos, and chatting with long-lost school friends, the bargain between using free services in exchange for personal data seemed relatively innocuous. It is now glaringly apparent that the true cost of free services is far steeper than anyone anticipated. Jaron Lanier, author of *Who Owns the Future?*, describes how the tension between Silicon Valley's advertising business model and its original social ideals led to what has become the most effective social manipulation tool in human history:

there's only one way to merge [socialism with libertarian ideals], which is what we call the advertising model, where everything's free but you pay for it by selling ads. But then because the technology gets better and better, the computers get bigger and cheaper, there's more and more data—what started out as advertising morphed into continuous behavior modification on a mass basis, with everyone under surveillance by their devices and

receiving calculated stimulus to modify them. So you end up with this mass behavior-modification empire, which is straight out of Philip K. Dick, or from earlier generations, from *1984*. It's this thing that we were warned about. It's this thing that we knew could happen.... And despite all the warnings, and despite all of the cautions, we just walked right into it, and we created mass behavior-modification regimes out of our digital networks. We did it out of this desire to be both cool socialists and cool libertarians at the same time.¹³

Advertising, according to Lanier's assessment, has become the implicit compromise that allows Silicon Valley to simultaneously embrace its socialist roots, entrepreneurial spirit, and libertarian ideals.

As the human and social costs continue to pile up, we might wonder how long it will take to change the "free" advertising-driven business model to something more socially responsible. Barring government regulation or a complete overhaul of the industry's values and psychology, the answer, I'm afraid, is never. As Lanier explains, there is simply no incentive for Facebook, Google, Twitter, or the legions of companies who make money from advertising to change of their own accord. In order to change the practices of the industry, the drivers of its behavior must change.

The Price of Free

The social costs of big tech's current priorities and motivations are plentiful, and examples of its failures in civic responsibility are vast. Umair Haque has argued that "social media's effects on social and civic well-being are worse than they are on emotional well-being: they last longer, do more damage," and require more substantial clean-up and rebuilding efforts,¹⁴ while Tristan Harris argues that society can no longer "afford the advertising business model."

The price of free is actually too high. It is literally destroying our society, because it incentivizes automated systems that have these inherent flaws. Cambridge Analytica is the easiest way of explaining why that's true. Because that wasn't an abuse by a bad actor—that was the inherent platform. The problem with Facebook is Facebook.¹⁵

Harris's argument that Facebook's issues are built into the company's DNA raises a problematic truth: that so long as the fundamental structures, financial incentives, and business models of such companies do not change, the issues we're uncovering will continue to pile up. We are only beginning to collectively realize that "free" services like Facebook and Google are never really free, that we simply offer payment in different ways, such as with our time, attention, and the cohesion and civility of our society.

Of all the problems and PR disasters that have transpired in Silicon Valley, one of the most unsettling is the role technology and social media companies have played in the erosion of democracy. The breakdown of civic discourse, the dissemination of false information online, the targeting of individuals with particular information, and the polarizing effect of social media platforms each contribute to the widening gap between what is needed for democracy to function effectively and the disruptive technological factors at play.

While there has always been inaccurate information in circulation, never before has there been so much, and never before has it been weaponized at scale. The modern phenomenon of misinformation and disinformation have proliferated on the internet, where there remains no system by which to determine the accuracy of information online. We may instinctively know that certain sites or sources, such as peer-reviewed journal articles, are based on controlled research and fact, just as we may know that other sites are meant to be satirical. Wading through the other two plus billion websites on the internet, however, it can often be difficult to tell a reputable site from a biased or intentionally fake news site. Even if we are aware of fact checking sources, such as snopes.com or mediabiasfactcheck.com, the majority of us rarely bother to check the quality of every site we visit. The lack of built-in quality assurance on the internet makes the web simultaneously a treasure trove of high-quality information—academic research, scholarly essays, investigative journalism, and books—and a landmine of bad and misleading information.

If the extent of the problem was simply how best to categorize and sort information based on its quality, we wouldn't have a terribly treacherous path ahead. Sure, it would be a pain to fact-check the whole of the internet, but it wouldn't be impossible. The problem facing democracy is

complicated by the ways in which the advertising ecosystem of the internet dictates the flows of information online. Algorithms, including those used on Twitter, Google, YouTube, Facebook, and Instagram, are designed to show you whatever will engage you as much as possible. This is because engagement—defined as more clicks, more interactions, and more time spent on apps or websites—is synonymous with greater ad revenue for ad brokers like Facebook and Google; the more engaged you are, the more ads you see, the more money they make. Engagement, according to technology consultant and writer Tobias Rose-Stockwell is “the currency of the attention economy,”¹⁶ which means it is in the financial interests of tech companies that are reliant on advertising to keep us online and engaged in whatever way possible for as long as possible. The impacts of this model are corrupting, to put it lightly. Having ads dictate the flows of information has resulted not only in the spread of misinformation and disinformation, but also in the prioritization of sensationalized content, filter bubbles, and the ability to micro-target individuals with particular information.

Given that the internet boasts to hold the sum total of the world’s knowledge, it has become the place we increasingly inhabit to gather (and in some cases mainline) information. Where facts about the world around us used to be mediated by a variety of factors and channels, such as local news (subsidized by the government and/or paid for by mass, untargeted TV commercials), national newspapers (paid for by the reader), and the time it took to report a story (lag time that ensured increased accuracy of information), this is no longer the case. The uncertain quality of the content we see online has spawned the phenomena of misinformation, disinformation, “alternative facts,” and “fake news,” each of which is either born of or intensified by the advertising ecosystem that drives the internet. At some point between its conception by Tim Berners-Lee as a mechanism to share ideas and information, and its commoditization and commercialization, the internet has become a place where truth is subjective.

Data from Pew Research Center suggests two-thirds of adults in the U.S. get news from social media, and about half rely on Facebook for news (half of these users get their news from Facebook alone, while just one-in-five rely on three or more sites for news).¹⁷ These findings suggest

several troubling implications: that the content prioritized by Facebook's algorithm contributes significantly to the average American's news consumption; that many users do not rely on a diverse sample of news from various online sources; and that advertising-driven businesses are feeding a substantial portion of our collective news diet. Combine this with the unchecked quality of most news online, and you get a veritable dumpster-fire of good, bad, and downright ugly information, all mixed up together in one unpoliced internet. The result, according to Rose-Stockwell, has been the normalization of propaganda and the evisceration of traditional journalism.

Today we have democratized propaganda—anyone can use these strategies to hijack attention and promote a misleading narrative, a hyperbolic story, or an outrageous ideology—as long as it captures attention and makes a profit for advertisers. Journalism—the historical counter to propaganda—has become the biggest casualty in this algorithmic war for our attention. And without it, we are watching the dissolution of a measured common reality.¹⁸

A 2016 Stanford study found that middle school through university students “could not distinguish between news and sponsored content, source evidence, or evaluate claims on social media.”¹⁹ The results suggest that as entertainment, click-bait, and opinions continue to intermingle with and masquerade as evidence-based information, many young people are not adequately equipped to question the source, accuracy, or quality of information they encounter online.

While there are positive effects unique to the rise of digital information, including the ease of research, the accessibility of information, and proliferation of traditionally marginalized voices and issues, there are also significant costs. The spread of sensationalized and false content has led to a less informed populace, more black and white thinking, and the phenomenon of denialism. The creation of filter bubbles has contributed to a more proliferated, angry, and fragmented society. Our hyper-connected world has allowed for the normalization of extremist content, as fringe or fanatical views—which would naturally be drowned out in a traditional community—can more easily come together in digital spaces

worldwide. Add to all this the ability to target information to specific individuals for nefarious purposes, and you get a fair bit of confusion and chaos.

Confusion and Chaos

Perhaps the first thing to understand about false information is its sheer ability to spread, virus-like, to the feeds of unsuspecting and unprepared consumers. A 2018 study on the spread of fake news, the largest ever of its kind, demonstrated that lies spread significantly faster than truth.²⁰ The researchers of the study explain the reason for this lies in “the degree of novelty and the emotional reactions of recipients” to false stories, as well as the tendency for bots, which are programmed to disseminate reactive content, to spread fake stories. Using 11 years of data from Twitter, which included over 4.5 million tweets, along with information from six different independent fact-checking organizations, researchers demonstrated just how profoundly facts fail compared to fictions:

A false story reaches 1,500 people six times quicker, on average, than a true story does. And while false stories outperform the truth on every subject—including business, terrorism and war, science and technology, and entertainment—fake news about politics regularly does best. Twitter users seem almost to *prefer* sharing falsehoods. Even when the researchers controlled for every difference between the accounts originating rumors—like whether that person had more followers or was verified—falsehoods were still 70 percent more likely to get retweeted than accurate news.²¹

There are no shortage of examples to illustrate the ways in which false, biased, and misrepresented information have spread on social networks like Facebook, Twitter, and YouTube. Just take Twitter’s worst offender, Donald Trump, whose lies and misrepresentations of facts spread like digital wildfire thanks to the digital bullhorn that the platform affords him.

