Section 230: A Reform Agenda for the Next Administration

Matt Perault

October 26, 2020
Day One Project Working Paper*

Summary

While Donald Trump and Joe Biden disagree about many things, they agree on at least one: Section 230 of the Communications Decency Act should be repealed. As tech journalist Casey Newton has said, "both candidates want to end the internet as we know it."

President Trump has called for "immediately" repealing Section 230, which provides tech platforms (such as Facebook, Twitter, Reddit, and many others) with immunity from liability for speech by their users. The President has repeatedly spoken out on this issue, deeming it so important that he tweeted "REPEAL SECTION 230!!!" within 24 hours of returning to the White House after his COVID-19 hospitalization. In May 2020, President Trump issued an Executive Order on Preventing Online Censorship aimed at curbing tech platforms that "engage in deceptive or pretextual actions stifling free and open debate by censoring certain viewpoints." Pursuant to this Executive Order, Federal Communications Commission (FCC) Chairman Ajit Pai issued a statement in mid-October announcing that he would initiate rulemaking proceedings to "clarify" the meaning of Section 230. The Trump Justice Department has also been active on Section 230 reform, including by sending proposed legislation to Congress after hosting a workshop on the topic as well as by publishing a set of recommendations for reform.

While former Vice President Biden argued in May that the President's Executive Order was problematic due to First Amendment concerns, he has also expressed deep concerns with Section 230. In January 2020, the Vice President called for Section 230 to be "revoked." Biden has emphasized that his focus is on moderating more content rather than less, and supports eliminating protections for tech platforms that knowingly host false content. Despite this difference between the two candidates' goals for narrowing the scope of Section 230 protections, they are charting a similar policy path for the next administration.

As the election has approached, proposals to reform Section 230 have been published almost weekly by legislators, academics, and think tanks. Multiple controversies have drawn additional attention to the topic. Even Supreme Court Justice Clarence Thomas urged the Court to take a case to determine the "correct" meaning of Section 230.

Across the political spectrum, calls for Section 230 reform have been rooted in the shared objectives of reducing the prevalence of problematic online content and improving platforms' content-moderation decisions. But what is considered problematic varies dramatically. Some criticisms of the status quo are focused on bias in platforms' content-moderation decisions, and others on the harassment of women and people of color. Some stakeholders target misinformation generally, while others are principally concerned with foreign election interference and voter suppression. Some people want platforms to be more aggressive in removing speech, while others want platforms to be more protective of it.

Conflicting motivations for reform have made it difficult to reach consensus on the path forward. A desire to score political points in an election year has given momentum to sweeping reforms—such as the total revocation of Section 230—that would radically alter the nature of online expression. Unfortunately, this focus on headline-grabbing reform has resulted in insufficient attention to more targeted changes that might incrementally but meaningfully improve the status quo.

The next administration needs a different blueprint for reform. This paper argues that the next administration should prioritize targeted changes that will deter some of the most harmful forms of speech while also preserving the features of tech platforms that are essential to online expression. Specifically, the next administration should work with Congress, industry, civil society, and academia to implement the following five reforms:

1. **Modernize federal criminal law for the digital age.** Congress should modernize federal criminal law by prohibiting some of the most egregious forms of online speech, such as intentional voter suppression on digital platforms. On Day One of the next administration, the Department of Justice should create a task force to develop a framework for this legislation.
2. **Require platforms to comply with court orders to remove illegal content.** Congress should pass a modified version of the [Platform Accountability and Consumer Transparency (PACT) Act](#), removing Section 230 as a defense if platforms refuse to remove content after a judge has ordered it to be removed.
3. **Define what it means for a platform to "develop" content.** By June 2021, the Federal Trade Commission (FTC) should hold three workshops to gather data from users, platforms, and experts to better understand how to delineate between content development and content hosting. Workshop findings should form the basis for non-binding FTC guidance for judges and platforms that outlines criteria for determining when platforms move from merely hosting content to actively developing it.
4. **Design products to facilitate individual accountability.** Platforms should make it easier for people to report problematic content so that offenders can be held accountable, where appropriate. The Department of Justice should convene state attorneys general, election-protection experts, specialists in defamation law from the American Bar Association, and industry representatives for a workshop on options for improving reporting flows on tech platforms.
5. **Share data to inform future policymaking.** Congress should pass legislation providing a safe harbor to companies that share data with researchers consistent with privacy best practices. Alternatively, the FTC and the Department of Justice could issue a notice of enforcement discretion, signaling that they will not pursue enforcement action against platforms that share data with researchers consistent with certain best practices. Finally, researchers and policymakers should work with platforms to identify data that they should publish in their transparency reports that would help to measure progress.

