

Twitter and Tear Gas

THE POWER AND FRAGILITY
OF NETWORKED PROTEST

Yale

UNIVERSITY PRESS

NEW HAVEN & LONDON

Platforms and Algorithms

I TRAVELED TO CAIRO IN THE SPRING OF 2011, a few months after the fall of President Hosni Mubarak. Egypt was unsettled but jubilant. One of the Egyptians I interviewed was a blogging pioneer whom I will call “Hani.”¹ In the early years of the twenty-first century, Hani had been among the first to take advantage of the internet’s revolutionary potential. Most Egyptian bloggers made it through the Mubarak era unscathed because the government could not keep up with or fully understand the new medium. Unfortunately, the government noticed Hani; he was tried and sentenced to years in prison for the crime of insulting Mubarak. At the time, there was little open dissent in Egypt. The public sphere was dominated by mass-media outlets controlled by the government, and Egyptians were in the early stages of experimenting with the use of the internet for sharing political information.² When he was released in November 2010 after six years in prison, Hani was still defiant. Before his prison term, Hani’s blog had been a bustling crossroads of discussion, with his voice reaching farther than he had ever thought possible. After his involuntary hiatus, Hani told me that he had resumed blogging with enthusiasm, but he found that his blog, which had formerly been abuzz with spirited conversations, as well as the rest of the Egyptian blogosphere, seemed deserted. “Where is everybody?” Hani asked me before answering himself, “They’re on Facebook.”

A few years later, I heard a very similar story from Hossein Derakshan, an Iranian blogger, who had become the primary actor in a similarly unfortu-

nate “natural” experiment. Before 2008, he operated a lively blog in Farsi with a large readership in Iran, gaining a reputation as Iran’s “blogfather.” Tragically, he was put in jail in 2008 for six years, missing the whole shift to Facebook. When he was finally released, in 2014, he started enthusiastically blogging again—to crickets. There was no response or readership. Assuming that he just had to keep blogging via Facebook, he took it up and wholeheartedly put his material there. Hossein told me the story in 2016: how his Facebook posts just disappeared into the site, his weighty subjects unable to garner the cheery “Likes” that are a key currency of the algorithm that runs on the platform. The web is all turning into a form of television, he sighed and pondered if, at this rate, the powers-that-be may not even have to censor it in Iran. Facebook’s algorithmic environment would bury them, anyway.³

For many of the Egyptian activists I talked with, especially in the early days of the revolution, Facebook’s ability to reach so many Egyptians felt empowering. Ordinary people who otherwise might not have taken to the internet were joining the site for social reasons: to keep in touch with family and friends. For many Egyptians, joining Facebook was the entry to becoming connected to their family and friends, but it also meant joining the networked public sphere. Exposure to the ideas and information circulated by political activists was a side effect of their Facebook membership. A study based on a survey of Tahrir Square protesters—that I co-authored—confirms that social media platforms like Facebook and Twitter drove the crucial early turnout of protesters in Tahrir Square that triggered the avalanche of dissent.⁴ More than a quarter of the protesters surveyed had first heard about the protests on Facebook, and Twitter users significantly more likely to were among the initial group that showed up in Tahrir Square on the first day of the protests. Overall, the study found that social media had played a crucial role.

During January and February, many Egyptians were riveted by the power struggle being played out between the Tahrir Square protesters and the country’s leadership, who had heretofore seemed invincible. Mubarak’s government did not grasp the power that the ability to document, communicate, and coordinate via social media placed in the hands of ordinary people. By the time Mubarak was forced to resign, Facebook had become a major player in the civic sphere, and its use continued to grow after the

initial uprising. Even the new military council that replaced Mubarak launched a Facebook page. But what did it mean that Facebook had become so central to the political life of the country? This was unclear at the time.

With the advent of social media platforms around 2005, the burgeoning civic space developing online, mostly through blogs, expanded greatly. In the same time period though, it also underwent a major transformation, shifting from individual blogs and web pages to massive, centralized platforms where visibility was often determined by an algorithm controlled by the corporation, often with the business model seeking to increase page-views.⁵ In many places, including the United States, the Middle East, Russia, Turkey, and Europe, the networked public sphere largely shifted to commercial spaces. The platforms were chiefly Facebook, Twitter, and YouTube, along with a few others that facilitated sharing content.⁶ Some countries had no prior infrastructure to build upon, or to transition away from. For example, Myanmar, just emerging from a military dictatorship under which there had been no active public sphere in the traditional sense, plunged straight into the networked public sphere.⁷

As these changes occurred, scholars and civic activists worried about how these new “sovereigns of cyberspace,” platforms like Facebook and Twitter, would wield their power.⁸ Would they censor and restrict freedoms to serve the interests of advertisers or governments? Would they turn over user information to repressive regimes? Internet-freedom advocate Rebecca MacKinnon was prescient in identifying the core problem: the growth of privately owned spaces that functioned as a new kind of public space, as if street corners or cafés where people gathered were owned by a few corporations.⁹

During the 1950s, when U.S. television networks showed images of the brutal acts of police encountered by civil rights protesters, their often belated editorial decisions to bring these issues to the attention of the American public opened possibilities for activists and ultimately helped shape the trajectory of the movement. During the next decade, when civil rights protesters were planning future actions, reaching network news audiences became one of their key strategic goals. Activists knew that television coverage (or the lack of it) could potentially make or break a movement.

Nowadays, the function of gatekeeping for access to the public sphere is enacted through internet platforms’ policies, algorithms, and affordances.

In some ways, this has empowered movements by reducing their dependency on traditional mass media and their editors. In other respects, the current digital communications gatekeeping ecosystem has been reduced to a very few but very powerful choke points. Social movements today are largely dependent on a very small number of corporate platforms and search engines (or, more accurately, one search engine, Google).

While billions of people use the internet, a small number of services capture or shape most of their activities. Facebook has 1.5 billion users, 1 billion of whom log in daily to see updates and news from the hundreds of people they have “friended” on the platform.¹⁰ Google processes more than three billion searches every day. The dominance of a few platforms online is not a historical coincidence; rather, it is the product of two important *structural* dynamics: network effects¹¹ and the dominance of the ad-financing model for online platforms.

The term “network effects” (or “network externalities”) is a shorthand for the principle that the more people who use a platform, the more useful that platform is to each user.¹² Such effects are especially strong for online social networking platforms since the main point is to access other users and the content they have posted. Think of a telephone that could talk only to telephones made by the same company: what good is a wonderful telephone if you cannot call anyone with it? You would want to get the one most of your friends used even if you liked another company’s model better. When network effects operate, potential alternatives are less useful simply because fewer people use them. Thus a platform that achieves early success can become dominant as more and more people flock to it. Network effects limit competition and thus the ability of the market to impose constraints on a dominant platform. This advantage is operative for Facebook (where most people know that their friends and family will have accounts) and Google (users provide it with data and resources to make its search better, and advertisers pay to advertise on Google knowing that it is where people will search, hence Google has even more money available to improve its products). This is true even for nonsocial platforms like eBay (where buyers know that the largest number of sellers are offering items, and sellers know that the largest number of buyers will see their items).