The second important thing to understand about false information is that there are different kinds. The person who knowingly concocts a false news article is sharing a different type of information than the person

who reads it, assumes it's accurate, and shares it with his or her online community. Johns Hopkins Sheridan Libraries separate the different types of information available online into four categories, depending on their source, purpose, and quality: information, misinformation, disinformation, and propaganda. Information is pretty straightforward: it's what we're mostly after when we go online, even if we don't always find it. Information communicates factual data or knowledge about a subject or event, with reference to the relevant context and evidence. Good information should both be accurate and free from bias; if it is in some way biased, it should make note of it clearly.²² The remaining types of information—disinformation, misinformation, and propaganda—deviate from this definition in one or more ways.

Johns Hopkins cites propaganda as a commonly misunderstood and misused term, due to its historical association with the Nazi government in Germany.

[M]any people associate propaganda with inflammatory speech or writing that has no basis in fact. In reality, propaganda may easily be based in fact, but facts represented in such a way as to provoke a desired response.²³

Propaganda, which is still commonly employed in campaign speeches and political statements, is information with some basis in reality, but which is presented in a misleading way to influence attitudes or behavior. Disinformation, by contrast, “refers to disseminating deliberately false information,” in which the sharer is aware of and complicit in spreading falsehoods. Johns Hopkins calls disinformation the “lowest of the low,” and explains that it is most typically circulated by “individuals or institutions [in order] to say or write whatever suits a particular purpose, even when it requires deliberate fabrication.”²⁴ (The authors describe the Internet as “an excellent vehicle for disinformation.”) A third type of inaccurate information is misinformation, which is similar to disinformation in that it is false, but differs in that the individual or group sharing it is unaware of its inaccuracy. Information may begin as disinformation—that is, information that is knowingly false—and then be circulated as misinformation by people who see it, assume its veracity, and share it with others in good faith. The authors cite misinformation as

the most difficult subtype of bad information to identify, suggesting it may also be the most dangerous, due to its obscurity and the fact that it is typically delivered with good intentions.

Examples of the swift and steady spread of misinformation, disinformation, and propaganda online are too numerous to list; so, too, are instances of individual harm, hate speech, and democratic regression that accompany them. Some of the most well-known and horrible examples of the effects of misinformation on democracy, social unrest, and human rights violations include:

1. The spread of false information via Facebook in Myanmar, in which Buddhist extremists spread hate speech and false information about the country's Rohingya population, which resulted in riots, executions, rape, the burning of hundreds of Rohingya villages, and the ethnic cleansing of the country's Muslim minority.^{25,26} As of September 2018, 725,000 Muslim Rohingya had fled Bangladesh and 10,000 were confirmed dead (widely considered a conservative estimate). A report by an independent Human Rights Council concluded in a UN report that Facebook had been central to the campaign against the Rohingya and had proved itself to be "a useful instrument for those seeking to spread hate."²⁷ Ashin Wiratho, one of the leaders of the anti-Rohingya movement, likewise credits Facebook with its "success": "If the internet had not come to [Myanmar], not many people would know my opinion and messages ... The internet and Facebook are very useful and important to spread my messages."²⁸
2. In Sri Lanka, disinformation circulated on Facebook by Sinhalese nationalists was used to incite violence and hatred toward the country's Muslim population. Posts included messages such as, "Kill all Muslims, don't even save an infant" and instructions to "reap without leaving an iota behind."²⁹ Those who reported such posts, using Facebook's on-site reporting tool, were often told that they did not violate the platform's standards.³⁰ One man behind the violence, Amith Weerasinghe, shared videos, hateful messages, and warnings with thousands of followers on Facebook, which were reported to the platforms by researchers in Colombo, but never taken down. "Over the next three days, mobs descended on several towns, burning mosques, Muslim-owned shops and homes. One of those towns was

Digana. And one of those homes, among the storefronts of its winding central street, belonged to the Basith family. Abdul Basith, a 27-year-old aspiring journalist, was trapped inside. ‘They have broken all the doors in our house, large stones are falling inside,’ Mr. Basith said in a call to his uncle as the attack began. ‘The house is burning.’ The next morning, the police found his body. In response, the government temporarily blocked most social media. Only then did Facebook representatives get in touch with Sri Lankan officials, they say. Mr. Weerasinghe’s page was closed the same day.”³¹

3. In Indonesia, false messages were spread via Facebook and Whatsapp (a subsidiary of Facebook) that gangs were stealing and killing children and selling their organs. Villagers in nine rural Indonesian communities, upon seeing outsiders enter their towns, lynched those they presumed were coming to murder their children.³²
4. India, the country with the largest Facebook user base in the world (220 million), experienced a string of lynchings similar to those in Indonesia,^{33,34} as well as a misinformation campaign by fringe political parties and religious extremists intended to sow discord and false information to the 49 million voters in India’s Karnataka region. Both phenomena were carried out primarily via Whatsapp, where false information was forwarded en masse, with no means to trace or stop the spread of untrue messages.³⁵ Two men who lost their lives in the wake of one such misinformation campaign were Abijeet Nath and Nilotpal Das, who were driving back through India’s Assam province from a waterfall they had visited. When the two men stopped in a village to ask for directions, they were pulled from their vehicle and beaten to death by a mob, who assumed they were coming after the area’s children.³⁶ False information circulated by right-wing Hindu groups were similarly responsible for inciting violence toward Muslim populations. One example included “a grisly video that was described as an attack on a Hindu woman by a Muslim mob but was in fact a lynching in Guatemala. One audio recording on the service from an unknown sender urged all Muslims in the state to vote for the Congress party ‘for the safety of our women and children.’.... Like the rest of India, Karnataka is a Hindu majority state. A staple of electoral politics here is pitting Muslims against Hindus, and various Hindu castes against each other.”³⁷

5. In Brazil, an outbreak of yellow fever was thought to be the result of misinformation spread on Whatsapp regarding anti-vaccine propaganda.³⁸
6. In the Philippines, Rodrigo Duterte, along with other candidates, was trained by several Facebook employees during the country's 2016 presidential campaign, which instructed candidates how to set up accounts, drive engagement, and attract followers. Using this knowledge, Duterte's office mobilized itself around an anti-drug and criminal justice campaign, many of which were seeped with "aggressive messages, insults, and threats of violence," as well as disinformation, such as the false claim that the Pope Francis endorsed Duterte.³⁹ Using fear and falsehoods to drive his message, Duterte was elected in 2016. "He told Filipinos the nation was being ruined by drug abuse and related crime and promised to bring to the capital the merciless strategy he had employed in Davao. Soon, Duterte's death squads prowled the streets at night in search of drug dealers and other criminals. Images of blood-smeared bodies slumped over on sidewalks, women cradling dead husbands, and corpses in satin-lined caskets went viral. As the bodies piled up—more than 7,000 people have been killed as part of Duterte's war on drugs—the social media war escalated."⁴⁰ Still, the relationship between Facebook and Duterte appears strong; in November 2017, the social network agreed to a partnership with Duterte's government, in which the company agreed to fund the development of underwater fiber cables in the Philippine's Luzon Strait and "provide a set amount of bandwidth to the government."⁴¹

The proliferation of propaganda, misinformation, and disinformation on the web is at the heart of the internet's collision with democracy, civic order, and the degradation of human rights. Driving this troubled dynamic is the prioritization of engagement and push for growth into foreign markets, driven by the profit-oriented motivations of the industry.

