

CHARTING A WAY FORWARD

Online Content Regulation

Monika Bickert

VP, CONTENT POLICY

FACEBOOK

Table of Contents

03 I. A New Governance Challenge

06 II. Four Key Questions

09 QUESTION 1

How can content regulation best achieve the goal of reducing harmful speech while preserving free expression?

10 QUESTION 2

How should regulation enhance the accountability of internet platforms to the public?

13 QUESTION 3

Should regulation require internet companies to meet certain performance targets?

16 QUESTION 4

Should regulation define which “harmful content” should be prohibited on internet platforms?

19 III. Principles for Future Regulators

21 IV. The Beginning of a Conversation

22 End Notes

A New Governance Challenge



Billions of people have come online in the past decade, gaining unprecedented access to information and the ability to share their ideas and opinions with a wide audience. Global connectivity has improved economies, grown businesses, reunited families, raised money for charity, and helped bring about political change. At the same time, the internet has also made it easier to find and spread content that could contribute to harm, like terror propaganda. While this type of speech is not new, society's increased reliance on the internet has shifted the power among those who control and are affected by mass communication.

For centuries, political leaders, philosophers, and activists have wrestled with the question of how and when governments should place limits on freedom of expression to protect people from content that can contribute to harm. Increasingly, privately run internet platforms are making these determinations, as more speech flows through their systems. Consistent with human rights norms, internet platforms generally respect the laws of the countries in which they operate, and they are also free to establish their own rules about permissible expression, which are often more restrictive than laws. As a result, internet companies make calls every day that influence who has the ability to speak and what content can be shared on their platform.

Problems arise when people do not understand the decisions that are being made or feel powerless when those decisions impact their own speech, behavior, or experience. People may expect similar due process channels to those that they

enjoy elsewhere in modern society. For example, people in liberal democracies may expect platforms to draft content standards with user and civil society input in mind, as well as the ability to seek redress if they feel a decision was made in error. People are able to register preferences through market signals by moving to other platforms, but they may not be satisfied by this as their only option.

As a result, private internet platforms are facing increasing questions about how accountable and responsible they are for the decisions they make. They hear from users who want the companies to reduce abuse but not infringe upon freedom of expression. They hear from governments, who want companies to remove not only illegal content but also legal content that may contribute to harm, but make sure that they are not biased in their adoption or application of rules. Faced with concerns about bullying or child sexual exploitation or terrorist recruitment, for instance, governments may find themselves doubting the effectiveness of a company's efforts to combat such abuse or, worse, unable to satisfactorily determine what those efforts are.

This tension has led some governments to pursue regulation over the speech allowed online and companies' practices for content moderation. The approaches under consideration depend on each country's unique legal and historical traditions. In the United States, for example, the First Amendment protects a citizen's ability to engage in dialogue online without government interference except in the narrowest of circumstances. Citizens in other countries often have different expectations about freedom of expression, and governments have different expectations about platform accountability. Unfortunately, some of the laws passed so far do not always strike the appropriate balance between speech and harm, unintentionally pushing platforms to err too much on the side of removing content.

Facebook has therefore joined the call for new regulatory frameworks for online content—frameworks that ensure companies are making decisions about online speech in a way that minimizes harm but also respects the fundamental right to free expression. This balance is necessary to protect the open internet, which is increasingly threatened—even walled off—by some regimes.

Facebook wants to be a constructive partner to governments as they weigh the most effective, democratic, and workable approaches to address online content governance. As Mark Zuckerberg [wrote in a recent op-ed](#):

It's impossible to remove all harmful content from the Internet, but when people use dozens of different sharing services—all with their own policies and processes—we need a more standardized approach ... Regulation could set baselines for what's prohibited and require companies to build systems for keeping harmful content to a bare minimum.¹

This paper explores possible regulatory structures for content governance outside the United States and identifies questions that require further discussion. It builds off recent developments on this topic, including legislation proposed or passed into law by governments, as well as [scholarship](#) that explains the various content governance approaches that have been adopted in the past and may be taken in the future.² Its overall goal is to help frame a path forward—taking into consideration the views not only of policymakers and private companies, but also civil society and the people who use Facebook’s platform and services.