It is true that network effects did not provide absolute protection early in the race to commercialize the internet: MySpace was beaten out by Facebook, for example, and Yahoo and Altavista by Google—they had gotten started earlier, but had not yet established in as dominant a position. Network effects doesn't protect companies from initial missteps, especially in the early years before they pulled way ahead of everyone else, and such dominance does not occur independent of the quality of the company's product. Google's new method of ranking web pages was clearly superior to the earlier competitors. Network effects may not mean that the very first companies to enter a new and rapidly growing market and achieve sizable growth will necessarily be the ones to emerge as dominant once the market has matured and growth has slowed. But at that point, whichever companies are dominant will be very difficult for competitors to unseat. Network effects are certainly apparent in the dynamics we see currently in the use of, for example, Facebook, Google, and eBay. Beyond network effects, the costs of entry into these markets have also become high because of the data these companies have amassed. A competitor to these behemoths would need to be massively financed and would still be at a huge disadvantage given the enormous amount of data about users' habits these companies have already collected.

Another key dynamic operating in this commercial, quasi-public networked sphere dominated by a few companies is that most platforms that activists use, the places where user-generated content is produced and shared, are financed by advertising.¹³ Ads on the internet are not worth as much to advertisers as print ads in traditional media because they are easily ignored in an online environment and because there are so many of them. This means that immense scale is paramount for the financial viability of an online platform. Platforms must be huge, or they will find themselves in financial trouble. Even Twitter, with hundreds of millions of users, is considered too small to be viable by Wall Street. That each internet ad is worth so little encourages corporations to surveil users' habits, actions, and interests. The only way for platforms to increase the price they are paid for ads is to create tailored ads that target particular users who are likely to buy specific products. The vast amounts of data that platforms collect about users are what allow this tailoring to be performed.

These pressures to achieve huge scale and to minutely monitor users promote the centralization and surveillance tendency of platforms like Facebook and Google and their interests in monopolizing both ad dollars and users. The enormous platforms in turn become even better resourced hubs of activity. These structural factors combine in a runaway dynamic that smothers smaller platforms: the huge platforms are the only ones that have enough surveillance data to profile their users so that the ads they display are worth something, which in turn means that they have even more resources and data on users as more and more people join them because that is where most of their friends are.¹⁴

Because of this spiral of network effects and ad financing, for an increasing number of people, Facebook and Google are the internet, or at least the framework that shapes their experience of it.¹⁵ For social movements, Facebook is the indispensable platform along with a very few others, like Twitter and Tumblr (owned by Yahoo), and Google is the ne plus ultra of search engines. The picture-sharing site Instagram and the messaging service WhatsApp, which are also important, have already been acquired by Facebook. These platforms own the most valuable troves of user data, control the user experience, and wield the power to decide winners and losers for people's attention by making small changes to their policies and algorithms in a variety of categories, including news, products, and books. These platforms also offer users other strengths and real benefits. For example, like Google provides better security against state snooping (except that of the U.S. government), and Facebook's WhatsApp is encrypted end-to-end, making it more secure than all the poorly financed alternatives while still being widely available and easy to use (a major issue plaguing niche platforms that cater to activists).

Communicating primarily in this networked public but privately owned sphere is a bit like moving political gatherings to shopping malls from public squares or sending letters via commercial couriers rather than the U.S. Postal Service; neither shopping malls nor Facebook nor any other private company guarantees freedom of speech or privacy. Now, one person can reach hundreds of thousands or even millions of people with a live feed on a cell phone but only as long as the corporate owners permit it and the algorithms that structure the platform surface it to a broad audience. Neither of these is always assured for political content.

Internet platforms are much more than gatekeepers to the broader publics, like the mass media of an earlier era. Facebook also serves other essential communication and assembly functions. Activists also use it as a coffee shop, which scholar Jürgen Habermas famously idealized as the cornerstone of a critical public sphere. For activists, the platform also takes on a resemblance to the office of an underground newspaper—a place to mingle and have back-channel conversations in ways that are reminiscent of their historical antecedents in the alternative print press.¹⁶ It also serves as a living room where families gather to socialize and, having usurped many of the functions of traditional telephones, as a tool that makes one-to-one conversations possible.¹⁷ Facebook thus combines multiple functions that are indispensable to social movements, from the public to the private, for access to large audiences and to facilitate intimate interpersonal transactions. Now all these functions are thus subject to the policies, terms, and algorithms of a single platform.

Despite what seems to be merely a transfer of the same type of dependency from one type of media to another, social media platforms filter, censor, and promote in ways that differ from those of earlier forms of mass media, so the dependencies are not transferred identically. Platforms' power over users rests largely in their ability to set the rules by which attention to content is acquired rather than by picking the winners directly, the way mass media had done in the past. These companies shape the rules, which give them real power, but they are also driven by user demand, creating a new type of networked gatekeeping.

In this chapter, I focus mostly on Facebook and the interaction between its policies and social movement dynamics because Facebook is crucial to many social movements around the world, and there is no real alternative because of its reach and scope. Its putative competitors, such as Twitter, capture a fraction of most populations or, like Instagram, are owned by Facebook. In country after country, Facebook has almost universal reach among internet users, dwarfing other platforms. Together, Google and Facebook capture the vast majority of the advertising money in the digital world.¹⁸ Even so, many of the issues raised in this chapter apply to other platforms as well, even ones with a much smaller reach.

In the past, much scholarship on social movements studied their interaction with mass media and probed the operations of mass media from

many angles, ranging from institutional studies to ethnographies of their employees.¹⁹ In the age of the digital public sphere, digital platforms are a similar topic: their policies, the ideologies of their founders and engineers, the specifics of their legal concerns, their financing models, their terms-of-service and algorithms all interact in important ways with social movement dynamics. I will highlight a few of the most pressing issues, but mine is not an exhaustive list, only a stark demonstration of the power of a few platforms and the reach of their choices.

At the height of Egypt's revolutionary movement in 2010 and early 2011, as I noted in chapter 1, public discontent coalesced around a Facebook page called "We Are All Khaled Said," named after a young man who had been brutally tortured and killed by Egyptian police. Sadly, his death at the hands of the police was not a rare occurrence in Egypt. But Said's story received a significant amount of attention when "before" and "after" photos of him—one showing a smiling young man, the other a mangled, tortured corpse—went viral. The images made the brutality of Egyptian police concrete and symbolized its horror. The Facebook page "We Are All Khaled Said" became the focal point for the agitation of hundreds of thousands of Egyptians. Eventually a call for protests on January 25 posted on that page roused people to action that turned into an uprising. However, that course of events was almost tripped up because of Facebook's "real-name" policy.