To their credit, Facebook, Twitter, and YouTube have made some efforts to address the spread of false information. In India, Facebook now limits Whatsapp groups to five people (previously capped at 256 members/group), in an attempt to mitigate the spread of misinformation and rumors on its popular messaging app. The company has also

conducted safety workshops in the Philippines for journalists, non-profit employees, and students focused on digital literacy and safety. YouTube released its plans to help ebb the flow of misinformation by investing \$25 million in an effort to promote more legitimate and trusted news sources on its platform, which is part of a larger \$300 million plan to address misinformation within its parent company, Google.⁴² The issue remains, however, that such efforts are at odds with the core business model of such companies. As Olivia Solon aptly points out, “how can [Silicon Valley] condemn the practice on which its business model depends?”⁴³

In addition to reconciling the tension between the profits generated by misinformation and the safety of their users, another question many tech companies must grapple with is their position on disinformation and those who spread it. Currently, unless something violates a platform’s specific policies, which tend to focus on hate speech and inappropriate content, there is no unified position on how to treat false information online. Even now that the disastrous effects of misinformation and disinformation campaigns have been exposed, platforms like Facebook, YouTube, and Twitter have failed to categorically denounce and expunge such information from their sites, seemingly unsure how to balance free speech with their users’ safety.

But why does it matter? Surely everyone has a right to scream and shout their opinion into the void of social media, even if the opinion they espouse or the article they post is false or misinformed. As Mark Zuckerberg has argued, there is merit in leaving such information on Facebook’s platform, even when that information is patently false and perpetuates hatred, conspiracy theories, and anti-Semitism.

[T]here’s a set of people who deny that the Holocaust happened. I find that deeply offensive. But at the end of the day, I don’t believe that our platform should take that down because I think there are things that different people get wrong. I don’t think that they’re *intentionally* getting it wrong... What we will do is we’ll say, “Okay, you have your page, and if you’re not trying to organize harm against someone, or attacking someone, then you can put up that content on your page, even if people might disagree with it or find it offensive.”⁴⁴

The suggestion that Holocaust denial should even be included in the realm of “information” is a startling assertion, with deeply troubling implications. Facebook’s goal, according to Zuckerberg, is not to be the judge of truth and accuracy, but “to prevent hoaxes from going viral and being widely distributed. The approach that we’ve taken to false news is not to say, you can’t say something wrong on the internet. I think that that would be too extreme.”⁴⁵ Facebook demotes identified hoaxes but does not necessarily remove them unless moderators believe content will “result in real harm, real physical harm, or if you’re attacking individuals, then that content shouldn’t be on the platform.”⁴⁶

The parameters Zuckerberg and his company cite constitute an extremely literal understanding of harm, free from nuance and more socially contextualized understandings. The most general understanding of harm is typically that which could reasonably cause either physical, mental, or emotional damage. What constitutes mental and emotional damage however, can be rather subjective sticking points. Moral philosopher Bernard Gert defined harm as that which causes pain, death, disability, or the loss of freedom or pleasure, while political and legal philosopher Joel Feinberg determined harm also included what he called “welfare interests,” which took into consideration harm that affected one’s

intellectual acuity, emotional stability, the absence of groundless anxieties and resentments, the capacity to engage normally in social intercourse and to enjoy and maintain friendships, at least minimal income and financial security, a tolerable social and physical environment, and a certain amount of freedom from interference and coercion.⁴⁷

The subversion of democracy via orchestrated disinformation campaigns can undoubtedly, then, be counted as harmful? The eradication of trust, also, is surely dangerous to ordered social institutions? The challenge of making democracy work amidst the backdrop of the digital age is a hugely complex task. In *The People vs. Tech*, Jamie Bartlett argues that the primary difficulty in reconciling democracy and technology comes back to the principles that govern each institution and the fact that the rules that guide Western democracies are fundamentally at odds with those that govern cyberspace.⁴⁸

Out of Sight, Out of Mind

Related to the problem of misinformation and disinformation is the phenomenon of denialism, which has really come into its own thanks to the propagation of fake and misleading information. Denialism can be defined as the “refusal to admit the truth of a concept or proposition that is supported by the majority of scientific or historical evidence.”⁴⁹ As Keith Kahn-Harris explains, while denialists were once confined to the fringes of public discourse, they now occupy a much more visible and central position as a result of the internet’s global reach and connectivity.

As information becomes freer to access online, as “research” has been opened to anyone with a web browser, as previously marginal voices climb on to the online soapbox, so the opportunities for countering accepted truths multiply. No one can be entirely ostracised, marginalised and dismissed as a crank anymore. The sheer profusion of voices, the plurality of opinions, the cacophony of the controversy, are enough to make anyone doubt what they should believe.⁵⁰

Khan-Harris goes on to explain the conflict between science and wishful thinking that underlies many denialist positions.

Denialism is not stupidity, or ignorance, or mendacity, or psychological pathology. Nor is it the same as lying. Of course, denialists can be stupid, ignorant liars, but so can any of us. But denialists are people in a desperate predicament ... a very modern predicament. Denialism is a post-enlightenment phenomenon, a reaction to the “inconvenience” of many of the findings of modern scholarship. The discovery of evolution, for example, is inconvenient to those committed to a literalist biblical account of creation. Denialism is also a reaction to the inconvenience of the moral consensus that emerged in the post-enlightenment world.⁵¹

The proliferation of extreme, inaccurate views, taken to its extreme, informs what Khan-Harris calls post-denialism, in which the world is fashioned to take any form the narrator desires.

[I]ts methods liberate a deeper kind of desire: to remake truth itself, to remake the world, to unleash the power to reorder reality itself and stamp one's mark on the planet. What matters in post-denialism is not the establishment of an alternative scholarly credibility, so much as giving yourself blanket permission to see the world however you like.... Whereas denialism explains—at great length—post-denialism asserts. Whereas denialism is painstakingly thought-through, post-denialism is instinctive. Whereas denialism is disciplined, post-denialism is anarchic. The internet has been an important factor in this weakening of denialist self-discipline. The intemperance of the online world is pushing denialism so far that it is beginning to fall apart. The new generation of denialists aren't creating new, alternative orthodoxies so much as obliterating the very idea of orthodoxy itself. The collective, institutional work of building a substantial bulwark against scholarly consensus gives way to a kind of free-for-all.⁵²

Denialism represents not only the erosion of information, but also the collective breakdown of order, truth, and the psychological orientation these provide. Two modern and devastating examples of denialism include climate change and the anti-vaccination movement. While certain minority groups have historically expressed scepticism at the idea of global warming, modern scientific evidence and research points unequivocally to the dire and urgent problem of climate change and its effects on our ecosystems, weather, and the inhabitability of the planet. It is a terribly inconvenient truth, to quote Al Gore's 2006 documentary on the same topic, that human industrial activities have caused such severe damage to the planet. No one enjoys thinking about the fact that we are responsible for raising global carbon dioxide, methane, and nitrous oxide levels to the point that they have collectively impacted earth's rapidly rising temperature or threatened millions of plant and animal species.^{53,54} It's not fun to think about, but it's true, and it remains our most pressing global problem. None of this, however, can stop tweets from Donald Trump claiming climate change was invented by the Chinese in order to weaken the U.S. manufacturing industry.⁵⁵ The assertion, like many Trump makes on Twitter, is accompanied by neither facts nor research; it is merely a succinct, grossly inaccurate claim, sent out on a whim and compressed to 280 characters or less. And Trump is not alone. Denialists, conspiracy theorists, and those who wish to sow discord on any topic can

now do so from anywhere, for any reason, to any end, with the push of a button.

Denialist movements have also been used to discredit and misinform global populations about vaccinations. The World Health Organization linked a 2018–2019 surge in measles cases across Europe and the U.S. to widespread erroneous information propagated online, regarding claims that the vaccine was ineffective or even harmful. The false information has contributed to a significant decrease in the uptake of the life-saving MMR vaccine, which has contributed to a staggering 60,000 cases of infection in 2018 (twelve times the number of cases reported in 2016) and dozens of deaths in Europe alone.⁵⁶

False information about MMR continues to be spread online, particularly on social media, giving a platform to the anti-vaccination movement to push erroneous claims. Some of the posts have hundreds of thousands of ‘likes’ and include false claims that healthcare professionals have been lying to the public or that immunisation injections amount to nothing more than ‘poison being pumped into people’s bloodstreams’.⁵⁷

Helen Stokes-Lampard, a professor of medicine and Chair of the Royal College of General Practitioners in the U.K., who has studied the disease’s recent progression and its ties to denialist propaganda, also cites the role of false information online in bolstering the re-emergence of what is a lethal but entirely preventable disease.⁵⁸ Stokes-Lampard explains that the underlying issue remains a “lack of regulation and enforcement around this material online,” which has allowed anti-vax “groups to build momentum without the opportunity for any form of meaningful evidence-based rebuttal.”⁵⁹

The epidemic of faulty information and the unwillingness to police it ties back to the tech industry’s focus on individualism, its motivation for profit, and its advertising-driven business model. Bold, ridiculous claims naturally attract our attention and trigger our emotions (rage, concern, fervent agreement); the natural result is that these are shared more, which in turn results in more engagement and more profit. While the leaders of YouTube, Twitter, and Facebook probably don’t consciously want to contribute to the biggest outbreak of preventable measles deaths in history,

the option to aggressively police false information is at odds with both their business models and belief in individual (if erroneous) expression.