Challenge and Opportunity

Section 230 reform has been an appealing target for government action because of the increasing bipartisan frustration about the existence of problematic online content and the content-moderation decisions made by platforms. It has also been an appealing target because of rampant misunderstandings about the law. Many people [confuse](#) Section 230 [protections](#) with the [protections](#) of the First Amendment. [Others inaccurately](#) assert that Section 230 protections are conditioned on platform neutrality.

Alongside the proposals from President Trump and former Vice President Biden, federal legislators have introduced a flurry of competing reform proposals. Two bipartisan bills have received significant attention: the Eliminating Abusive and Rampant Neglect of Interactive Technologies Act (EARN IT Act), introduced by Senators Graham, Blumenthal, Hawley, and Feinstein, and the Platform Accountability and Consumer Transparency Act (PACT Act), introduced by Senators Schatz and Thune. Senator Hawley has introduced numerous other proposals, including one that would make Section 230 protection conditional upon platforms demonstrating that they are politically neutral. In the last two months alone, two other bills have been introduced in the Senate (the Online Freedom and Viewpoint Diversity Act, introduced by Senators Wicker, Graham, and Blackburn, and the See Something, Say Something Act, introduced by Senators Manchin and Cornyn) and one in the House (the Protecting Americans from Dangerous Algorithms Act, introduced by Representatives Eshoo and Malinowski).

Many academics and think tanks have introduced reform proposals as well. Benjamin Wittes and Danielle Citron proposed reforming Section 230 so that its immunity would be conditional on a platform taking "reasonable steps to prevent or address unlawful uses of its services." Citron also authored a paper with Mary Anne Franks that described adverse impacts of Section 230 on victims of online abuse. The Citron/Franks paper proposed a clearer delineation between online speech and online conduct, and also proposed excluding "Bad Samaritans" from Section 230 protections. Daphne Keller has written extensively on Section 230, both detailing the current state of the law to separate myth from reality and outlining a series of "doctrinal knobs and dials" that reformers can use in crafting a new approach to liability protections. TechFreedom founder Berin Szoka has proposed a First Amendment test for Section 230 reforms, encouraging lawmakers to ensure that reform proposals are as minimally restrictive and "narrowly tailored" as possible while still achieving desired outcomes. Other scholars who have written extensively on Section 230 reform include Eric Goldman at Santa Clara University and Jeff Kosseff, author of *The Twenty Six Words that Created the Internet*.

Two academic research centers recently published notable reform proposals. A report from the Shorenstein Center at Harvard advocated for a Digital Platform Agency that would enforce a "duty of care" for platforms, while a report from the NYU Stern Center for Business and Human Rights called for a "quid pro quo" approach to liability protection where immunity would be contingent upon compliance with certain new responsibilities. Both concepts were discussed in previous reports by the Stigler Center at the University of Chicago's Booth School of Business. The "quid pro quo" idea was previously proposed by Ellen Goodman, a law professor at Rutgers.