One of the most consequential decisions that social media platforms make for their users is whether people can use pseudonyms—and easily create multiple accounts—or whether there is a formal (legal "terms-of-service") requirement that they use their "real" name, however defined. Few platforms require "real names," but Facebook does. Although its policy is something of an exception for internet platforms, it is hugely consequential for social movements because Facebook's dominant size and extent mean that it is used by the ordinary people whom activists want to reach. Facebook acts as a *de facto* public sphere reaching large sections of the population in countries that heavily censor mass media news, leaving platforms like Facebook and Twitter as the only alternatives outside the direct control of the state.

Facebook's policy on real names is not an accident. Trying to force or nudge people to use their "real names" is part of the articulated ideology of Facebook and is central to its business model. The rule is also part of the expressed ideology of its founder (who still controls the platform), Mark Zuckerberg. In reference to pseudonym use, Zuckerberg once said, "Having two identities for yourself is an example of lack of integrity"—a statement ignoring the obvious function of social roles: people live in multiple contexts and they do not behave the same way in each of them.²⁰ A student is not the same way at home, in class, or at a party. For a commercial platform making money from advertising, the advantages of requiring real names are obvious because traceable names allow advertisers to target real people, and to match their information across different settings and databases—following them from voter files to shopping records to their travel and locations. Facebook's policy on names and its method of enforcing its rule have entangled many movements and activists in its web.

The Khaled Said episode, centering as it did on graphic and therefore controversial photographs, echoes an earlier incident in U.S. history, the murder of Emmett Till. Till was a black teenager who had been lynched for allegedly talking to a white woman in Mississippi. His devastated mother held an open-casket funeral for him in Chicago, Till's hometown, that drew tens of thousands of mourners. The inhumanity of the people who had lynched him was exposed in the visage of the mutilated, broken body of the murdered youth. A few newspapers and magazines published grim pictures of Till in the casket. Seeing those images was a galvanizing moment for many persons and exposed many white people to the reality of the ongoing lynchings at a time when the civil rights movement was poised to expand nationally. (The Montgomery bus boycott began within four months of Till's murder.)²¹

Khaled Said's case played a similar role in Egypt. A young Egyptian activist told me about Khaled Said's story and the pictures moved him from being a political bystander to being an activist: "He [Said] wasn't even political. Yet the police tortured and killed him. If it could happen to him, it could happen to anyone, even me."

Wael Ghonim, the administrator of the "We Are All Khaled Said" Facebook page, told me that he had focused on Said's case because it was repre-

sentative and was not tied to a particular political brand or leadership.²² From the stance of an activist, it was a good case to make a point because it was easy to identify with this unlucky young man who had done little more than fall prey to police. Ghonim chose to remain anonymous as the administrator of the page rather than using his legal name to keep attention on political issues rather than himself, but also, importantly, to protect himself and his family from retaliation by Egypt's repressive government. Soon, hundreds of thousands of people began conversing with one another on the page, yet unaware of either the essential role it and they would play in toppling the thirty-year autocracy of Hosni Mubarak or the challenges they were to face just to keep the page open.

In November 2010, a couple of months before the uprising to come, Facebook abruptly deactivated the "We Are All Khaled Said" page. There was immediate speculation that this might be an act of censorship by the Egyptian government. But how had the censorship been accomplished? How was Facebook pressured by the government? An intense discussion raged as puzzled people—including activists around the world—tried to make sense of why the page was yanked.

A Facebook spokesperson confirmed that Facebook made the decision without pressure from the Egyptian government. Facebook deactivated the page because the account holder, Wael Ghonim, had used a pseudonym.²³ Facebook said that his use of a fictitious name was "a violation of our terms," reason enough to delete the page despite its huge following and political significance. Just like that, through its internally decided naming policy, Facebook had censored one of the most important spots for political gathering in Egypt, at the height of political activity, without even a request by the Egyptian government.

The international human rights community pleaded with Facebook to reverse the takedown. In the end, the page was reactivated after a courageous Egyptian woman living abroad offered to allow her real name to be used in connection with the page. Her offer to publicly associate herself with the Said Facebook page, which she made simply to satisfy Facebook's terms of service, meant that she risked permanent exile from her native country and reprisals against members of her family. If she had not stepped up, the page might never have returned and might never have played the

major role it did just a few months later, on January 25, 2011, as one of the top coordination and information sources for Egyptian protesters. Even this reactivation was only possible after employees inside Facebook also stepped up to pressure the company. A page without such visibility might have simply disappeared.

This is far from the only such example. Michael Anti is a Chinese journalist and a former reporter for the Beijing bureau of the *New York Times* who goes by that name in his offline life. He was awarded fellowships at Harvard and Cambridge, and is well known as a democracy activist. Anti specializes in using new media to write about Chinese censorship. In March 2011, he was thrown off Facebook, the place where he stayed in touch with thousands of people. The reason? Even though Michael Anti is what his Chinese friends call him and is his byline in the *New York Times*, the name is a pen name. Anti never uses his legal name, Zhao Jing, which is completely unknown to his circle of friends and colleagues, let alone his readers. Anti angrily decried the contrast between his treatment and that of Facebook cofounder Mark Zuckerberg's puppy, named Beast, which is allowed its own page. Because of Facebook's real-name policy, to this day, Anti does not have a Facebook page.

Even in developed nations where people are not necessarily hiding from the authorities, Facebook's policies cause problems for social movements. LGBTQ people have been some of the sharpest and most vocal critics of Facebook's real-name policies. LGBTQ people may go by names that are different from their legal ones as a preference or as a protection against family members who are angry about their sexual orientation or gender identity and who may act abusively toward them. There have been numerous incidents where LGBTQ people with public visibility—either as activists or as performers—were pushed off Facebook by vengeful people who reported them for not using their real names.²⁴

If you use Facebook, you may be surprised by the preceding stories, and also by my claim that activists regularly encounter problems with the real-name policies, because you may have noticed that some Facebook friends do not use their real names. The vast majority of people use their real name on Facebook. Although a significant minority do not, they never encounter problems as long as they are not political activists.

It is true that a sizable portion of Facebook's users avoid using legal names on the site. In my surveys of college students, I often find that about 20 percent use a nonlegal name—often nicknames known only to their friends. A quick search reveals that there are many individuals on Facebook who use names like "Santa Claus" or "Mickey Mouse" and continue to have a perfectly normal Facebook experience. Why, then, did Wael Ghonim run into such trouble?

Facebook's real-name policy, like most policies of almost all social media platforms, is implemented through "community policing"—a method with significantly different impacts depending on the community involved. Community policing means that the company acts only if and when something is reported to it and mostly ignores violations that have not been flagged by members of the community. This model, also called "report and takedown," is encouraged by U.S. laws that declare that these platforms are not legally responsible for content that gets posted unless they fail to take down items that they are told violate the law. Community policing puts social movement activists—indeed, anyone with visibility—at a distinct disadvantage. The more people who see you—especially if you are commenting on or advocating for social movements or on politically sensitive issues, which makes you more of a target—the more opportunities there are for someone to report you.