Filter Bubbles

Related to the resurgence of denialism is the phenomenon of filter bubbles. The basic concept underlying filter bubbles is that we tend to gravitate towards ideas that confirm what we already believe, which eventually blocks out information that is not in line with our existing opinions. Given the breadth of information available to us, we might assume that the spoils of knowledge would make us more, rather than less informed; after all, the central attraction of the internet lies in its promise of information. With all that knowledge at our fingertips, surely we would become more informed, more knowledgeable, wiser versions of ourselves? Not so, according to historian Timothy Snyder, who explains that

[I]n assuming that the Internet would make us more rather than less rational, we have missed the obvious danger: that we can now allow our browsers to lead us into a world where everything we would like to believe is true.⁶⁰

The fact that we can find any and all information online means that there is likely to be support for every position, regardless of accuracy. Though the web was built in the name of academic research exchange, the quality of the information that populates the modern internet is, quite obviously, of different calibers. As a result, we are constantly presented with conflicting information, of different qualities, in unfathomable quantities, all the time. The task of sorting through this mass of information is too much for most of us to bear, and the path of least cognitive resistance is often to gravitate towards what already feels comfortable and familiar, rather than challenge ourselves to explore or consider a new position.

The reason we are more likely to see news and information that corresponds to our existing worldview has everything to do with the algorithms and the business models of tech companies who rely on advertising revenue. The algorithms that run the sites we use to peruse information,

such as Facebook and YouTube, gather our data in an effort to predict what they think we are most interested in or want to see, in order to drive engagement, show us “relevant” content, and ultimately generate more time spent on the platform. The result is that we start to see only what we want, rather than a range of accurate, balanced information. The more refined our algorithms become, the less likely we are to come across things outside our comfort zone, which might challenge our beliefs. And voila! There you have a filter bubble.

The implication here is actually quite scary. Staying inside our bubbles may hinder our ability to think differently, consider another’s perspective, or intelligently defend our own. The more refined the algorithms dictating our newsfeeds become, the more embedded we become in our virtual echo chambers, largely unaware of points of view that differ from our own. Living solely in the company of our own opinions may feel good, but it’s not doing us any favors as rational human beings. Research suggests that the more content we have to sort through, the worse we become at differentiating between high- and low-quality information. A 2017 study found that while confirmation bias had originally evolved to help us quickly dispel useless and false information, in the context of the internet, and particularly with social media, “such a bias easily leads to ineffective discrimination.”⁶¹ Gene Demby of National Public Radio explains that instead of bridging opinions and encouraging conversation as social media platforms claim to do, these sites in fact make us less aware of different opinions and more insulated in our own.⁶²

The other troubling consequence of self-selecting our information is the phenomenon of algorithmic extremism. John Naughton’s interview with Zeynep Tufekci outlines the means by which YouTube’s algorithm impels us down rabbit holes of increasingly extreme content:

Tufekci tried watching videos on non-political topics such as vegetarianism (which led to videos about veganism), and jogging (which led to items about running ultramarathons). “It seems,” she reflected, “as if you are never ‘hardcore’ enough for YouTube’s recommendation algorithm. It promotes, recommends and disseminates videos in a manner that appears to constantly up the stakes. Given its billion or so users, YouTube may be one of the most powerful radicalising instruments of the 21st century.”⁶³

The mechanism of YouTube's algorithm works in two ways: first, it takes a topic we have already expressed interest in and gives us even more information on that topic, knowing we're already intrigued. Second, it pushes us towards greater and greater extremes of our chosen topic. Interested in tea? Why not learn how to produce and distribute organic homemade kombucha? Want to know more about the Israeli-Palestinian conflict? Here, have some conspiracy theory videos about Al Jazeera being a secret tool of the Israeli government. Before we know it, we may begin incorporating new, more extreme information into our worldview that, while vaguely related to our original query, is a far cry from where we started. When it comes to the discrepancy between YouTube's financial incentives and its moral obligations to curb extremism, as Chandra Steele points out, it is obvious which is winning: "YouTube has no interest in curbing this, because it quite literally pays: The more extreme the content, the more watchers are hooked, and the more revenue is generated."⁶⁴

Us and Them

If we quietly absorbed and did not act on the information we came across online—be it extreme, biased, or just misinformed—the practical problems of extremist content would perhaps be relatively benign. Unfortunately, the innate desire to share our findings ("you *have* to see this video," "OMG, you should read this," "have you seen this study?") usually outweighs our ability to keep quiet. As social beings, we constantly seek connection with others; particularly in heightened states of emotion—be it outrage, shock, or inspiration—we are especially likely to reach out to other people. In our digital worlds, this translates into a drive to post, comment, engage, and share the radical content we find with others online (a realization social media ad brokers have exploited in full).

Sharing can go one of two ways: If our digital comrades agree with us, we become further entrenched in our position or existing thinking; if people disagree with our ideas, however, we may find ourselves defensive, angry, or indignant, and tempted to distance ourselves from such dissenters. In the real world, we would (hopefully) be unlikely to abuse, belittle,

or chastise people with views different from our own. Online, however, behind the safety of our screens, it can be a slippery slope from reasoned debate to outrage and moral fury. While the desire to arrange ourselves into factions based on similarities is a natural human tendency, Thomas Friedman explains that this propensity is drastically exacerbated on the web.

[W]hile it's true that polarization is primarily driven by our human behavior, social media shapes this behavior and magnifies its impact. Say you want to say something that is not based on fact, pick a fight, or ignore someone that you don't like. These are all natural human impulses, but because of technology, acting on these impulses is only one click away.⁶⁵

The dynamic Friedman describes can be seen in the increased callousness and lack of civility we see online: Twitter wars, bitter Facebook arguments, and endless hateful comments sections, which often descend into ridicule, name-calling, and even threats of violence. The volatility of these examples is heightened by the reduced social costs of such exchanges. As Molly Crockett, a neuroscientist and Assistant Professor of Psychology at Yale explains, “digital media may promote the expression of moral outrage by magnifying its triggers, reducing its personal costs and amplifying its personal benefits”⁶⁶; in other words, the highs we get from expressing ourselves online are amplified, while any potential social costs are diminished.⁶⁷

Offline, moralistic punishment carries a risk of retaliation. But online social networks limit this risk. They enable people to sort themselves into echo chambers with sympathetic audiences. The chance of backlash is low when you're only broadcasting moral disapproval to likeminded others. Moreover, they allow people to hide in a crowd. Shaming a stranger on a deserted street is far riskier than joining a Twitter mob of thousands. Another cost of outrage expression is empathic distress: punishing and shaming involves inflicting harm on other human beings, which for most of us is naturally unpleasant. Online settings reduce empathic distress by representing other people as two dimensional (sic) icons whose suffering is not readily visible. It's a lot easier to shame an avatar than someone whose face you can see.⁶⁸

We may feel smug and satisfied having belittled our idiot aunt who keeps posting flat earth videos, however, the pain or anger our response elicits in her (or others) and the divisiveness to our relationship is less immediately apparent.

The polarization that results from this dynamic is rooted in our tribalistic instincts, which results in both the separation from and dehumanization of others who do not share our views. Crockett writes that

there is a serious risk that moral outrage in the digital age will deepen social divides. A recent study suggests a desire to punish others makes them seem less human. Thus, if digital media exacerbates moral outrage, in doing so it may increase social polarization by further dehumanizing the targets of outrage. Polarization in the US is accelerating at an alarming pace, with widespread and growing declines in trust and social capital. If digital media accelerates this process further still, we ignore it at our peril.⁶⁹

In Crockett's estimation, the ring-fencing and protectionism that social platforms encourage is not only negative but dangerous. In an interview with Chris Cox, Facebook's former Chief Product Officer, Nicholas Thompson, editor in chief of *Wired* magazine, called filter bubbles and the spread of misinformation "the biggest problem with Facebook." Thompson asked Cox whether the radicalization and polarization of Facebook's users was tied to its business model, which Cox quickly denied. This position, of course, is overwhelmingly refuted by research, which continues to find that social media networks like Facebook and Twitter encourage rather than quell digital manifestations of tribalism, which researchers now classify as a systemic global risk.

