



PROJECT MUSE®

## Making the Internet Safe for Democracy

Francis Fukuyama

Journal of Democracy, Volume 32, Number 2, April 2021, pp. 37-44 (Article)

Published by Johns Hopkins University Press

DOI: <https://doi.org/10.1353/jod.2021.0017>



➔ *For additional information about this article*

<https://muse.jhu.edu/article/787834>

# MAKING THE INTERNET SAFE FOR DEMOCRACY

*Francis Fukuyama*

*Francis Fukuyama is Mosbacher Director of Stanford's Center on Democracy, Development and the Rule of Law. This article reflects the findings of the Stanford Working Group on Platform Scale and its report at <https://cyber.fsi.stanford.edu/publication/report-working-group-platform-scale>.*

Many people have come to see the internet as one of the chief threats to contemporary democracy. The internet, and large platforms such as Google, Facebook, and Twitter in particular, have been blamed for the rise of Donald Trump and the populism he represents, the proliferation of conspiracy theories and fake news, and the intense political polarization afflicting the United States as well as many other democracies. Across the world, politicians with authoritarian leanings, such as Rodrigo Duterte in the Philippines and Narendra Modi in India, have made effective use of Facebook and Twitter to reach their followers and attack opponents.

There is, nonetheless, a great deal of confusion as to where the real threat to democracy lies. This confusion begins with a question of causality: Do the platforms simply reflect existing political and social conflicts, or are they actually the cause of such conflicts? The answer to that question will in turn be key to finding the appropriate remedies.

This issue came to a head in the aftermath of the 6 January 2021 mob assault on the U.S. Congress that was instigated by the outgoing President Trump. In the wake of that violence, Twitter shut down Trump's account, cutting him off from the primary channel that he had used to communicate with his followers. While many people applauded this decision and even saw it as overdue, others worried about the sheer power that Twitter had amassed. President Trump was indeed effectively muzzled in the days following the ban. Conservatives immediately castigated the move—and the parallel actions by Facebook, Google, and Amazon that soon followed—for what they labeled “censorship.” And while one may approve of Twitter's decision as a short-run response to

the danger of violent incitement, conservative critics of this move raise legitimate points about the dangers of platform power.

Legally speaking, the censorship charge falls flat. In U.S. law, the First Amendment's prohibition of censorship applies only to government actions; the Amendment actually protects the right of private parties such as Twitter and Facebook to publish whatever content they want. Beyond these protections, online platforms have been shielded from certain forms of liability by Section 230 of the 1996 Communications Decency Act. The problem we face today, however, is one of scale: These platforms are so large that they have come to constitute a "public square" within which citizens contest issues and ideas. There are plenty of private corporations that curate the information they publish; these are media companies, with names such as the *New York Times* or the *Wall Street Journal*. But none of these legacy media companies is as dominant or reaches as many people as Twitter, Facebook, and Google. The scale of these internet platforms is great enough that decisions made by their owners could impact the outcome of democratic elections in a way that legacy media companies' decisions could not.

The other big problem with the large internet platforms is one of transparency. While Twitter publicly announced its ban of President Trump, it, Facebook, and Google make literally thousands of content-curation decisions each day. The great mass of takedowns are relatively uncontroversial, as with those targeting terrorist incitement, child pornography, or overt criminal conspiracies. But some decisions to flag or remove posts have been either more contentious or simply erroneous, particularly since the platforms began to rely increasingly on artificial-intelligence (AI) systems to moderate content during the covid-19 pandemic. An even more central question concerns not what content social-media platforms remove, but rather what they display. From among the vast number of posts made on Twitter or Facebook, the content we actually see in our feeds is selected by complex AI algorithms that are designed primarily not to protect democratic values, but to maximize corporate revenues. It is thus unsurprising that these platforms have been blamed for propagating conspiracy theories, slander, and other toxic forms of viral content: This is what sells. Users do not know why they are seeing what they see on their feeds, or what they are *not* seeing because of the decisions of an invisible AI program.

## Harms

We thus need to be precise about the nature of the threat that the large platforms pose to modern liberal democracy. It does not lie in the mere fact that they carry "fake news" or conspiracy theories or other kinds of harmful political content. The U.S. First Amendment protects the right of citizens to say whatever they want, short of promoting violence or

sedition. Other democracies are less absolute in their free-speech protections, but nonetheless agree on the underlying principle that there should be an open marketplace of ideas in which the government should play a minimal role.