Industries that compete with tech platforms have also joined the reform fight. News publishers, who have increased their policy advocacy efforts in recent years to try to reclaim a portion of the advertising pie from tech platforms, have argued that Section 230 is a "license to build rage machines," and that tech platforms are filled with "made-up garbage." News Corp has claimed that "[p]ower without responsibility, without accountability, is a recipe for disaster," and has advocated for the creation of an Algorithm Review Board. AT&T lobbied for changes that would "reduce the gross disparities in legal treatment that have emerged between the dominant online platforms and the traditional purveyors of third-party content," claiming that platform neutrality deserves more attention from policymakers than net neutrality. Oracle, Disney, 20th Century Fox, and IBM have also advocated for weakening Section 230 protections. Oracle has become so associated with lobbying efforts against Section 230 that in the wake of its recently announced technology

partnership with TikTok, some commentators speculated about whether the deal would cause the company to [revisit](#) its position.

<u>Obstacles to reform</u>

Despite widespread concern with Section 230 and the numerous proposals for changing it, the road to reform is paved with obstacles. First, previous Section 230 reforms have had unintended consequences, including harming the very communities those reforms aimed to protect. For instance, the 2018 [SESTA/FOSTA](#) legislation narrowed Section 230 protections with the goal of curbing online trafficking and protecting sex workers. But the sex-worker community has [claimed](#) that it had the opposite [effect](#), [increasing](#) violence toward sex workers and shuttering online forums that helped workers to identify abusive clients.

Second, any reform that significantly increases litigation costs for tech platforms will result in product changes that constrain user expression. If platforms become liable for a wider range of content posted by users, they will be incentivized to limit the ability of users to post any content that might generate that liability. Reforms could have this effect even without dramatically changing the likelihood that platforms are ultimately held liable in court. If platforms need to devote more legal resources to proving that they behaved responsibly, for instance, then they will change the design of their products even if they would likely prevail in most cases. The result is likely to be more friction in sharing and more friction in advertising, with platforms playing a gatekeeper role similar to the role played by broadcasters and publishers. Some people might be happy with a version of YouTube that feels like Netflix or a version of Facebook that feels like a newspaper, in which a user can't share content unless that content is explicitly approved by the platform. But many would not.

Third, increasing the costs of platform moderation might reduce the competitiveness of the tech sector. Section 230 makes it possible for startups to build and grow a user-generated content product, even if they're not able to allocate immense resources to content moderation or legal fees. Large platforms like Google, Facebook, and Twitter can afford to hire tens of thousands of people to moderate content and to build tools (such as artificial-intelligence classifiers) that make content moderation more efficient. Smaller companies often cannot. As a result, expanding liability for platforms will [likely](#) make the tech sector less competitive overall.

Fourth, the tech industry isn't the only industry that will be affected by Section 230 reform. It is a common misunderstanding that only large tech platforms benefit from the protections of Section 230. In reality, the provision provides immunity to [any entity](#) that hosts content online. For instance, even though News Corp has advocated against Section 230, it benefits from the protections when it provides [comments sections](#) on its websites. The repercussions of broad Section 230 reform could be more disruptive to non-tech industries than is currently anticipated.

Fifth and finally, there is still tremendous uncertainty about the scope and scale of the problem that advocates for Section 230 reform aim to address. Research has found that the impact of problematic online content may be relatively small. One study reviewed 49 field experiments on the impact of political advertising on voter preferences and [found](#) that "the best estimate of the effects...is zero." Other research has shown that exposure to misinformation is relatively [limited](#), that a very small percentage of platforms' user bases are [responsible](#) for spreading misinformation, that sharing fake news was "[quite rare](#)" in the 2016 election, and that interacting with accounts run by the Russian Internet Research Agency had [limited effect](#). Studies have also called into question the [prevalence](#) and impact of "filter bubbles," [finding](#) there is a "substantial amount

of overlap (51%) in the ideological distributions of accounts followed by users on opposite ends of the political spectrum." Other [research](#) suggests that the "second-order effects" of media narratives about misinformation may have a more significant negative impact than the misinformation itself. A recent study by the Berkman Klein Center at Harvard [found](#) that voting disinformation was "an elite-driven, mass-media led process" and that "social media played only a secondary and supportive role."