This body of work suggests that, paradoxically, our behavioural mechanisms to cope with information overload may make online information markets less meritocratic and diverse, increasing the spread of misinformation and making us vulnerable to manipulation. Anecdotal evidence of hoaxes, conspiracy theories and fake news in online social media is so abundant that massive digital misinformation has been ranked among the top global risks for our society, and fake news has become a major topic of debate in the United States and Europe.⁷⁰

The more platforms corral the information we see, the less likely we are to engage thoughtfully, rationally, and kindly with others on a range of important issues, such as politics, climate change, and economic inequality. Instead, we are more likely to split into factions, unencumbered by many of the social norms that previously held society together.

What social media companies failed to take account of, with their mission to connect the world, is that the human brain is wired to collaborate locally (within-group cooperation) and instinctively dislike or act with hostility towards strangers (between-group competition).⁷¹ Rather than furthering its mission to “build communities,” Facebook’s ambition to bring two billion people into one gigantic virtual common room, without thought or forward-planning, has actually driven intense polarization, distrust, and prejudice. Moving fast and breaking things in the name of growth has been accomplished to startling effect; unfortunately, what has been broken are communities, trust, and informed discussion, along with the evolution of a new brand of tribalism, which spreads more easily and is more difficult to immobilize.

Our tendency towards polarization is born from our tendency to categorize. Our brains have evolved to observe something, label it, and store that information away, to be retrieved later. This saves us the trouble of assessing each situation, person, and object from scratch each time we encounter something or someone new, which helps us navigate and, in some cases, survive our environment. Most of the time, this cognitive functionality is helpful. If you see someone brandishing a gun and screaming, it’s nice your brain instinctively tells you to run away. There are, however, very real downsides to our tendency to think in categorical terms, particularly those that divide people into categories of “us” and “them.” In his book *Factfulness*, Hans Rosling describes how each of us unconsciously and “automatically categorizes and generalizes all the time” and the harm this can engender.

It is not a question of being prejudiced or enlightened. Categories are absolutely necessary for us to function. They give structure to our thoughts... The necessary and useful instinct to generalize... can also distort our worldview. It can make us mistakenly group together things, or people, or

countries that are actually different. It can make us assume everything or everyone in one category is similar. And, maybe most unfortunate of all, it can make us jump to conclusions about a whole category based on a few, or even just one, unusual example.⁷²

The act of psychological categorization is problematized by online dynamics, according to Peter Bazalgette, which tends to heighten our biases and debase our capacity for empathy. In our online worlds, tribal behavior and “unbridled prejudices” can quickly amplify, as they are “given free rein in [the] empathy-free, digital dystopia” of the web.⁷³ Former human rights lawyer Amanda Taub and journalist Max Fischer similarly suggest that Facebook’s biggest flaw and “most consequential impact may be in amplifying the universal tendency toward tribalism.”

Posts dividing the world into “us” and “them” rise naturally, tapping into users’ desire to belong. Its gamelike interface rewards engagement, delivering a dopamine boost when users accrue likes and responses, training users to indulge behaviors that win affirmation. And because its algorithm unintentionally privileges negativity, the greatest rush comes by attacking outsiders: The other sports team. The other political party. The ethnic minority. ... by supercharging content that taps into tribal identity, [Facebook] can upset fragile communal balances.⁷⁴

In addition to the hit of dopamine Taub and Fischer describe, Amy Chua, an expert on tribalism and its social impacts, explains that in-group instincts that demonize others also raise our oxytocin levels, meaning we are influenced by not one but two of the most powerful neurological motivators in our body’s arsenal. The combination, Chua says, “physiologically ‘anesthetizes’ the empathy one might otherwise feel” towards others,⁷⁵ pushing us to unconsciously act differently and more cruelly than we normally would towards others.

Add to the problems of misinformation, disinformation, propaganda, denialism, and tribalism, the rise of troll farms, automated and semi-automated accounts (“bots”), and the ability of any group to target any message to anyone in the world, and you have a recipe for a veritable political catastrophe, the effects of which we began to see in 2016.

2016: U.S. Presidential Election & Brexit

In addition to the divisive misinformation campaigns around the world in countries that rely heavily on Facebook for news and communication, two more internet bombs of extremism and misinformation dropped in 2016: the U.S. presidential election and Britain's decision to leave the European Union. Within the Facebook/Instagram, Google/YouTube, and Twitter ecosystems, the ability to advertise products and messages to target audiences took a dark turn in the years leading up to both 2016 elections. Foreign and domestic agents utilized the micro-targeting technology available on each platform in order to influence and disrupt the political systems of the U.S. and U.K., encouraging "Leave" votes in the U.K. and the election of Donald Trump across the pond. Central to the success of both campaigns was Cambridge Analytica, a data analytics company responsible for hoarding and hijacking information extracted from Facebook's unsuspecting user base.

Way back in 2008, two Cambridge University researchers, Michal Kosinski and David Stillman, discovered that online behaviors and psychometric data were incredibly useful in predicting users' personality, traits, and demographics, such as race, sexual orientation, political affiliation, intelligence, substance use, and even if they were the children of divorced parents.⁷⁶

A few years later, in 2011, the real fun started. Another Cambridge researcher named Aleksandr Kogan collaborated with Facebook on a study published in the journal *Personality and Individual Differences* on friendships.⁷⁷ The data for the study was supplied to Kogan by Facebook, which included information on 57 billion Facebook friendships. The same year, Facebook began offering a feature on its platform called "friends permissions," which allowed third-party developers to collect masses of personal information about users and their friends (without their friends' permission). During this time, while approximately 9 million apps were integrated with Facebook, a vast quantity of user data was extracted and harvested by various companies using the feature. Kogan has since estimated that tens of thousands of developers extracted data in the same way he did, and that Facebook was very much aware of this, saying the company considered it "a feature, not a bug."⁷⁸

In 2013, Cambridge Analytica was founded by Christopher Wylie and Alexander Nix as a subsidiary of SCL Group, which described itself as a strategic communications company focusing on “global election management.” Using sophisticated data mining and analysis techniques, SCL focused primarily on advising governments and military organizations on behavioral change programs. Later that year, Nix and Wylie demoed Cambridge Analytica’s capabilities to billionaire Robert Mercer, a Trump supporter, who put up an initial \$15 million in funding; Steve Bannon, who would later become Trump’s chief strategist, invested an estimated \$1 to 5 million in the company.

The following year, Kogan and his colleague, Joseph Chancellor, founded a company called Global Science Research (GSR) and signed a contract with Cambridge Analytica to create an app that would harvest users’ psychometric Facebook data. Kogan and Chancellor extracted data from 270,000 Facebook users and their friends—including their status updates, likes, and private messages—which amounted to a data set of more than 87 million people. Cambridge Analytica then used Kogan and Chancellor’s data to create over 30 million user profiles, identify target voter groups, and design specific targeted messaging to influence voters’ opinions and behaviors.

In June 2016, the Trump campaign hired Cambridge Analytica for \$6 million and began to use Facebook as both its primary fundraising and propaganda vehicle.

The campaign uploaded its voter files—the names, addresses, voting history, and any other information it had on potential voters—to Facebook. Then, using a tool called Lookalike Audiences, Facebook identified the broad characteristics of, say, people who had signed up for Trump newsletters or bought Trump hats. That allowed the campaign to send ads to people with similar traits. Trump would post simple messages like “This election is being rigged by the media pushing false and unsubstantiated charges, and outright lies, in order to elect Crooked Hillary!” that got hundreds of thousands of likes, comments, and shares. The money rolled in. Clinton’s wonkier messages, meanwhile, resonated less on the platform. Inside Facebook, almost everyone on the executive team wanted Clinton to win; but they knew that Trump was using the platform better. If he was the candidate for Facebook, she was the candidate for LinkedIn.⁷⁹

According to Bloomberg reporters Joshua Green and Sasha Issenberg, in addition to spreading false and inflammatory information, Cambridge Analytica data was used to encourage voter suppression: “the Trump campaign used so-called dark posts—nonpublic posts targeted at a specific audience—to discourage African Americans from voting in battleground states.”⁸⁰ As Tufekci points out, however, Trump’s campaign “wasn’t deviantly weaponizing an innocent tool. It was simply using Facebook exactly as it was designed to be used.”