This debate will be central to shaping the character of the internet for decades to come. If designed well, new frameworks for regulating harmful content can contribute to the internet’s continued success by articulating clear, predictable, and balanced ways for government, companies, and civil society to share responsibilities and work together. Designed poorly, these efforts may stifle expression, slow innovation, and create the wrong incentives for platforms.

Four Key Questions

02

Since the invention of the printing press, developments in communications technology have always been met with calls for state action. In the case of internet companies, however, it is uniquely challenging to develop regulation to ensure accountability. There are four primary challenges.

1. Legal environments and speech norms vary.

Many internet companies have a global user base with a broad spectrum of expectations about what expression should be permitted online and how decision makers should be held accountable. The cross-border nature of communication is also a defining feature of many internet platforms, so companies generally maintain one set of global policies rather than country-specific policies that would interfere with that experience.

2. Technology and speech are dynamic.

Internet services are varied and ever changing. Some are focused on video or images, while some are primarily text. Some types of internet communications are one-to-one and analogous to the private sphere, while others are more like a town square or broadcasting—where thousands or millions of people can access the content in question. Some services are ephemeral, and some are permanent. Among these different interaction types, norms for acceptable speech may vary. Just as people may say things in the privacy of a family dinner conversation that they would not say at a town hall meeting, online communities cultivate their own norms that are enforced both formally and informally. All are constantly changing to compete and succeed.

3. Enforcement will always be imperfect.

Given the dynamic nature and scale of online speech, the limits of enforcement technology, and the different expectations that people have over their privacy and their experiences online, internet companies' enforcement of content standards will always be imperfect. Their ability to review speech for policy or legal violations will always be limited. Even in a hypothetical world where enforcement efforts could perfectly track language trends and identify likely policy violations, companies would still struggle to apply policies because they lack context that is often necessary, such as whether an exchange between two people is in jest or in anger, or whether graphic content is posted for sadistic pleasure or to call attention to an atrocity.

4. Companies are intermediaries, not speakers.

Publisher liability laws that punish the publication of illegal speech are unsuitable for the internet landscape. Despite their best efforts to thwart the spread of harmful content, internet platforms are intermediaries, not the speakers, of such speech, and it would be impractical and harmful to require internet platforms to approve each post before allowing it.

The ability of individuals to communicate without journalistic intervention has been central to the internet's growth and positive impact. Imposing traditional publishing liability would seriously curtail that impact by limiting the ability of individuals to speak. Companies seeking assurance that their users' speech is lawful would need to review and approve each post before allowing it to appear on the site. Companies that could afford to offer services under such a model would be incentivized to restrict service to a limited set of users and err on the side of removing any speech close to legal lines.

Legal experts have cautioned that holding internet companies liable for the speech of their users could lead to the end of these services altogether. As explained by Jennifer Granick, surveillance and cybersecurity counsel at the American Civil Liberties Union, "There is no way you can have a YouTube where somebody needs to watch every video. There is no way you can have a Facebook if somebody needs to watch every post. There would be no Google if someone had to check every search result."³ Such liability would stifle innovation as well as individuals' freedom of expression. This model's flaws appear all the more clear in light of the significant efforts many companies make to identify and remove harmful speech—efforts that often require the protective shield of intermediary liability protection laws. (Instead, as this paper explains, other models for regulating internet platforms would better empower and encourage them to effectively address harmful content.)

These challenges suggest that retrofitting the rules that regulate offline speech for the online world may be insufficient. Instead, new frameworks are needed. This paper will put forward four questions that underpin the debate about that framework:

- 01 How can content regulation best achieve the goal of reducing harmful speech while preserving free expression?
- 02 How should regulation enhance the accountability of internet platforms?
- 03 Should regulation require internet companies to meet certain performance targets?
- 04 Should regulation define which “harmful content” should be prohibited on internet platforms?

Question 1: How can content regulation best achieve the goal of reducing harmful speech while preserving free expression?