More research is needed to better understand the role of online content and platform decision-making in persuading and mobilizing platform users for or against various causes, as well as the effects of online content and platform decision-making on our democratic institutions. Yet research to date suggests that the problems may be narrower in scope and scale than many people think. If so, then targeted solutions are likely preferable to broad reforms that could radically disrupt the internet ecosystem and adversely affect content hosts, creators, and consumers.

Proposed Actions

On Day One, the next administration should work with Congress, industry, civil society, and academia to implement five reforms to reduce some of the most harmful online speech, clarify the scope of platform liability without significantly increasing litigation costs, facilitate individual accountability, and provide data to inform product and policy design. The administration should establish a task force within the White House Office of Science & Technology Policy charged with overseeing these reforms. The task force should include experts with the necessary experience to guide engagement across the executive branch, Congress, industry, civil society, and academia.

Reform 1. Modernize federal criminal law for the digital age

On Day One, the next administration should take steps to modernize federal criminal law for the digital age. Because Section 230 does not provide a defense against violations of federal criminal law, passing new federal criminal law will enable Congress to have a meaningful effect on the scope of Section 230 protections without altering the text of Section 230 itself. Specifically, legislators should pass new federal criminal laws that prohibit some of the most harmful forms of online speech, such as voter suppression and incitement to riot.

Current legislative landscape

Federal criminal law has not kept pace with technological change. Using federal criminal law to moderate online speech is an approach that has been underutilized because the First Amendment restricts Congress's power to limit expression. But the First Amendment does not prohibit speech restrictions entirely. The Supreme Court has [held](#) that Congress may pass laws to restrain inflammatory speech that is "directed to inciting or producing imminent lawless action and is likely to incite or produce such action." Congress [may](#) also restrict the time, place, and manner of speech so long as those restrictions are content-neutral, narrowly tailored to serve a significant governmental interest, and leave open alternatives for communication. The Court has also [upheld](#) the constitutionality of laws restricting election-related speech when those laws are needed to "protect[] voters from confusion and undue influence" and to "ensur[e] that an individual's right to vote is not undermined by fraud in the election process."

Consistent with First Amendment jurisprudence, existing law already bars harmful speech in very specific circumstances that are likely to result in harm, such as [organizing a riot](#). Federal criminal law also prohibits certain types of voter intimidation, and has been used in [cases](#) where phone lines were jammed to prevent people from coordinating transportation to the polls. In 2017, the Justice Department [stated](#) its interpretation

of how federal civil rights law might be applied in the context of voting: "It is the Criminal Division's position that Sections 241 and 242 may be used to prosecute schemes to intimidate voters in federal elections through threats of physical or economic duress, or to prevent otherwise lawfully qualified voters from getting to the polls in elections when a federal candidate is on the ballot." Federal law also prohibits certain types of voter fraud.

But gaps remain. While many states prohibit deceptive practices in voting, no equivalent federal law exists. The Justice Department has been explicit about the existence of this gap:

> *"Voter suppression schemes are designed to ensure the election of a favored candidate by blocking or impeding voters believed to oppose that candidate from getting to the polls to cast their ballots. Examples include providing false information to the public – or a particular segment of the public – regarding the qualifications to vote, the consequences of voting in connection with citizenship status, the dates or qualifications for absentee voting, the date of an election, the hours for voting, or the correct voting precinct. Currently there is no federal criminal statute that expressly prohibits this sort of voter suppression activity."*

In the absence of federal law in this area, organizations have called on social-media companies to combat voting misinformation, and have threatened to take action if those companies are not sufficiently responsive. But at a time when many policymakers have expressed concerns about the power of tech companies, platforms should not be given the primary power and responsibility for policing voter suppression. Instead, policymakers should fill the gap in federal law. One option would be to reform Section 230 to permit platforms to be held liable under existing state law on voting practices. However, interstate use of online tools to engage in voter suppression is precisely the type of national issue where federal leadership and standardization is warranted, and Congress has the authority to act.