The campaign did it cheaply, with Facebook staffers assisting right there in the office, as the tech company does for most large advertisers and political campaigns. Who cares where the speech comes from or what it does, as long as people see the ads? The rest is not Facebook’s department.⁸¹

After Trump’s victory in November 2016, Zuckerberg described the idea that his platform may have been used to influence the results of the Presidential election as “a pretty crazy idea.” Sixteen months later, Zuckerberg was finally convinced. Facebook suspended SCL and Cambridge Analytica, as well as Wylie and Kogan, from the platform. Chancellor, by contrast, has been gainfully employed by Facebook since 2015. This was followed by a series of events: the suspension of CEO Nix from Cambridge Analytica, Cambridge Analytica’s bankruptcy, the suspension of 200 apps from Facebook’s platform, the plunging of Facebook’s stock by 24 percent, and the company’s admission that activities on its service indicated “coordinated inauthentic behavior” from the Russian, Kremlin-linked group, the Internet Research Agency. Cambridge Analytica’s parent company, SCL, on the other hand, continues to capitalize on the data obtained from Facebook. In early 2017, armed with the psychometric information of 230 million U.S. citizens, SCL “had won contracts with the US State Department and was pitching to the Pentagon.”⁸²

Around the same time, British citizens were dealing with their own election nightmare. In a session before the Digital, Culture, Media and Sport Committee, Wylie explained to MPs how the EU referendum vote “was won through fraud” via the Vote Leave campaign “improperly channelling money through a tech firm with links to Cambridge Analytica.”⁸³

Wylie said it was striking that Vote Leave and three other pro-Brexit groups—BeLeave, which targeted students; Veterans for Britain, and Northern Ireland’s Democratic Unionist party—all used the services of the little-known firm Aggregate IQ (AIQ) to help target voters online. He told MPs that AIQ was effectively the Canadian arm of Cambridge Analytica/SCL, deriving the majority of its income by acting as a sub-contractor.⁸⁴

The shuffling of funds between the Vote Leave and BeLeave campaigns, which was spent on AIQ’s services, is currently under investigation by the UK Electoral Commission. The EU referendum won by a small margin (2%) of the vote, a result that Wylie believes would have been very different had it not been for AIQ’s involvement, combined with the possible violation of campaign spending limits.

Wylie, who was just 24-years-old when he helped Nix form Cambridge Analytica, now describes the company as a “full service propaganda machine.”⁸⁵ He told Carole Cadwalladr, the Guardian reporter who broke the scandal, he believed the methods employed by Cambridge Analytica and the campaigns that had hired them were “worse than bullying. Because people don’t necessarily know it’s being done to them.”

At least bullying respects the agency of people because they know. . . . if you do not respect the agency of people, anything that you’re doing after that point is not conducive to a democracy. And fundamentally, information warfare is not conducive to democracy.⁸⁶

The importance of sound, accurate information is essential to the institution of democracy. United States District Judge Amy Berman Jackson, who sentenced Trump campaign manager Paul Manafort in 2019 for multiple charges, including tax fraud and conspiracy, stated at the sentencing, if “people don’t have the facts, democracy can’t work.” In an article for *Boston Review*, Clara Hendrickson specifically calls Facebook and Instagram’s priorities antithetical to democracy, noting that its policies have “proved to fragment, polarize, and threaten liberal democracy.”⁸⁷

According to the University of Gothenburg's 2018 annual Democracy Report, democracy began declining in 2006 and 2007 across a number of regions, including Latin America, the Caribbean, Eastern Europe, Central America, the Middle East, North Africa, Western Europe and North America.⁸⁸ These years are, coincidentally, considered to be seminal in tech; 2006 was the year Twitter was launched, Facebook was released to the public, and Google acquired YouTube. In June of the following year, the iPhone made its debut. The only two regions whose democracies the report found to be improving rather than regressing were sub-Saharan Africa and Asia, which, according to the report analysis, are the only two regions whose internet penetration rates fall below the world average (in other words, these two regions don't use the internet as much as those whose democracies are in decline).⁸⁹

In addition to domestic manipulation of U.S. and U.K. elections, foreign influence on social media has also played an important role, particularly Russian interference. In one of its 2018 Congressional hearings, Facebook admitted that 170 Instagram accounts and 120 Facebook pages “were found to have spread propaganda from Russia’s Internet Research Agency.”⁹⁰ If 120 doesn’t sound too bad, consider data journalist and research director at Columbia University’s Tow Center for Digital Journalism Jonathan Albright’s findings: posts from only six of the Russian accounts suspended by Facebook had been shared a whopping 340 million times.⁹¹ Investigations into Trump’s victory, Brexit, and Russian interference are currently well under way at the time of this writing and, once complete, will likely illustrate one of the most comprehensive, devastating, and unthinkable attacks on democracy in modern history.

Russia, Facebook, Trump, Mercer, Bannon, Brexit. Every one of these threads runs through Cambridge Analytica. Even in the past few weeks, it seems as if the understanding of Facebook’s role has broadened and deepened. The Mueller indictments were part of that, but Paul-Olivier Dehaye—a data expert and academic based in Switzerland, who published some of the first research into Cambridge Analytica’s processes—says it’s become increasingly apparent that Facebook is “abusive by design”. If there is evidence of collusion between the Trump campaign and Russia, it will be in the platform’s data flows.⁹²

The mechanism by which “Facebook was hijacked, repurposed to become a theatre of war,” and “how it became a launchpad for what seems to be an extraordinary attack on the U.S.’s democratic process,”⁹³ is both scary and complex. The motivation that allowed it to persist, however, is much more straightforward: revenue generated by advertising.

Notes

1. Brin, S., & Page, L. (1998). The Anatomy of a Large-Scale Hypertextual Web Search Engine. *Computer Networks and ISDN Systems*, 30(1), 107–117. [https://doi.org/10.1016/S0169-7552\(98\)00110-X](https://doi.org/10.1016/S0169-7552(98)00110-X)
2. Schmidt, E. (2016, June 9). *Interview with Eric Schmidt* (C. Rose, Interviewer). Retrieved from <https://charlirose.com/videos/28222>
3. Cohen, N. (2017, October 13). Silicon Valley is Not Your Friend. *The New York Times*. Retrieved from <https://www.nytimes.com/interactive/2017/10/13/opinion/sunday/Silicon-Valley-Is-Not-Your-Friend.html>
4. Rosenberg, E. (2015, February 5). How Google Makes Money. Retrieved August 27, 2018, from *Investopedia* website: <https://www.investopedia.com/articles/investing/020515/business-google.asp>
5. Google: Ad Revenue 2001–2018. (n.d.). Retrieved May 27, 2019, from *Statista* website: <https://www.statista.com/statistics/266249/advertising-revenue-of-google/>
6. Tufekci, Z. (2018, January 16). It’s the (Democracy-Poisoning) Golden Age of Free Speech. *Wired*. Retrieved from <https://www.wired.com/story/free-speech-issue-tech-turmoil-new-censorship/>
7. Tarnoff, B. (2017, August 23). Silicon Valley Siphons Our Data like Oil. But the Deepest Drilling has Just Begun. *The Guardian*. Retrieved from <https://www.theguardian.com/world/2017/aug/23/silicon-valley-big-data-extraction-amazon-whole-foods-facebook>
8. Zuckerberg, M. (2018, July 18). *Zuckerberg: The Recode Interview* (K. Swisher, Interviewer). Retrieved from <https://www.recode.net/2018/7/18/17575156/mark-zuckerberg-interview-facebook-recode-kara-swisher>
9. Facebook’s Gollum Will Never Give Up Its Data Ring. (2018, May 5). Retrieved September 9, 2018, from *Newco Shift* website: <https://shift.newco.co/2018/05/05/facebooks-gollum-will-never-give-up-its-data-ring/>