While governments have different thresholds for limiting speech, the general goal of most content regulation is to reduce harmful speech while preserving free expression. Broadly speaking, regulation could aim to achieve this goal in three ways: (1) holding internet companies accountable for having certain systems and procedures in place, (2) requiring companies meet specific performance targets when it comes to content that violates their policies, or (3) requiring companies to restrict specific forms of speech, even if the speech is not illegal. The discussion of Questions 2, 3, and 4 will discuss these approaches in more detail.

Facebook generally takes the view that the first approach—requiring companies to maintain certain systems and procedures—is the best way to ensure an appropriate balancing of safety, freedom of expression, and other [values](#). By requiring systems such as user-friendly channels for reporting content or external oversight of policies or enforcement decisions, and by requiring procedures such as periodic public reporting of enforcement data, regulation could provide governments and individuals the information they need to accurately judge social media companies' efforts.

Governments could also consider requiring companies to hit specific performance targets, such as decreasing the prevalence of content that violates a site's hate speech policies, or maintaining a specified median response time to user or government reports of policy violations. While such targets may have benefits, they could also create perverse incentives for companies to find ways to decrease enforcement burdens, perhaps by defining harmful speech more narrowly, making it harder for people to report possible policy violations, or stopping efforts to proactively identify violating content that has not been reported.

Finally, regulators could require that internet companies remove certain content—beyond what is already illegal. Regulators would need to clearly define that content, and the definitions would need to be different from the traditional legal definitions that are applied through a judicial process where there is more time, more context, and independent fact-finders.

These approaches are not mutually exclusive nor are they the only options. However, they do represent some of the most common ideas that we have heard in our conversations with policymakers, academics, civil society experts, and regulators. The approach a government may pursue will depend on what sort of platform behavior it wants to incentivize. Ultimately, the most important elements of any system will be due regard for each of the human rights and values at stake, as well as clarity and precision in the regulation.

Question 2: How should regulation enhance the accountability of internet platforms to the public?

People in democratic societies are accustomed to being able to hold their governments accountable for their decisions. When internet companies make decisions that have an impact on people's daily lives, those people expect the same accountability. This is why internet companies, in their conversations with people using their services, face questions like, "Who decides your content standards?" "Why did you remove my post?," and "What can I do to get my post reinstated when you make a mistake?"

Regulation could help answer those questions by ensuring that internet content moderation systems are consultative, transparent, and subject to meaningful independent oversight. Specifically, procedural accountability regulations could include, at a minimum, requirements that companies publish their content standards, provide avenues for people to report to the company any content that appears to violate the standards, respond to such user reports with a decision, and provide notice to users when removing their content from the site. Such regulations could require a certain level of performance in these areas to receive certifications or to avoid regulatory consequences.

Regulation could also incentivize—or where appropriate, require—additional measures such as:

- Insight into a company's development of its content standards.
- A requirement to consult with stakeholders when making significant changes to standards.
- An avenue for users to provide their own input on content standards.
- A channel for users to appeal the company's removal (or non-removal) decision on a specific piece of content to some higher authority within the company or some source of authority outside the company.
- Public reporting on policy enforcement (possibly including how much content was removed from the site and for what reasons, how much content was identified by the company through its own proactive means before users reported it, how often the content appears on its site, etc.).

Regulations of this sort should take into account a company's size and reach, as content regulation should not serve as a barrier to entry for new competitors in the market.

An important benefit of this regulatory approach is that it incentivizes an appropriate balancing of competing interests, such as freedom of expression, safety, and privacy. Required transparency measures ensure that the companies' balancing efforts are laid bare for the public and governments to see.

If they pursued this approach, regulators would not be starting from scratch; they would be expanding and adding regulatory effect to previous efforts of governments, civil society, and industry. For example, the Global Network Initiative Principles⁴ and other civil society efforts have helped create baselines for due process and transparency linked to human rights principles. The GNI has also created an incentives structure for compliance. Similarly, the European Union Code of Conduct on Countering Illegal Hate Speech Online has created a framework that has led to all eight signatory companies meeting specific requirements for transparency and due process.