Proposed legislation

Specifically, Congress should pass legislation prohibiting deceptive practices related to voting, such as intentional efforts to perpetuate election fraud or to mislead voters about voting time and location. This legislation should explicitly prohibit the use of online platforms to engage in these deceptive practices. Congress might also consider including broader restrictions on other problematic speech that receives less protection in U.S. law, such as misinformation campaigns by foreign actors or incitements to riot. Such restrictions might address bipartisan concerns about the use of online speech to organize violence against individuals and businesses, and should be narrowly constructed to focus on likely imminent harm. Online incitement to riot would be a particularly easy gap to address. While it is currently unlawful to use "the mail, telegraph, telephone, radio, or television" with intent to incite, organize, or aid and abet a riot, federal criminal law does not explicitly apply to online communications.

To survive a First Amendment challenge, the government would need to show that such legislation was "necessary to serve a compelling state interest and that it is narrowly drawn to achieve that end." This is a high hurdle, but not an impossible one. The Supreme Court has upheld some speech-restrictive voter protection laws. Several state laws that prohibit deceptive practices in voting have not yet been struck down on First Amendment grounds. For instance, Virginia prohibits communication of "false information" related to the "the date, time, and place of the election or the voter's precinct, polling place, or voter registration status."

If passed and upheld, federal criminal law prohibiting deceptive practices related to voting would enable law-enforcement officials to prosecute bad actors, deter people from using online platforms to engage in voter suppression, and create liability for platforms that intentionally engage in suppressive speech directly or that aid and abet individuals who do. Even when platforms do not directly violate this new criminal standard, the legislation may compel platforms to turn over user data that can be used to identify individuals who have violated the law and assist in their prosecution. This is already the case for other types of law-enforcement [requests](). Existence of a criminal standard might also encourage platforms to provide user information [voluntarily]() when they see evidence of illegal activity. If the goal is to reduce the amount of voter misinformation on platforms, then expanded legal liability for voter misinformation would likely produce this result.

Task force recommendations

To help craft a law that survives First Amendment review, the next administration should create a task force to review existing state law on voter suppression and to develop recommendations for federal legislation that would expand criminal liability for misconduct—including misconduct through online platforms. The task force should focus on developing legislation that would address three aspects of voter suppression: (1) voter intimidation and harassment; (2) intentional dissemination of inaccurate information about voting dates, voter eligibility, and the voting process; and (3) foreign interference with voting.

Since its focus would be on criminal law, the task force should be led by the Department of Justice and should include state attorneys general and representatives from the Federal Trade Commission. To tackle foreign interference, the task force should also include national-security experts, including representatives from the National Security Council and the Defense Department. To tackle election issues, the task force should include representatives of the Federal Elections Commission. The task force should also consult broadly with non-government stakeholders, including nonprofit organizations focused on voter fraud, civil rights organizations, academic experts, and industry. The task force should issue recommendations for legislation no later than January 20, 2022—one year from the first day of the next administration—to increase the likelihood that legislation is passed before the 2022 midterm elections.

Reform 2. Require platforms to comply with court orders to remove illegal content

Currently, tech platforms are rarely obligated to remove content that is found to be illegal by a federal judge. Although most platforms have terms of service that [prohibit]() use of the platform to engage in illegal activity and thus often remove illegal content on their own, they do not typically have a legal obligation to do so. This current best practice should be codified in law.

Congress should pass legislation requiring tech platforms to remove content that a court has deemed illegal, and explicitly stating that Section 230 cannot be used as a defense when platforms host that content. The PACT Act provides a helpful model for such legislation. Although the PACT Act might be challenged on First Amendment [grounds](), its provisions on removal of content pursuant to court orders offer a targeted approach to online misinformation. Importantly, the relevant PACT Act provisions do not change the scope of underlying legal liability, since a judge must find the content to be illegal under existing law before a platform would be obligated to remove it. Moreover, the PACT Act doesn't condition Section 230 protections on an unrelated policy objective like [encryption](), and doesn't impose amorphous obligations of "reasonableness" or a duty of care that would dramatically increase litigation costs for platforms.