10. Battelle, J. (2017, September 15). Lost Context: How Did We End Up Here? Retrieved September 2, 2018, from *Newco Shift* website: <https://shift.newco.co/2017/09/15/lost-context-how-did-we-end-up-here/>
11. Carr, C. (2010, August 8). The Meaning of Mad Men: Philosophers Take on TV. *Time*. Retrieved from <https://content.time.com/time/arts/article/0,8599,2009261,00.html>
12. Chollet, F. (2018, March 28). What Worries Me about AI. Retrieved September 7, 2018, from *François Chollet* website: <https://medium.com/@francois.chollet/what-worries-me-about-ai-ed9df072b704>
13. Lanier, J. (2018, April 17). *We Won, and We Turned into Assholes* (N. Kulwin, Interviewer) [*New York Magazine*]. Retrieved from <http://nymag.com/selectall/2018/04/jaron-lanier-interview-on-what-went-wrong-with-the-internet.html>
14. Haque, U. (2017, September 15). Is Social Media a Failure? Retrieved August 27, 2018, from *Eudaimonia and Co* website: <https://eand.co/is-social-media-a-failure-f4f970695d17>
15. Kulwin, N. (2018, April 13). An Apology for the Internet—From the Architects Who Built It. *Select All*. Retrieved from <http://nymag.com/intelligencer/2018/04/an-apology-for-the-internet-from-the-people-who-built-it.html>
16. Rose-Stockwell, T. (2017). This is How Your Fear and Outrage are Being Sold for Profit. *Medium*. Retrieved February 18, 2019, from <https://medium.com/@tobiasrose/the-enemy-in-our-feeds-e86511488de>
17. Shearer, E., & Gottfried, J. (2017). *News Use across Social Media Platforms 2017*. Retrieved from Pew Research Center website: <https://www.journalism.org/2017/09/07/news-use-across-social-media-platforms-2017/>
18. Rose-Stockwell, T. (2017). This is How Your Fear and Outrage are Being Sold for Profit. *Medium*. Retrieved February 18, 2019, from <https://medium.com/@tobiasrose/the-enemy-in-our-feeds-e86511488de>
19. Wineburg, S., McGrew, S. Breakstone, J., & Ortega, T. (2016). Evaluating Information: The Cornerstone of Civic Online Reasoning. *Stanford Digital Repository*. Retrieved September 2, 2018, from <https://purl.stanford.edu/fv751yt5934>
20. Vosoughi, S., Roy, D., & Aral, S. (2018). The Spread of True and False News Online. *Science*, 359(6380), 1146–1151. <https://doi.org/10.1126/science.aap9559>
21. Meyer, R. (2018, March 8). The Grim Conclusions of the Largest-Ever Study of Fake News. *The Atlantic*. Retrieved from <https://www.theatlantic.com/technology/archive/2018/03/largest-study-ever-fake-news-mit-twitter/555104/>

22. Johns Hopkins Sheridan Libraries. (2018, June 25). Evaluating Information: Propaganda vs. Misinformation. Retrieved August 22, 2018, from <http://guides.library.jhu.edu/evaluate/propaganda-vs-misinformation>
23. Ibid.
24. Ibid.
25. Taub, A., & Fisher, M. (2018, April 21). Where Countries are Tinderboxes and Facebook is a Match. *The New York Times*. Retrieved from <https://www.nytimes.com/2018/04/21/world/asia/facebook-sri-lanka-riots.html>
26. Facebook Chief Fires Back at Apple Boss. (2018, April 2). *BBC News*. Retrieved from <https://www.bbc.com/news/technology-43619410>
27. Human Rights Council. (2018). *Report of the Independent International Fact-finding Mission on Myanmar* (No. A/HRC/39/64; pp. 1–21). Retrieved from https://www.ohchr.org/EN/HRBodies/HRC/RegularSessions/Session39/_layouts/15/WopiFrame.aspx?sourcedoc=/EN/HRBodies/HRC/RegularSessions/Session39/Documents/A_HRC_39_64.docx&action=default&DefaultItemOpen=1
28. Larson, C. (2017, November 7). Facebook Can't Cope with the World It's Created. *Foreign Policy*. Retrieved from <https://foreignpolicy.com/2017/11/07/facebook-cant-cope-with-the-world-its-created/>
29. Taub, A., & Fisher, M. (2018, April 21). Where Countries are Tinderboxes and Facebook is a Match. *The New York Times*. Retrieved from <https://www.nytimes.com/2018/04/21/world/asia/facebook-sri-lanka-riots.html>
30. Ibid.
31. Ibid.
32. Ibid.
33. Harris, J. (2018, May 6). In Sri Lanka, Facebook's Dominance has Cost Lives. *The Guardian*. Retrieved from <https://www.theguardian.com/commentisfree/2018/may/06/sri-lanka-facebook-lives-tech-giant-poor-countries>
34. Hern, A. (2018, July 20). WhatsApp to Restrict Message Forwarding after India Mob Lynchings. *The Guardian*. Retrieved from <https://www.theguardian.com/technology/2018/jul/20/whatsapp-to-limit-message-forwarding-after-india-mob-lynchings>
35. Goel, V. (2018, May 16). In India, Facebook's WhatsApp Plays Central Role in Elections. *The New York Times*. Retrieved from <https://www.nytimes.com/2018/05/14/technology/whatsapp-india-elections.html>

36. Waterson, J. (2018, June 17). Fears Mount Over WhatsApp's Role in Spreading Fake News. *The Guardian*. Retrieved from <https://www.theguardian.com/technology/2018/jun/17/fears-mount-over-whatsapp-role-in-spreading-fake-news>
37. Goel, V. (2018, May 16). In India, Facebook's WhatsApp Plays Central Role in Elections. *The New York Times*. Retrieved from <https://www.nytimes.com/2018/05/14/technology/whatsapp-india-elections.html>
38. Waterson, J. (2018, June 17). Fears Mount over WhatsApp's Role in Spreading Fake News. *The Guardian*. Retrieved from <https://www.theguardian.com/technology/2018/jun/17/fears-mount-over-whatsapp-role-in-spreading-fake-news>
39. Etter, L. (2017, December 7). Rodrigo Duterte Turned Facebook into a Weapon, with a Little Help from Facebook. *Bloomberg.Com*. Retrieved from <https://www.bloomberg.com/news/features/2017-12-07/how-rodrigo-duterte-turned-facebook-into-a-weapon-with-a-little-help-from-facebook>
40. Ibid.
41. Ibid.
42. YouTube to Invest \$25 Million to Boost "Trusted" News Sources (2018, July 10). *ABS-CBN News*. Retrieved from <https://news.abs-cbn.com/business/07/10/18/youtube-to-invest-25-million-to-boost-trusted-news-sources>
43. Solon, O. (2018, March 24). 'A Grand Illusion': Seven Days that Shattered Facebook's Facade. *The Guardian*. Retrieved from <https://www.theguardian.com/technology/2018/mar/24/cambridge-analytica-week-that-shattered-facebook-privacy>
44. Zuckerberg, M. (2018, July 18). *Zuckerberg: The Recode Interview* (K. Swisher, Interviewer). Retrieved from <https://www.recode.net/2018/7/18/17575156/mark-zuckerberg-interview-facebook-recode-kara-swisher>
45. Ibid.
46. Ibid.
47. Feinberg, J., *The Moral Limits of the Criminal Law, Volume 1: Harm to Others* (New York: Oxford University Press, 1984), p. 37.
48. Bartlett, J. (2018). *The People vs Tech: How the Internet is Killing Democracy (and How We Save It)*. New York: Penguin.
49. Definition of denialist. (n.d.). In *Oxford Dictionaries*. Retrieved from <https://en.oxforddictionaries.com/definition/denialist>