This model also allows the companies to have a very small staff compared with their user base, significantly lowering their expenses. For example, at its height, General Motors employed hundreds of thousands directly and perhaps millions indirectly through its supply chain. In contrast, Facebook directly employs a little more than 12,600 people despite a user base of 1.5 billion. This combination of legal shelter for "report and takedown" and dramatically lower costs means that the model of a tiny employee base compared with the number of users, and indifference to terms-of-service violations of ordinary users, is common among software companies.

However, activists are not ordinary users of social media. People active in social movements tend to be more public, focus on outreach to people beyond their immediate social networks, and hold views that might be minority perspectives, polarizing stances, or opinions targeted by govern-

ments. Activists are more likely to be targeted for reporting by social media “community” members, people who oppose their ideology, or the authorities or people in the pay of the authorities. If your Facebook friends are close friends and acquaintances, you generally mind your own business, and do not comment publicly on potentially controversial matters, no one is likely to report you for calling yourself Mickey Mouse. Activists behave exactly the opposite way on Facebook. Activists ruffle feathers and challenge authorities. Most activists I have interacted with over the years make many of their political posts public (visible to everyone, not just their Facebook friends) to try to attract attention to their ideas. Activists also often try to broaden their social networks as much as they can in order to spread their message. Many activists I know maintain thousands of friends on social media and in many other ways stand out from the crowd.

All this leaves movements vulnerable to being targeted directly through community policing because their opponents seek to report them for infractions, real or imagined. Often, such reporting takes place in an organized manner, which means that companies are more likely to take it seriously as if it were a real infraction since the number of complaints is high. For example, on Turkish Twitter, there are often calls for reporting political opponents as “spam” to the degree that spam has now become a verb: “Can we please spam Myopponent99?” (meaning not “Let’s send spam to Myopponent99” but “Let’s all falsely report Myopponent99 as a spammer and hope that results in the account getting suspended”). Such mass reporting of opponents as spam or abusive is often successful in getting accounts suspended temporarily or even permanently. And this does not happen only in other countries; even in the United States, false reports of violations of terms of service are routinely attempted and sometimes successful—often targeting feminists, LGBTQ people, or political dissidents.

Activists, especially in repressive countries, use nicknames on Facebook for a variety of reasons. For example, I have seen activists use pseudonyms to keep random vigilantes from finding their home addresses—they are not necessarily hiding who they are, but just making it not too easy for people with low motivation or competence to quickly find them. If opponents report them, their accounts are in jeopardy unless they begin using their

legal names, which must be verified by submitting documents like images of a driver’s license or passport in what can be a risky and time-consuming process. Just the verification process may endanger their lives, depending on the severity of the repression in the country. I have seen this happen repeatedly but will not list examples—it would put these activists at further risk.

Even activists who use their real names are at risk of having their accounts suspended when political opponents and authorities make false accusations against them. When activists are reported, even if the report is false, they often must go through the verification process anyway, which sometimes disables their profile for weeks, especially in cases when their non-English but accurate, real names appear plausibly fake to Facebook’s employees or algorithms.

After a great deal of harsh criticism, Facebook has slightly modified its policy, shifting to “first and last names” people use in everyday life. However, the documents that they accept for account verification are almost overwhelmingly legal documents such as checks, credit cards, medical records, and bank statements. Some of the choices they accept for identity verification, such as a yearbook photo, may work for Western activists, but activists or LGBTQ people in developing countries rarely have these options. Ironically, implementing these slight modifications to the real-name policy may have taken some of the heat off Facebook because LGBTQ communities in Western nations, those in the best position to make noise about their plight, have found ways to work with the company, but non-Western activists and affected communities elsewhere around the world, who have a lot less power vis-à-vis Facebook, continue to suffer.

In one instance, a politically active Facebook friend of mine who lives in a Middle Eastern country racked by violence was caught in a catch-22. Facebook’s terms of service mandate “no vulgar names.” But vulgar in what language? Her very real and legal non-English name corresponds to a vulgar word in English—which ended up with her account getting suspended. To get around this cultural imbroglio, she tried to use a nickname, but Facebook then asked her to verify that it was her legal name. She could not because it was not. She ended up having to send many copies of her passport over Facebook’s system, a process that put her at risk of identity theft. She repeated the process many times, getting suspended on and off, sometimes

because of her “vulgar” name other times because her replacement name was a nickname. She was finally able to reinstate her account after much effort, largely because she was connected to people who could alert Facebook to the issue. For others, such an ordeal might mean that they are, in effect, banned from the biggest public square in the world, which is also the biggest private social network. The stakes could hardly be higher.

What determines the kind of content that is allowed on platforms and the kind that is removed, censored, or suppressed? There is no simple answer. Platforms operate under a variety of pressures, ranging from commercial and legal ones to their own business models, as well as their ideological choices. Many platforms have explicit policies about the type of content that is allowed or forbidden. These policies are partly driven by financial concerns but are also influenced by a company's own vision of its platform. Many of these companies are quite young and are run by founders who own large amount of stock. Therefore, the role of individual ideology is greater than it is in an established, traditional company that is fully securitized and subject only to Wall Street considerations. Platforms are also subject to a multitude of different legal regimes because they operate in countries with dissimilar and sometimes conflicting free-speech, hate-speech, libel, and slander laws. Tellingly, intellectual property laws are a prominent exception to the rule “Let the community handle it.” Copyright, an aspect of intellectual property law, is generally implemented in a much more proactive and comprehensive manner. Somewhat unsurprisingly, social media platforms, which are corporate entities, are far more concerned about intellectual property rights that corporations care most about, and where they have more legal remedies, than about individual privacy or political agency.²⁵

The most important social media platforms for social movements, Facebook and Twitter, and the video-sharing service YouTube, owned by Google, have significantly different terms of service reflecting various editorial policies as well as the norms adopted by users. In the more freewheeling Twitterverse, fairly little is banned by the terms of service, although Twitter has been making some of its rules stricter (or at least applying them more strictly). In particular, Twitter has been pressured to act because of concerns about abuse, especially of female and/or minority people and ac-

tivists, the use of the platform by groups seeking or inciting violence, racism, hate speech (illegal in much of Europe), and lately the rise of ISIS in the Middle East.

Facebook, on the other hand, has stricter rules and is more trigger-happy in deleting content for terms-of-service violations. Facebook has removed content ranging from breast-feeding pictures to posts considered to denigrate a community, often with little recourse for the people whose posts are removed. In September 2016, Facebook removed a post by a Norwegian journalist because it included a picture of a naked child. The picture was the Pulitzer Prize-winning 1972 photo showing a nine-year-old Vietnamese girl, Phan Thi Kim Phuc, running naked and screaming “Too hot, too hot,” having just been badly burned by a napalm attack. The picture had been published on the front page of the *New York Times* and seared into many people's memories as a symbol of the brutality of the war in Vietnam. It had been reprinted countless times as an iconic photo showing the tragedy of war.