Although PACT Act provisions governing content removal are promising, they could be improved through four changes. First, platforms should be required to reinstate content if the judge's decision is reversed on appeal. Second, the legislation should not require platforms to remove content within 24 hours, since that timeline may not be feasible for all platforms. Third, removal orders should apply only to existing content and accounts that can be identified through a specific URL or location, and should not be used to impose ongoing content monitoring and filtering obligations. Fourth, the orders should be directed to the speaker, rather than the platform, so as to avoid dragging platforms into thousands of court disputes. To ensure that these speaker-directed orders are effective, platforms should create a submission mechanism for complainants to submit court orders requiring content removal. Platforms, legislators, and courts should work together to develop mechanisms to validate those orders. These four changes will strengthen the PACT Act, and may help build stronger bipartisan support for it.

Reform 3. Define what it means for a platform to "develop" content

The applicability of Section 230 liability protections depends on whether a platform is "hosting" or "creating" content. If the platform is "responsible, in whole or in part, for the creation or development of information," then it is an "information content provider" and cannot use Section 230 as a defense against liability. If it is merely hosting content—providing "computer access by multiple users to a computer server"—then it is an "interactive computer service" that may plead Section 230 as a defense.

Despite the importance of this dividing line between content hosting and content creation, the law is remarkably vague on the difference between the two. The definitions in Section 230 provide little guidance on what it means for a platform to "develop" content "in part," and there have been few court decisions on the question. The Roommates.com decision by the Ninth Circuit in 2008 is the leading case in this arena. That decision held that a "website helps to develop unlawful content, and thus falls within the exception to section 230, if it contributes materially to the alleged illegality of the conduct."

But jurisprudence has not kept pace with emerging questions about the role of platforms in developing content. Could a platform's algorithm-design decisions ever play such significant role in content creation that a platform could be held liable for "development in part"? In advertising, could a platform's targeting controls rise to the level of "development" if they play a meaningful role in shaping the content of an ad, such as by incentivizing political advertisers to run more polarizing campaigns?

One option for resolving these ambiguities is for Congress to add language to the definitions of "interactive computer service" and "information content provider" to clarify what it means to "develop" content "in part." Using legislation to provide this clarity would provide significantly more certainty for platforms and for users, but would also likely take a long time to weave its way through the political process, and would be difficult to adapt over time to evolving social and technological developments. These pitfalls make legislative reform a fraught path.

A second option is to lean entirely on courts to develop relevant standards over time. But relying on decisions in individual cases may also mean that the standards will evolve too slowly to keep pace with technology and are problematically influenced by the arbitrariness of which cases end up making it through the court system. Judges may also not be well-suited to lead the development of these norms because they often lack the requisite technical expertise.

The third and preferred option is for the FTC to lead in gathering information that will help distinguish between content hosting and content creation. By June 2021, the FTC should hold three workshops to gather input from users, platforms, and experts on what it means for a platform to "develop" content. Platforms should share information about the design choices they face, and how these decisions shape user expression. They should provide information about a wide range of products and design elements, from user posts to advertising to algorithmic preferencing.

Workshop findings should be used for the basis of non-binding FTC guidance for judges and platforms on when platforms move from merely hosting content to actively developing it. The guidance might be similar to the advisory opinions that the FTC regularly issues. The FTC should issue this guidance no later than January 20, 2022. It should revisit the guidance annually to consider whether revisions are necessary to accommodate new product developments. Over time, collected FTC guidance might facilitate better jurisprudence on content development—jurisprudence that responds appropriately to changing product realities, and that helps platforms better understand their potential liability when designing products and services.