The real problem centers around the platforms' ability to either amplify or silence certain messages, and to do so at a scale that can alter major political outcomes. Any policy response should *not* aim at silencing speech deemed politically harmful. The notion that Donald Trump won the 2020 presidential vote in a landslide and that the Democrats stole the election through massive fraud is false and terribly damaging to U.S. democracy. But it is also sincerely believed by tens of millions of Americans, and it is neither normatively acceptable nor practically possible to prevent them from expressing opinions to this effect. For better or worse, people holding such views need to be persuaded, and not simply suppressed.

What policy needs to target instead is the dominant platforms' power to either amplify or silence certain voices in the political sphere. Up to now we have been relying on people such as Twitter's CEO Jack Dorsey or Facebook's Mark Zuckerberg to "do the right thing" and curate harmful political content. This is a response that may work in the short run, when the nation is faced with an imminent threat of political violence. But it is not a long-term solution to the underlying problem, which is one of excessively concentrated power.

No democracy can rely on the good intentions of particular powerholders. Numerous strands of modern democratic theory uphold the idea that political institutions need to check and limit arbitrary power regardless of who wields it. This principle is implicit in John Rawls's concept of the "veil of ignorance," according to which fair rules in a liberal society must be drawn up without regard to knowledge of the person or persons to whom they apply. The 1780 Constitution of the State of Massachusetts, drafted by John Adams, Samuel Adams, and James Bowdoin, stated that "the executive shall never exercise the legislative [or] judicial powers . . . to the end it may be a government of laws and not of men." James Madison's famous *Federalist* 51 lays the ground for a system of divided powers by arguing that "in framing a government which is to be administered by men over men, the great difficulty lies in this: you must first enable the government to control the governed; and in the next place oblige it to control itself." The only practical solution to this problem was to comprehend "in the society so many separate descriptions of citizens as will render an unjust combination of a majority of the whole very improbable, if not impracticable." In other words, power could be controlled only by dividing it, through a system of checks and balances.

The authors of these strictures were taking aim at state power, but their concerns apply doubly to concentrations of private power. Private power faces no checks comparable to popular elections; it can be controlled only by the government (through regulation) or by competition among power

holders. Due to a traditional suspicion of state power, market competition has generally been the preferred means of controlling and limiting private power in the United States. Fear of monopoly power's economic and political consequences, among other concerns, inspired passage of the legislation making up the backbone of U.S. antitrust law—the Sherman (1890), Clayton (1914), and Federal Trade Commission (1914) Acts.

## Remedies

How can we reduce the underlying power of today's internet platforms? I believe that a potential solution to this problem lies in using both technology and regulation to outsource content curation from the dominant platforms to a competitive layer of “middleware companies.” I advance this proposal not because I am certain that it will work, but because the alternative approaches that have been suggested are likely to be less effective.

The first and most obvious of these approaches is to use antitrust law to break up Facebook and Google, much as the telephone giant AT&T was broken up in the 1970s. After a prolonged period of lax enforcement of antitrust laws, there is a growing consensus that they need to be applied to the big tech companies, and suits have been brought against these platforms by the European Commission, the Justice Department, the Federal Trade Commission, and a coalition of state attorneys-general.

Breaking up these companies would indeed reduce their power over politics. But under current U.S. and EU laws, reaching a decision in the courts could take over a decade, as past antitrust cases against IBM and Microsoft did. More important, network externalities suggest that a baby Facebook emerging out of such a breakup could grow much faster than AT&T did when it was divided, and quickly reach the size of its parent. Antitrust law in any case is designed primarily to remedy the familiar harms stemming from concentrations of economic power, not the novel political risks produced by social media. What might realistically come out of current antitrust initiatives will be constraints on the platforms' acquisition of startups, or on their recourse to vertical-tying agreements (policies that compel users of a product offered by one of the tech giants to procure a related service from that same company). Yet outcomes of this kind will not address the political problems posed by platform scale.