Facebook was criticized for censoring the post and was rebuked by the prime minister of Norway, who also had posted the photo to the platform. Facebook then responded by deleting the prime minister's post as well. After global expressions of outrage, including stories in leading traditional newspapers, Facebook finally backed down and reinstated the post. It's worth pondering what might have happened if Facebook had been the dominant channel of distribution in 1972. Except for publicity campaigns to pressure Facebook to reverse its decisions, users have little or no recourse against the actions Facebook takes.

Making these types of decisions is not straightforward, nor are there easy answers—especially ones that scale with the low employment business model of technology giants. Google, too, has struggled, especially because its video platform, YouTube, is a major means of propaganda for both activists and terrorists, ranging to ISIS beheadings in the Middle East and rampaging mass shooters in the United States. An activist in Egypt once recounted to me his battles with Google about taking down content that depicted violence. A policy against depictions of violence might seem to make sense when the video depicts an ISIS beheading or a mass shooting. But what about a video that documents misconduct of the police or the

army? Some of the videos were horrifying, but, as the activist told me, "That was the only way we could get the word out." In response to the pressure, Google decided to allow such videos to remain on the site because of their "news value." Only a few years later, other antiviolence activists tried to pressure Google to take down videos showing beheadings carried out by ISIS. This policy too, was applied inconsistently. Videos of Westerners murdered at the hands of ISIS were removed fairly quickly, while similar videos of executions of local Syrians, Iraqis, or Libyans often remained on YouTube. As this example shows, there is no simple, easy-to-implement answer or method that applies uniformly to all cases, which means such decisions can neither be easily automated nor outsourced to lowly-paid, harried employers.

To get a better grasp of the complexities of the policies and practices that govern what content is allowed or disallowed on social media platforms, let us look at the example of activists and political parties in Turkey aligned with a particular perspective on the Kurdish issue in the country. The military coup of 1980 in Turkey unleashed a brutal wave of repression that was especially harsh in Kurdish southeastern Turkey. In the same period, an armed militant group called the Kurdistan Workers' Party (PKK) launched what would become a multidecade insurgency. The conflict claimed forty thousand lives, mostly in the 1980s and 1990s. I lived in Turkey for most of those years but knew few details about the situation—besides the fact that something awful was going on—because coverage was heavily censored both on state television and in privately held newspapers.

Change came in 2002, when a new Islamist-leaning party without the same historical commitment to Turkish ethnic identity, the Justice and Development Party (AKP), came to power. After a few years, the AKP government initiated a peace process with the PKK, resulting in a fragile cease-fire and improved laws that allowed Kurdish identity to be expressed more explicitly. At the same time, a mostly Kurdish political party also flourished in the region, capturing a majority of the votes in many Kurdish cities, often overwhelmingly. But even though there was no longer just one, state-owned, television station in Turkey, the mass media remained indirectly constrained through pressures on the media's corporate owners.²⁶ At the time,

the southeastern Kurdish region was generally calm (a situation that would change around 2013 and significantly worsen after 2015), and censorship of the mass media was not the primary problem, at least for Kurds.

Despite this more open political environment, for years Kurdish politicians were censored on Facebook. The Facebook page of the mayor of the biggest majority Kurdish city in the region was banned even though almost four hundred thousand people had "liked" his page before it was taken down. The list of Kurds who were banned from Facebook ranged from prominent authors to elected deputies (parliamentary officials). The suppression encompassed an assortment of pages such as a site for Kurdish music and other popular, even mundane pages with hundreds of thousands of followers and likes. Yet Facebook did not provide clear explanations of the reasons for prohibiting the pages. Most of the time, it offered a terse statement about "violations of community guidelines." Some Facebook messages claimed that the proscribed sites had hosted pornography (which, given the traditional nature of the community, seemed quite unlikely). Sometimes no explanation was given. Administrators of these sites appealed, but written requests to Facebook for explanations often went unanswered.

People asked whether the censorship was a result of government pressure. This did not make sense because the same Kurdish officials appeared on traditional news media even as their Facebook pages were blocked and banned.

Curious about the censorship mechanism, I started following these pages, and asked many people in Turkey, including free-speech activists and lawyers, as well as officials, whether they were aware of court orders or backchannel pressures from the government on Facebook to ban Kurdish politics. I knew that many suspected that the government was behind these closures, because Kurdish content had often been suppressed in earlier years. However, all the people I spoke with, including sources close to the government, said that they were not lobbying or communicating with Facebook about these pages. I could find neither motive nor evidence of government interference. It was a mystery.

Some light was shed on the matter when I talked to high-level employees from Facebook, including Richard Allan, Facebook's vice president for

public policy, who oversees European and Middle Eastern countries. Allan, a friendly, sharp, and knowledgeable Englishman, listened as I voiced my concerns, and he then walked me through the process. He explained that Facebook had adopted the U.S. State Department's list of "terrorist organizations," which included the Kurdish insurgent group, the PKK. He also assured me that Facebook was taking down only content that promoted violence.

His statement would suggest that Facebook was banning only PKK content. But this did not fully solve the mystery since deputies who had been legally elected, established journalists, and even some Kurdish culture pages were also censored, their pages shut down on and off. There was also much banning of items such as journalists' reports of public events, even when the events were written about in Turkish newspapers without issues. After examining the banned Facebook pages, I realized that the trouble seemed to be that Facebook was failing to distinguish PKK propaganda from ordinary content that was merely about Kurds and their culture, or news about the group or the insurgency. It was like banning any Irish page featuring a shamrock or a leprechaun as an Irish Republican Army page, along with BBC reports about "the troubles" in Northern Ireland.

For example, in March 2015, during the Kurdish New Year celebrations, a Turkish journalist posted on Instagram—a site owned by Facebook—a picture she had taken showing, ironically, elderly Kurdish women who had symbolically taped their mouths shut, wearing T-shirts with the PKK's imprisoned leader's visage suggested by a distinctive outline of black hair and mustache overlain on their white shirts. The reporting suggested that they were protesting the fact that the imprisoned leader of the group had not met with his lawyers recently. Instagram quickly censored the picture, taking the whole post down, and Facebook did the same on the journalist's page. The same thing happened to pictures of the same rally from another prominent Turkish journalist.

Both journalists were known to be sympathetic to Kurdish rights, and Turkish nationalists had long targeted them on social media. But all they had done was post a picture from a public, legal rally of some women wearing a t-shirt with a suggestive outline of the jailed leader. The picture was

clearly newsworthy; similar photos from the rally had even appeared on pro-government outlets in Turkey. The outraged journalists loudly took to Twitter, where they had a large following, and complained about the censorship of their pictures. Facebook and Instagram later reinstated the pictures, as well as pictures from the same rally posted by other journalists. But Facebook's reversal occurred only after the journalists' public protests achieved a substantial amount of attention, which less prominent people might not have garnered.