Reform 4. Recommend ways for platforms to facilitate individual accountability

While Section 230 immunizes platforms for certain speech by its users, it doesn't immunize users. In many areas, such as copyright, platforms have designed reporting flows to facilitate private resolution of disputes. In others, platforms recommend that users report behavior to law-enforcement officials who may then use the information as a basis for investigation and prosecution. Platforms also sometimes facilitate reporting to trusted members of a community, such as in bullying cases where a potential victim might prefer to report an incident to a teacher instead of to a platform or to law enforcement.

Similar design features might make it easier to hold people accountable for unlawful online activity that causes real-world harm. As discussed above, many states prohibit deceptive practices in voting. Platforms could make it easier to hold violators accountable by providing functionality that enables people to report such content not only to the platform, but also to the offices of state attorneys general. These offices could then review the content in question and determine whether prosecution is warranted. To enable resolution of these disputes (e.g., to identify posters), platforms may also need to retain certain data, such as basic subscriber information. Platforms might also provide alternative options for reporting voting misinformation for those who prefer not to involve law enforcement. For instance, platforms could provide options to report false information to an election-monitoring organization or to contact a nonpartisan voting-information organization for accurate information about voting times and locations. Platforms could also provide functionalities that assist with reporting and enforcement for other areas of state civil and criminal law, such as options for reporting defamation to lawyers who specialize in defamation law.

To help platforms design reporting mechanisms and workflows, the Department of Justice should convene state attorneys general, election protection experts, specialists in defamation law from the American Bar Association, and industry representatives for a workshop on reporting options that will be most likely to address illegal conduct and, where appropriate, to facilitate private resolution of complaints.

Reform 5. Share data to inform future policymaking

Like many other debates in tech policy, the debate about Section 230 has presumed that we know more about the costs and benefits of certain policy solutions than we do. This false sense of certainty stems in part from constant reporting on "numerators": the existence of individual examples of problematic content or estimates of the total amount of problematic content. These examples are often harrowing, since people use

online communication tools to share deeply disturbing text, images, and video. And the total amount of problematic content is staggering, since the scale of user-generated content platforms is immense.

But contextualizing harms necessitates the inclusion of denominators alongside numerators. How pervasive is problematic content within the entire online ecosystem? How much problematic content is removed relative to the content that remains? What is the volume of content that empowers vulnerable communities compared to the volume of content that harms and harasses? What are the respective removal rates for liberal and conservative content—is there systemic partisan bias? Have platforms gotten better over time at identifying and removing problematic content? Has their error rate in mistakenly removing acceptable content improved? These critical questions are difficult to answer.

It's also difficult to quantify the impact of potential reforms. If platforms were more liable for problematic content, how much valuable speech would be censored? If platforms played more of a gatekeeper role (like newspapers and TV networks do), who would be locked out? Would reforms designed to protect vulnerable communities also compromise their ability to speak and organize?

Accurately weighing the tradeoffs between freedom of online expression and protections against harmful expression is essential to informed policymaking. Overstating the harms of the status quo may result in more support for policy frameworks that radically depart from our current approach, even if they have a detrimental impact on online speech and safety. Systematically overvaluing the benefits of certain reforms and undervaluing their costs may lead to solutions that have a negative effect overall. Conversely, understating the harms may decrease willingness to consider robust reform options that would improve user experiences on tech platforms and reduce social harm.

As such, the next administration should ensure that platforms share data that enables academics and other experts to assess the effects of different policy options.

Consider Twitter's decision to place public-interest warnings against glorifying violence on President Trump's tweets about shooting "looters" in Minneapolis. If our societal goal is to minimize the number of people who see misinformation, who believe it, and who engage in harmful real-world behavior because of it, then we should evaluate Twitter's decision based on two factors. First, did Twitter's warnings have educational value for people who saw the President's tweets. Second, did Twitter's actions reduce the dissemination of that content?

Twitter did not make data available that would enable researchers to evaluate these questions. As a result, it's difficult to say whether Twitter's decision was the right one and whether other platforms should have followed its lead. Instead of a debate rooted in facts about educational value and virality, the public conversation about Twitter's approach was largely based on conjecture, generalities about the value of free speech, and disgust at the substance of President Trump's comments.