50. Kahn-Harris, K. (2018, August 3). Denialism: What Drives People to Reject the Truth. *The Guardian*. Retrieved from <https://www.theguardian.com/news/2018/aug/03/denialism-what-drives-people-to-reject-the-truth>
51. Ibid.
52. Ibid.
53. NASA. (2018). *Climate Change Causes: A Blanket Around the Earth* [Vital Signs of the Planet]. Retrieved from <https://climate.nasa.gov/causes>
54. Díaz, S., Settele, J., & Brondízio, E. (2019). *IPBES Global Assessment Summary for Policymakers*. Retrieved from IPBES website: <https://www.ipbes.net/news/ipbes-global-assessment-summary-policymakers-pdf>
55. Kahn-Harris, K. (2018, August 3). Denialism: What Drives People to Reject the Truth. *The Guardian*. Retrieved from <https://www.theguardian.com/news/2018/aug/03/denialism-what-drives-people-to-reject-the-truth>
56. Measles and Rubella Surveillance Data. (2018, August 13). Retrieved August 22, 2018, from World Health Organization website: http://www.who.int/immunization/monitoring_surveillance/burden/vpd/surveillance_type/active/measles_monthlydata/en/
57. Stokes-Lampard, H. (2018, August 21). Anti-vaxxers are Still Spreading False Claims as People Die of Measles. *The Guardian*. Retrieved from <https://www.theguardian.com/commentisfree/2018/aug/21/anti-vaxxers-measles-mmr-vaccine-gp-online>
58. Ibid.
59. Ibid.
60. Snyder, T. (2018, May 21). Fascism is Back. Blame the Internet. *Washington Post*. Retrieved from <https://www.washingtonpost.com/news/posteverything/wp/2018/05/21/fascism-is-back-blame-the-internet/>
61. Qiu, X., Oliveira, D. F. M., Shirazi, A. S., Flammini, A., & Menczer, F. (2017). Limited Individual Attention and Online Virality of Low-quality Information. *Nature Human Behaviour*, 1(7), 0132. <https://doi.org/10.1038/s41562-017-0132>
62. Demby, G. (2016, July 9). *How Social Media Impacts the Conversation on Racial Violence* (L. Neary, Interviewer) [NPR]. Retrieved from <https://www.npr.org/2016/07/09/485356145/how-social-media-impacts-the-conversation-on-racial-violence>

63. Naughton, J. (2018, March 18). Extremism Pays. That's Why Silicon Valley Isn't Shutting It Down. *The Guardian*. Retrieved from <https://www.theguardian.com/commentisfree/2018/mar/18/extremism-pays-why-silicon-valley-not-shutting-it-down-youtube>
64. Steele, C. (2018, November 28). Extremism Pays for YouTube, but at What Cost? Retrieved January 16, 2019, from *PC Magazine* website: <https://www.pcmag.com/news/365140/extremism-pays-for-youtube-but-at-what-cost>
65. Friedman, T. L. (2016). *Thank You for Being Late: An Optimist's Guide to Thriving in the Age of Accelerations* (p. 274). New York: Farrar, Straus and Giroux.
66. Crockett, M. J. (2017). Moral Outrage in the Digital Age. *Nature Human Behaviour*, 1(11), 769–771. <https://doi.org/10.1038/s41562-017-0213-3>
67. Crockett also notes that expressing our moral outrage online may make us less likely to act in the real world in a meaningful and socially engaged way, such as by donating or volunteering our time, as digital platforms “the way we experience outrage, and limiting how much we can actually change social realities.”
68. Crockett, M. J. (2017). Moral Outrage in the Digital Age. *Nature Human Behaviour*, 1(11), 769–771. <https://doi.org/10.1038/s41562-017-0213-3>
69. Ibid.
70. Qiu, X., Oliveira, D. F. M., Shirazi, A. S., Flammini, A., & Menczer, F. (2017). Limited Individual Attention and Online Virality of Low-quality Information. *Nature Human Behaviour*, 1(7), 0132. <https://doi.org/10.1038/s41562-017-0132>
71. Bazalgette, P. (2017). *The Empathy Instinct: How to Create a More Civil Society* (p. 221). London: John Murray.
72. Rosling, H., Rosling, O., & Rosling Rönnlund, A. (2018). *Factfulness: Ten Reasons We're Wrong about the World—And Why Things are Better Than You Think* (p. 146). London: Sceptre.
73. Bazalgette, P. (2017). *The Empathy Instinct: How to Create a More Civil Society* (p. 117). London: John Murray.
74. Taub, A., & Fisher, M. (2018, April 21). Where Countries are Tinderboxes and Facebook is a Match. *The New York Times*. Retrieved from <https://www.nytimes.com/2018/04/21/world/asia/facebook-sri-lanka-riots.html>

75. Chua, A. (2018, June 14). Tribal World. *Foreign Affairs*. Retrieved July/August 2018, from <https://www.foreignaffairs.com/articles/world/2018-06-14/tribal-world>
76. The researchers would go on to publish two papers on the subject: “Private traits and attributes are predictable from digital records of human behavior” in 2013, and “Psychological targeting as an effective approach to digital mass persuasion” in 2017. They would also eventually be approached by Christopher Wylie to discuss using their data.
77. Yearwood, M. H., Cuddy, A., Lamba, N., Youyou, W., van der Lowe, I., Piff, P. K., ... Spectre, A. (2015). On Wealth and the Diversity of Friendships: High Social Class People Around the World have Fewer International Friends. *Personality and Individual Differences*, 87, 224–229. <https://doi.org/10.1016/j.paid.2015.07.040>
78. Stahl, L. (2018, April 22). *Aleksandr Kogan: The Link between Cambridge Analytica and Facebook*. Retrieved from <https://www.cbsnews.com/news/aleksandr-kogan-the-link-between-cambridge-analytica-and-facebook/>
79. Vogelstein, F., & Thompson, N. (2018, February 12). Inside Facebook’s Two Years of Hell. *Wired*. Retrieved from <https://www.wired.com/story/inside-facebook-mark-zuckerberg-2-years-of-hell/>
80. Tufekci, Z. (2018, January 16). It’s the (Democracy-Poisoning) Golden Age of Free Speech. *Wired*. Retrieved from <https://www.wired.com/story/free-speech-issue-tech-turmoil-new-censorship/>
81. Ibid.
82. Cadwalladr, C. (2018, March 18). ‘I Made Steve Bannon’s Psychological Warfare Tool’: Meet the Data War Whistleblower. *The Guardian*. Retrieved from <https://www.theguardian.com/news/2018/mar/17/data-war-whistleblower-christopher-wylie-faceook-nix-bannon-trump>
83. Hern, A., & Sabbagh, D. (2018, March 27). EU Referendum Won Through Fraud, Whistleblower Tells MPs. *The Guardian*. Retrieved from <https://www.theguardian.com/uk-news/2018/mar/27/brexit-groups-had-common-plan-to-avoid-election-spending-laws-says-wylie>
84. Ibid.
85. Memoli, M., & Schecter, A. (2018, April 25). Bannon Turned Cambridge into “Propaganda Machine,” Whistleblower Says. Retrieved January 17, 2019, from *NBC News* website: <https://www.nbcnews.com/politics/donald-trump/bannon-turned-cambridge-propaganda-machine-whistleblower-says-n869126>

86. Cadwalladr, C. (2018, March 18). 'I Made Steve Bannon's Psychological Warfare Tool': Meet the Data War Whistleblower. *The Guardian*. Retrieved from <https://www.theguardian.com/news/2018/mar/17/data-war-whistleblower-christopher-wylie-faceook-nix-bannon-trump>
87. Hendrickson, C. (2018, April 12). Democracy vs. the Algorithm. Retrieved September 3, 2018, from *Boston Review* website: <http://bostonreview.net/science-nature/clara-hendrickson-democracy-vs-algorithm>
88. University of Gothenburg Varieties of Democracy Institute. (2018). *V-Dem Annual Democracy Report 2018* (pp. 1–96). Retrieved from https://www.v-dem.net/media/filer_public/3f/19/3f19efc9-e25f-4356-b159-b5c0ec894115/v-dem_democracy_report_2018.pdf
89. World Internet Users Statistics and 2019 World Population Stats. (2019, March 31). Retrieved May 26, 2019, from *Internet World Stats* website: <https://www.internetworldstats.com/stats.htm>
90. Volpicelli, G. (2018, April 28). Can Instagram Keep Its Nose Clean? *The Observer*. Retrieved from <https://www.theguardian.com/technology/2018/apr/28/instagram-at-the-crossroads-profits-facebook-data-scandal-politics-influencers-mental-health>
91. Vogelstein, F., & Thompson, N. (2018, February 12). Inside Facebook's Two Years of Hell. *Wired*. Retrieved from <https://www.wired.com/story/inside-facebook-mark-zuckerberg-2-years-of-hell/>
92. Cadwalladr, C. (2018, March 18). 'I Made Steve Bannon's Psychological Warfare Tool': Meet the Data War Whistleblower. *The Guardian*. Retrieved from <https://www.theguardian.com/news/2018/mar/17/data-war-whistleblower-christopher-wylie-faceook-nix-bannon-trump>
93. Ibid.

Open Access This chapter is licensed under the terms of the Creative Commons Attribution 4.0 International License (<http://creativecommons.org/licenses/by/4.0/>), which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence and indicate if changes were made.

The images or other third party material in this chapter are included in the chapter's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the chapter's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder.