A leaked document from Facebook's monitoring team provided a key insight; it showed that Facebook instructed employees to ban "any PKK support or [PKK-related] content with no context" or "content supporting or showing PKK's imprisoned founder."²⁷ One possible explanation of what was happening was that Turks who held strong nationalist views were using the community-policing mechanism to report Kurdish pages on which pictures from rallies or other political events from Kurdish cities appeared, even when the image was merely a photo taken in public or as part of a news story, and that Facebook employees who oversaw Turkish content monitoring were targeting those pages, either out of ignorance or perhaps because they were also Turkish nationalists opposed to Kurdish rights—a potential problem for platforms such as Facebook in a country with so much domestic polarization. In fact, in almost any country with deep internal conflict, the types of people who are most likely to be employed by Facebook are often from one side of the conflict—the side with more power and privileges.

Facebook's team overseeing the monitoring for Turkey is also located in Dublin, likely disadvantaging anyone who could not relocate to a European country, or does not speak English. Although I do not have statistics, I have, for example, heard from other sources that this puts women at a distinct disadvantage in the Middle East because their families are less likely to locate outside their home country for the benefit of employment at Facebook. The moderation teams—already pretty small—represent thus but a privileged slice of the countries that they oversee.

It is also possible that workers who knew little about the Turkish or Kurdish context and, possibly, who were not even formally employees of Facebook, did much of this monitoring. Journalists who have investigated

the content-monitoring industry have often found that these decisions are outsourced to low-paid workers in countries like the Philippines, who must look at a vast amount of content and make rapid-fire decisions under strict time constraints, sometimes barely a few seconds per decision.²⁸ Could these workers wade through the nuances of an already-difficult decision-making process and adequately judge the items with news value, those protected by freedom of speech, and those that were an incitement to violence—especially about countries where they had never been and where they did not understand the language? Or did they mostly make decisions in response to the volume of complaints received, something that is easy to quantify and organize?

These are complex situations without easy solutions. In July 2015, a few months after that picture of elderly Kurdish women engaged in a symbolic protest was censored, the cease-fire between Kurdish militants and the Turkish government collapsed again, and the insurgency picked up steam, resulting in more deaths. When reporters cover conflicts, the line between news value and propaganda is not always clear, especially when dealing with the huge numbers of user-generated images. In a nationalist, armed insurgency, where is the line between freedom of the press and images that might fuel a war or be propaganda for acts of terrorism that result in many deaths? And who is qualified to make those decisions?

In the United States, where the First Amendment of the Constitution guarantees broader freedom of speech than in almost any other major country, it may seem that the straightforward answer is to allow all types of content. However, even with the First Amendment as a legal framework, a zero-interference policy would run into problems. The U.S. government sometimes seeks to ban content that it considers a threat to itself. This includes posts by ISIS, which uses social media to recruit disaffected people or incite them to commit acts of terrorism. The United States also has strong copyright protections, and thus these platforms are under legal constraints to remove copyrighted content. What about other real cases, such as a graphic picture of someone's death posted on the internet for the purpose of harassing that person's loved ones? What about revenge porn, when a jilted ex-boyfriend releases or steals nude pictures of his ex-girlfriend or

wife and posts them as a malicious act (many real cases)? There are many other examples.

Governments, too, have increasingly learned to use these mechanisms to silence dissent. They can encourage or even pay crowds to purposefully "report" dissidents to get them banned or at least make them struggle to stay on a platform. In these cases, the authorities count on taking advantage of the thinly staffed, clunky enforcement mechanisms of major platforms. Officials can also directly pressure the companies.

Michael Anti's problems with technology companies did not begin with Facebook's real-name policies. In 2006, Anti had a popular Microsoft blogging platform that drew the ire of the Chinese government. Microsoft, which does much business in China, shut down his blog at the government's behest.²⁹ In another case, the internet giant Yahoo provided the details of the e-mail account of Chinese journalist and poet Shi Tao. Shi had used a Yahoo account to pseudonymously release a Communist Party document to a pro-democracy website. The authorities had no easy way to track down the whistleblower, so they turned to Yahoo. After Yahoo turned over information identifying Shi, he was sentenced to ten years in prison and forced labor. The case attracted widespread attention after Amnesty International declared Shi a prisoner of conscience and Shi received an International Press Freedom Award from the Committee to Protect Journalists. After the human rights backlash, Yahoo's CEO apologized to Shi's family. Still, the damage was done. Shi spent almost nine years in prison, and his family members were harassed by the authorities.³⁰ In 2016, it was also revealed that Yahoo secretly scanned user e-mails at the behest of the U.S. intelligence services, raising questions about the Fourth Amendment, which protects against search and seizure without due process.³¹

Activists trying to reach broader publics find themselves waging new battles, beyond those that involve conflict and negotiation with large media organizations. There is a new era for the dynamics of gatekeeping in the new, digital public sphere, and it is far from a simple one. I have discussed the downsides to social movements of these policies; but this doesn't mean that there is a perfect, easy answer to the question, nor a means to do this both ethically and at scale through automation or poorly-paid contractors

judging content in countries not their own. Major platforms could do a lot better by investing resources and giving more attention to the issue, but that their business model, their openness to government pressure, and sometimes their own mindset, often works against this.

Social media platforms increasingly use algorithms—complex software—to sift through content and decide what to surface, prioritize, and publicize and what to bury. These platforms create, upload, and share user-generated content from hundreds of millions, if not billions, of people, but most platforms do not and cannot show everything to everyone. Even Twitter, which used to show content chronologically—content posted last is seen first—is increasingly shifting to algorithmic control.

Perhaps the most important such algorithm for social movements is the one Facebook uses which sorts, prioritizes, and filters everyone's "news feed" according to criteria the company decides. Google's success is dependent on its page-ranking algorithm that distills a page of links from the billions of possible responses to a search query.

Algorithmic control of content can mean the difference between widespread visibility and burial of content. For social movements, an algorithm can be a strong tailwind or a substantial obstacle.³² Algorithms can also shape social movement tactics as a movement's content producers adapt or transform their messages to be more algorithm friendly.

Consider how the Black Lives Matter movement, now nationwide in the United States, encountered significant algorithmic resistance on Facebook in its initial phase. After a police officer killed an African American teenager in Ferguson, Missouri, in August 2014, there were protests in the city that later sparked nationwide demonstrations against racial inequalities and the criminal justice system. However, along the way, this burgeoning movement was almost tripped up by Facebook's algorithm.

The protests had started out small and local. The body of Michael Brown, the black teenager shot and killed by Ferguson police officer Darren Wilson on August 9, had been left in the street for hours. The city was already rife with tensions over race and policing methods. Residents were upset and grieving. There were rumors that Brown's hands had been up in the air when he was shot.

When the local police in Ferguson showed up at the first vigils with an aggressive stance, accompanied by dogs, the outrage felt by residents spread more broadly and brought in people who might not have been following the issue on the first day. The Ferguson situation began to attract some media attention. There had been tornadoes in Missouri around that time that had drawn some national journalists to the state. As reports of the use of tear gas during nightly protests started pouring in, journalists went to Ferguson. Ferguson residents started live-streaming video as well, although at this point, the protests were mostly still a local news story.