A second obvious remedy is government regulation, something that both the EU and individual EU member states have already sought to put in place. Germany's NetzDG law, for example, imposes hefty fines on companies which fail to remove content that is illegal in that country within a day once it has been reported. There are precedents in the United States for government regulation of the content distributed by major media platforms. Back in the 1960s, when the television networks enjoyed an oligopolistic control over political discussion somewhat similar to the growing

dominance of today's social-media platforms, the Federal Communications Commission (FCC) used its licensing power to enforce the Fairness Doctrine, which required large media outlets to present competing points of view. The Fairness Doctrine's constitutionality was upheld in the 1969 Supreme Court decision in *Red Lion Broadcasting Co. v. FCC*, but was relentlessly attacked thereafter by Republicans who felt that the FCC was biased against conservatives. The Fairness Doctrine was rescinded in 1987 through an administrative decision by the FCC, and attempts by Democrats to restore it were unsuccessful. While some European democracies retain enough of a social consensus to muster a mandate for content regulation, the United States today is far too polarized to be able to authorize the FCC or any other government body to determine what is "fair and balanced" and enforce such strictures against the internet platforms. Regulation therefore seems to be a dead end in the United States at the present moment.

A third approach to reducing platform power that has been put forward is data portability. The idea is that individual users own their data and should be able to move it to alternate platforms, just as they transfer their mobile-phone numbers from one carrier to another. While this approach sounds like an appealing way to increase competition among platforms, it runs into immediate difficulties involving both property rights and technical feasibility. For the platforms' purposes, the most important data that they hold is not personal data voluntarily surrendered to them by users, but the mountains of metadata created by the users' interaction with their platforms. It is legally not clear who owns metadata, and the platforms will fight to keep control over such data since this is the bedrock of their business models. Moreover, these data are hugely heterogeneous and platform-specific. Data portability is therefore not a way of addressing the political threat that platform power poses.

Finally, some have suggested that platform power might be kept in check by applying privacy legislation to keep the platforms from using data collected in one sphere, such as book retailing, in another, such as selling groceries or diapers (something that Amazon has done), without getting explicit consent from users. Such restrictions are already built into Europe's General Data Protection Regulation (GDPR). Experience with that law, however, indicates that such rules are very hard to enforce; in any event, the United States does not have a privacy regime comparable to GDPR in place at the national level. Moreover, when it comes to the power of existing tech giants, the cat is already out of the bag, so to speak: Google and Facebook have already amassed huge databases on their users which privacy restrictions limiting future data collection would not touch.

## Middleware

Given the inadequacy of these various approaches, it is worth taking a closer look at the alternative remedy that the Stanford Working

Group on Platform Scale has labeled “middleware.” Middleware is software that rides on top of a platform and affects the way in which users interact with the data that the platform carries. A properly constructed

---

*Middleware could reduce the platforms’ power by taking away their ability to curate content, giving this function to a wide variety of competitive firms.*

---

middleware intermediary could, for example, filter platform content not just to label but to eliminate items deemed false or misleading, or could certify the accuracy of particular data sources. At one extreme, middleware could take over the entire user interface of a Facebook or Google, relegating those platforms to the status of “dumb pipes” that simply serve up raw data, much like the telephone companies. At the other extreme, middleware could operate with a light touch, labeling but otherwise not affect-

ing the content-curation decisions being made by the platforms. This would resemble steps that Twitter has already taken to label certain types of content deemed misleading, including election news in the runup to the November 2020 U.S. elections, but would allow users to choose from a broader menu of labeling options. There currently exist third-party services, such as NewsGuard, that plug into web browsers to offer users ratings of the credibility of news sources that they encounter. Middleware could perform a similar function while plugging directly into the social-media platforms. It could also transform the relationship between users and platforms in more fundamental ways.

Middleware could reduce the platforms’ power by taking away their ability to curate content, and outsourcing this function to a wide variety of competitive firms which in effect would provide filters that would be tailorable by individual users. When you signed up to Facebook or Google, you would be given a choice of middleware providers that would allow you to control your feed or searches, just like you now have a choice of browsers. In place of a nontransparent algorithm built into the platform, you could decide to use a filter provided by a nonprofit coalition of universities that would vouch for the reliability of data sources, or one that limited the display of products to those manufactured in the United States, or those that are environmentally friendly.

One of the likely objections to the middleware concept is that it will simply reinforce the “filter bubbles” that already exist on the platforms. Alt-right ideologues and conspiracy theorists could construct filters of their own that would keep out contrary views, leading to a further fragmentation of the political space. But as noted above, the objective of policy should not be to suppress harmful content. The latter, if it falls short of calling for violence, is constitutionally protected. In any event, it will be technologically very hard to eliminate such content. After the

January 6 attack on the U.S. Capitol, extremists began to move to the new platform Parler (which prided itself on a minimalist approach to moderation), and then, when Parler was temporarily offline after being dropped by Amazon's web-hosting service, to encrypted messaging services such as Telegram or Signal.