One barrier to increased data sharing is the risk of data misuse and the likelihood that platforms will bear responsibility for this risk. After the FTC levied a $5 billion fine against Facebook for making data available to a Cambridge researcher in the Cambridge Analytica scandal, platforms have taken steps to reduce third-party data access, determining that the value of data sharing for researchers is outweighed by the risks of data sharing for platforms.

To address this barrier, policymakers should provide [safe harbors](#) for sharing aggregated, anonymized data to help experts study the impact of content policies. The legislation should condition the safe harbor on platform compliance with a set of best practices for responsible data sharing. While legislation codifying this safe harbor would be the most enduring approach to incentivizing increased data sharing, an alternative option is for the Department of Justice or the FTC to issue a [notice](#) of enforcement discretion on data sharing for research. The notice would state that neither agency would pursue enforcement action against platforms that share data with researchers on the impacts of their content policies, provided that they share this data consistent with a set of privacy best practices. The FTC could develop these best practices in consultation with researchers, industry, and privacy professionals.

Finally, researchers and policymakers should work with platforms to identify data that they should publish data in their transparency reports to enable researchers, the [Office of Information and Regulatory Affairs](#), and other stakeholders to measure progress and evaluate whether the benefits of potential Section 230 reforms would outweigh their costs. For instance, platforms could publish annual data on removals mandated by PACT Act provisions (Reform 2 recommended in this paper) and on reports of problematic content received through new reporting flows (Reform 4). Attorneys general could publish data on prosecutions initiated pursuant to these new reporting flows. The Department of Justice could work with election-protection experts to provide data on prosecutions under new federal criminal law on voter suppression (Reform 1). Each of these datasets would help stakeholders evaluate the successes (and unintended consequences) of the reforms and identify areas for future policymaking.

Conclusion

As debates over the future of Section 230 rage on, policymakers should consider targeted reforms that will deter certain problematic types of online expression, such as voter suppression. They should also consider reforms that require platforms to comply with court orders to remove illegal content, provide guidance on what constitutes content development, facilitate individual accountability, and provide data to inform future decision making. While these reforms will not address all problematic online content, they will put the next administration on a path toward improving the state of online discourse without imposing significant new constraints on expression.

Frequently Asked Questions

What are the implications of defining when platforms are considered to have "developed" content?
Platforms cannot plead Section 230 as a defense if they create or develop content, even if they only develop it "in part." A clearer definition of the line between content hosting and content development will clarify when platforms are able to use Section 230 as a defense and when they cannot.

What obstacles do you expect the next administration to encounter with Congress, industry, academia, and civil society in the effort to push forward these issues?
Section 230 has been a polarizing issue, and stakeholders on all sides hold strong views on the right path forward. The reforms proposed in this working paper are aimed at building bipartisan consensus on measures that will help to mitigate some of the most pernicious online speech, while preserving the key components of the user-generated content product model that have been so essential to online expression.

About the Author

Matt Perault is Director of the Center on Science & Technology at Duke University and Associate Professor of the Practice at Duke's Sanford School of Public Policy. He previously served as a director of public policy at Facebook, where he led the company's global public policy planning efforts on issues such as competition, law enforcement, and human rights and oversaw public policy for WhatsApp, Oculus, and Facebook Artificial Intelligence Research. Prior to joining Facebook, Matt was Counsel at the Congressional Oversight Panel. He previously worked as a consultant at the World Bank and served as a law clerk for the U.S. District Court for the District of Columbia. Matt holds a law degree from Harvard Law School, a Master's degree in Public Policy from Duke University's Sanford School of Public Policy, and a Bachelor's degree in political science from Brown University.

About the Day One Project

The Day One Project is dedicated to democratizing the policymaking process by working with new and expert voices across the science and technology community, helping to develop actionable policies that can improve the lives of all Americans, and readying them for Day One of a future presidential term. For more about the Day One Project, visit dayoneproject.org.