On the evening of August 13, the police appeared on the streets of Ferguson in armored vehicles and wearing military gear, with snipers poised in position and pointing guns at the protesters. That is when I first noticed the news of Ferguson on Twitter—and was startled at such a massive overuse of police force in a suburban area in the United States. The pictures, essentially showing a military-grade force deployed in a small American town, were striking. The scene looked more like Bahrain or Egypt, and as the Ferguson tweets spread, my friends from those countries started joking that their police force might have been exported to the American Midwest.

Later that evening, as the streets of Ferguson grew tenser, and the police presence escalated even further, two journalists from prominent national outlets, the *Washington Post* and the *Huffington Post*, were arrested while they were sitting at a McDonald's and charging their phones. The situation was familiar to activists and journalists around the world because McDonald's and Starbucks are where people go to charge their batteries and access Wi-Fi. The arrest of the reporters roused more indignation and focused the attention of many other journalists on Ferguson.

On Twitter, among about a thousand people around the world that I follow, and which was still sorted chronologically at the time, the topic became dominant. Many people were wondering what was going on in Ferguson—even people from other countries were commenting. On Facebook's algorithmically controlled news feed, however, it was as if nothing had happened.³³ I wondered whether it was me: were my Facebook friends just not talking about it? I tried to override Facebook's de-

fault options to get a straight chronological feed. Some of my friends were indeed talking about Ferguson protests, but the algorithm was not showing the story to me. It was difficult to assess fully, as Facebook keeps switching people back to an algorithmic feed, even if they choose a chronological one.

As I inquired more broadly, it appeared that Facebook's algorithm—the opaque, proprietary formula that changes every week, and that can cause huge shifts in news traffic, making or breaking the success and promulgation of individual stories or even affecting whole media outlets—may have decided that the Ferguson stories were lower priority to show to many users than other, more algorithm-friendly ones. Instead of news of the Ferguson protests, my own Facebook's news feed was dominated by the “ice-bucket challenge,” a worthy cause in which people poured buckets of cold water over their heads and, in some cases, donated to an amyotrophic lateral sclerosis (ALS) charity. Many other people were reporting a similar phenomenon.

There is no publicly available detailed and exact explanation about how the news feed determines which stories are shown high up on a user's main Facebook page, and which ones are buried. If one searches for an explanation, the help pages do not provide any specifics beyond saying that the selection is “influenced” by a user's connections and activity on Facebook, as well as the “number of comments and likes a post receives and what kind of a story it is.” What is left unsaid is that the decision maker is an algorithm, a computational model designed to optimize measurable results that Facebook chooses, like keeping people engaged with the site and, since Facebook is financed by ads, presumably keeping the site advertiser friendly.

Facebook's decisions in the design of its algorithm have great power, especially because there is a tendency for users to stay within Facebook when they are reading the news, and they are often unaware that an algorithm is determining what they see. In one study, 62.5 percent of users had no idea that the algorithm controlling their feed existed, let alone how it worked.³⁴ This study used a small sample in the United States, and the subjects were likely more educated about the internet than many other

populations globally, so this probably underestimates the degree to which people worldwide are unaware of the algorithm and its influence. I asked a class of 20 bright and inquisitive students at the University of North Carolina, Chapel Hill, a flagship university where I teach, how they thought Facebook decided what to show them on top of their feed. Only two knew it was an algorithm. When their friends didn't react to a post they made, they assumed that their friends were ignoring them, since Facebook does not let them know who did or didn't see the post. When I travel around the world or converse with journalists or ethnographers who work on social media, we swap stories of how rare it is to find someone who understands that the order of posts on her or his Facebook feed has been chosen by Facebook. The news feed is a world with its own laws, and the out-of-sight deities who rule it are Facebook programmers and the company's business model. Yet the effects are so complex and multilayered that it often cannot be said that the outcomes correspond exactly to what the software engineers intended.

Our knowledge of Facebook's power mostly depends on research that Facebook explicitly allows to take place and on willingly released findings from its own experiments. It is thus only a partial, skewed picture. However, even that partial view attests how much influence the platform wields.

In a Facebook experiment published in *Nature* that was conducted on a whopping 61 million people, some randomly selected portion of this group received a neutral message to “go vote,” while others, also randomly selected, saw a slightly more social version of the encouragement: small thumbnail pictures of a few of their friends who reported having voted were shown within the “go vote” pop-up.³⁵ The researchers measured that this slight tweak—completely within Facebook's control and conducted without the consent or notification of any of the millions of Facebook users—caused about 340,000 additional people to turn out to vote in the 2010 U.S. congressional elections. (The true number may even be higher since the method of matching voter files to Facebook names only works for exact matches.³⁶) That significant effect—from a one-time, single tweak—is more than four times the number of votes that determined that Donald

Trump would be the winner of the 2016 election for presidency in the United States.

In another experiment, Facebook randomly selected whether users saw posts with slightly more upbeat words or more downbeat ones; the result was correspondingly slightly more upbeat or downbeat posts by those same users. Dubbed the “emotional contagion” study, this experiment sparked international interest in Facebook’s power to shape a user’s experience since it showed that even people’s moods could be affected by choices that Facebook made about what to show them, from whom, and how.³⁷ Also, for many, it was a revelation that Facebook made such choices at all, once again revealing how the algorithm operates as a hidden shaper of the networked public sphere.

Facebook’s algorithm was not prioritizing posts about the “Ice Bucket Challenge” rather than Ferguson posts because of a nefarious plot by Facebook’s programmers or marketing department to bury the nascent social movement. It did not matter whether its programmers or even its managers were sympathetic to the movement. The algorithm they designed and whose priorities they set, combined with the signals they allowed users on the platform to send, created that result.

Facebook’s primary signal from its users is the infamous “Like” button. Users can click on “Like” on a story. “Like” clearly indicates a positive stance. The “Like” button is also embedded in millions of web pages globally, and the blue thumbs-up sign that goes with the “Like” button is Facebook’s symbol, prominently displayed at the entrance to the company’s headquarters at One Hacker Way, Menlo Park, California. But there is no “Dislike” button, and until 2016, there was no way to quickly indicate an emotion other than liking.³⁸ The prominence of “Like” within Facebook obviously fits with the site’s positive and advertiser-friendly disposition.

But “Like” is not a neutral signal. How can one “like” a story about a teenager’s death and ongoing, grief-stricken protests? Understandably, many of my friends were not clicking on the “Like” button for stories about the Ferguson protests, which meant that the algorithm was not being told that this was an important story that my social network was quite inter-

ested in. But it is easy to give a thumbs-up to a charity drive that involved friends dumping ice water on their heads and screeching because of the shock in the hot August sun.

From press reporting on the topic and from Facebook’s own statements, we know that Facebook’s algorithm is also positively biased toward videos, mentions of people, and comments. The ALS ice-bucket challenge generated many self-made videos, comments, and urgings to others to take the challenge by tagging them with their Facebook handles. In contrast, Ferguson protest news was less easy to comment on. What is one supposed to say, especially given the initial lack of clarity about the facts of the case and the tense nature of the problem? No doubt many people chose to remain silent, sometimes despite intense interest in the topic.