Much as we may regret this fact, hate speech and conspiracy theories are embedded in the broader society, and middleware will do little to stamp them out. But that is not a proper policy objective in a society that values free speech. What middleware might do instead is dramatically dilute the power of the platforms to amplify fringe views and take them mainstream. We might think of this in terms of an infectious-disease analogy: Instead of encouraging infected people to mingle in the broader society, we should seek to isolate them in spaces they share with the already infected.

Middleware will not spontaneously arise out of market forces. While there is demand for such services, there is no clear business model that will make them viable today. The platform owners may be happy to be relieved of responsibility for making controversial political decisions in their content moderation; in fact, Twitter's Jack Dorsey himself has recently suggested "giving more people choice around what relevance algorithms they're using," adding: "You can imagine a more market-driven and marketplace approach to algorithms."<sup>1</sup> On the other hand, big tech will not like the loss of control that middleware intermediation creates. This means that the creation of a vibrant and competitive middleware sector will depend on government regulation, both to establish rules for the application programming interfaces (APIs) by which such companies would plug into the platforms, and to set revenue-sharing mandates that will ensure a viable business model for middleware purveyors. These are all issues that need to be fleshed out in greater detail as we think through the consequences of the political crisis we have faced.

## Prospects

More and more people are coming to the realization that modern technology has created something of a monster, a communications system which bypasses the once-authoritative institutions that used to structure democratic discourse and provide citizens with a common base of factual knowledge over which they could deliberate. The private companies that are responsible for this outcome are now among the largest in the world. They possess not only enormous wealth which they can use to protect their interests, but also something of a chokehold over the communications channels that facilitate democratic politics. They benefit from economies of scale that are inherent in networked systems, and there is no easy way to prevent them from getting even larger. The covid-19 pandemic that struck the world in 2020 has vastly increased their power and importance.



Up to now, the large platforms have not seen it as in their interests to deliberately manipulate political outcomes or electoral results. Their commercial interests have, however, motivated them to privilege certain forms of viral content that more often than not are fake, conspiracy-laden, and harmful to democratic practice. What we should be worried about in terms of democratic health is the underlying power that these platforms possess. Public policy needs to be deployed to reduce that power, which otherwise might well one day come under the control of owners who *do* want to deliberately manipulate elections.

The objective of public policy should not be to control speech. Modern democracies abjured such control when they committed themselves to protecting freedom of expression. What we want, rather, are public policies that prevent private actors from using their power to artificially amplify or suppress certain types of speech, and that maintain a level playing field on which ideas can compete.

While much of the discussion here has focused on the United States and the current crisis in U.S. democracy, excessive platform power has worldwide repercussions. Facebook and Twitter are even more politically important in smaller countries around the globe, where they have become the major channel of public and private communication. In the wake of Twitter's de-platforming of Donald Trump, critics immediately asked why similar decisions were not being made to curtail the anti-democratic behavior of other politicians around the world, from elected populists to rulers in autocracies, who have used incendiary rhetoric online. In India, for example, Facebook has been singled out for its failure to take down posts decried for fomenting violence against Muslims.

It is clear that these giant U.S. companies do not have anywhere near the capacity to make nuanced political judgements about the acceptability of speech in the roughly 150 countries in which they operate. It is very hard to see what would give them the incentive to acquire such capacity in the future. More important, they do not have the legitimacy to control speech in their home country, the United States, much less in other countries around the world.

This is why the diminution of platform power is critical for the survival of democracy around the world. While Europeans have made efforts to curb platform power, Americans up to now have been complacent about the issue. Now that there is a general consensus that the large platforms pose a danger to U.S. democracy, it is vital to understand precisely where that threat lies, and what remedies are both politically and technologically realistic.

#### NOTE

1. "Edited Transcript of TWTR.N Earnings Conference Call or Presentation 9-Feb-21 11:00pm GMT," Reuters, 9 February 2021, [www.yahoo.com/now/edited-transcript-twtr-n-earnings-230000867.html](http://www.yahoo.com/now/edited-transcript-twtr-n-earnings-230000867.html).