The platforms’ algorithms often contain feedback loops: once a story is buried, even a little, by the algorithm, it becomes increasingly hidden. The fewer people see it in the first place because the algorithm is not showing it to them, the fewer are able to choose to share it further, or even to signal to the algorithm that it is an important story. This can cause the algorithm to bury the story even deeper in an algorithmic spiral of silence.

The power to shape experience (or perhaps elections) is not limited to Facebook. For example, rankings by Google—a near monopoly in searches around the world—are hugely consequential. A politician can be greatly helped or greatly hurt if Google chooses to highlight, say, a link to a corruption scandal on the first page of its results or hide it in later pages where very few people bother to click. A 2015 study suggested that slight changes to search rankings could shift the voting preferences of undecided voters.³⁹

Ferguson news managed to break through to national consciousness only because there was an alternative platform without algorithmic filtering and with sufficient reach. On the chronologically organized Twitter, the topic grew to dominate discussion, trending locally, nationally, and globally and catching the attention of journalists and broader publics.⁴⁰ After three million tweets, the national news media started covering the story too, although not until well after the tweets had surged.⁴¹ At one point, before mass-media coverage began, a Ferguson live-stream video

had about forty thousand viewers, about 10 percent of the nightly average on CNN at that hour.⁴² Meanwhile, two seemingly different editorial regimes, one algorithmic (Facebook) and one edited by humans (mass media), had simultaneously been less focused on the Ferguson story. It's worth pondering if without Twitter's reverse chronological stream, which allowed its users to amplify content as they choose, unmediated by an algorithmic gatekeeper, the news of unrest and protests might never have made it onto the national agenda.⁴³

The proprietary, opaque, and personalized nature of algorithmic control on the web also makes it difficult even to understand what drives visibility on platforms, what is seen by how many people, and how and why they see it. Broadcast television can be monitored by anyone to see what is being covered and what is not, but the individualized algorithmic feed or search results are visible only to their individual users. This creates a double challenge: if the content a social movement is trying to disseminate is not being shared widely, the creators do not know whether the algorithm is burying it, or whether their message is simply not resonating.

If the nightly television news does not cover a protest, the lack of coverage is evident for all to see and even to contest. In Turkey, during the Gezi Park protests, lack of coverage on broadcast television networks led to protests: people marched to the doors of the television stations and demanded that the news show the then-widespread protests. However, there is no transparency in algorithmic filtering: how is one to know whether Facebook is showing Ferguson news to everyone else but him or her, whether there is just no interest in the topic, or whether it is the algorithmic feedback cycle that is depressing the updates in favor of a more algorithm-friendly topic, like the ALS charity campaign?

Algorithmic filtering can produce complex effects. It can result in more polarization and at the same time deepen the filter bubble.⁴⁴ The bias toward "Like" on Facebook promotes the echo-chamber effect, making it more likely that one sees posts one already agrees with. Of course, this builds upon the pre-existing human tendency to gravitate toward topics and positions one already agrees with—confirmation bias—which

is well demonstrated in social science research. Facebook's own studies show that the algorithm contributes to this bias by making the feed somewhat more tilted toward one's existing views, reinforcing the echo chamber.⁴⁵

Another type of bias is "comment" bias, which can promote visibility for the occasional quarrels that have garnered many comments. But how widespread are these problems, and what are their effects? It is hard to study any of this directly because the data are owned by Facebook—or, in the case of search, Google. These are giant corporations that control and make money from the user experience, and yet the impact of that experience is not accessible to study by independent researchers.

Social movement activists are greatly attuned to this issue. I often hear of potential tweaks to the algorithm of major platforms from activists who are constantly trying to reverse-engineer them and understand how to get past them. They are among the first people to notice slight changes. Groups like Upworthy have emerged to produce political content designed to be Facebook algorithm friendly and to go viral. However, this is not a neutral game. Just as attracting mass-media attention through stunts came with political costs, playing to the algorithm comes with political costs as well. Upworthy, for example, has ended up producing many feel-good stories, since those are easy to "Like," and thus please Facebook's algorithm. Would the incentives to appease the algorithm make social movements gear towards feel-good content (that gets "Likes") along with quarrelsome, extreme claims (which tend to generate comments?)—and even if some groups held back, would the ones that played better to the algorithm dominate the conversation? Also, this makes movements vulnerable in new ways. When Facebook tweaked its algorithm to punish sites that strove for this particular kind of virality, Upworthy's traffic suddenly fell by half.⁴⁶ The game never ends; new models of virality pop up quickly, sometimes rewarded and other times discouraged by the central platform according to its own priorities.

The two years after the Ferguson story saw many updates to Facebook's algorithm, and a few appeared to be direct attempts to counter the biases that had surfaced about Ferguson news. The algorithm started taking into

account the amount of time a user spent hovering over a news story—not necessarily clicking on it, but looking at it and perhaps pondering it in an attempt to catch an important story one might not like or comment on—and, as previously noted, programmers implemented a set of somewhat harder-to-reach but potentially available Facebook reactions ranging from “sad” to “angry” to “wow.” The “Like” button, however, remains preeminent, and so does its oversized role in determining what spreads or disappears on Facebook.

In May 2016, during a different controversy about potential bias on Facebook, a document first leaked to *The Guardian* and then released by Facebook showed a comparison of “trends” during August 2014. In an indirect confirmation of how the Ferguson story was shadowed by the ALS ice-bucket challenge, the internal Facebook document showed that the ALS ice-bucket challenge had overwhelmed the news feed, and that posts about Ferguson had trailed.⁴⁷

Increasingly, pressured by Wall Street and advertisers, more and more platforms, including Twitter, are moving toward algorithmic filtering and gatekeeping. On Twitter, an algorithmically curated presentation of “the best Tweets first” is now the default, and switching to a straight chronological presentation requires navigating to the settings menu. Algorithmic governance, it appears, is the future and the new overlords that social movements must grapple with.

The networked public sphere is not a flat, open space with no barriers and no structures. Sometimes, the gatekeepers of the networked public sphere are even more centralized and sometimes even more powerful than those of the mass media, although their gatekeeping does not function in the same way. Facebook and Google are perhaps historically unprecedented in their reach and their power, affecting what billions of people see on six continents (perhaps seven; I have had friends contact me on social media from Antarctica). As private companies headquartered in the United States, these platforms are within their legal rights to block content as they see fit. They can unilaterally choose their naming policies, allowing people to use pseudonyms or not. Their computational processes filter and prioritize content, with significant consequences.

This means a world in which social movements can potentially reach hundreds of millions of people after a few clicks without having to garner the resources to challenge or even own mass media, but it also means that their significant and important stories can be silenced by a terms-of-service complaint or by an algorithm. It is a new world for both media and social movements.