In adopting this approach, however, regulators must be cautious. Specific product design mandates—"put this button here," "use this wording," etc.—will make it harder for internet companies to test what options work best for users. It will also make it harder for companies to tailor the reporting experience to the device being used: the best way to report content while using a mobile phone is likely to differ from the best way to report content while using a laptop. In fact, we have found that certain requests from regulators to structure our reporting mechanisms in specific ways have actually reduced the likelihood that people will report dangerous content to us quickly.

Another key component of procedural accountability regulation is transparency. Transparency, such as required public reporting of enforcement efforts, would certainly subject internet companies to more public scrutiny, as their work would be available for governments and the public to grade. This scrutiny could have both positive and negative effects.

The potential negative effects could arise from the way transparency requirements shape company incentives. If a company feared that people would judge harshly of its efforts and would therefore stop using or advertising on its service, the company might be tempted to sacrifice in unmeasured areas to help boost performance in measured areas. For instance, if regulation required the measuring and public annual reporting of the average time a company took to respond to user reports of violating content, a company might sacrifice quality in its review of user reports in order to post better numbers. If regulation required public reporting of how many times users appealed content removal decisions, a company might make its appeal channel difficult to find or use.

On the other hand, companies taking a long-term view of their business incentives will likely find that public reactions to their transparency efforts—even if painful at times—can lead them to invest in ways that will improve public and government

confidence in the longer term. For example, spurred by multi-stakeholder initiatives like GNI or the EU Code of Conduct mentioned above, Facebook has increased efforts to build accountability into our content moderation process by focusing on consultation, transparency, and oversight. In response to reactions to our reports, we have sought additional input from experts and organizations.

Going forward, we will continue to work with governments, academics, and civil society to consider what more meaningful accountability could look like in this space. As companies get better at providing this kind of transparency, people and governments will have more confidence that companies are improving their performance, and people will be empowered to make choices about which platform to use—choices that will ultimately determine which services will be successful.

Question 3: Should regulation require internet companies to meet certain performance targets?

Governments could also approach accountability by requiring that companies meet specific targets regarding their content moderation systems. Such an approach would hold companies accountable for the ends they have achieved rather than ensuring the validity of how they got there.

For example, regulators could say that internet platforms must publish annual data on the “prevalence” of content that violates their policies, and that companies must make reasonable efforts to ensure that the prevalence of violating content remains below some standard threshold. If prevalence rises above this threshold, then the company might be subject to greater oversight, specific improvement plans, or—in the case of repeated systematic failures—fines.

Regulators pursuing this approach will need to decide which metrics to prioritize. That choice will determine the incentives for internet companies in how they moderate content. The potential for perverse incentives under this model is greater than under the procedural accountability model, because companies with unsatisfactory performance reports could face direct penalties in the form of fines or other legal consequences, and will therefore have a greater incentive to shortcut unmeasured areas where doing so could boost performance in measured areas. Governments considering this approach may also want to consider whether and how to establish common definitions—or encourage companies to do so—to be able to compare progress across companies and ensure companies do not tailor their definitions in such a way to game the metrics.

With that in mind, perhaps the most promising area for exploring the development of performance targets would be the *prevalence* of violating content (how much violating content is actually *viewed* on a site) and the *time-to-action* (how long the content exists on the site before removal or other intervention by the platform).

Prevalence: Generally speaking, some kinds of content are harmful only to the extent they are ever actually *seen* by people. A regulator would likely care more about stopping one incendiary hateful post that would be seen by millions of people than twenty hateful posts that are seen only by the person who posted them. For this reason, regulators (and platforms) will likely be best served by a focus on reducing the prevalence of views of harmful content. Regulators might consider, however, the incentives they might be inadvertently creating for companies to more narrowly define harmful content or to take a minimalist approach in prevalence studies.

Time-to-Action: By contrast, some types of content could be harmful even without an audience. In the case of real-time streaming of child sexual

exploitation, coordination of a violent attack, or an expressed plan to harm oneself, quick action is often necessary to save lives, and the weighing of values may be so different as to justify an approach that incentivizes fast decision-making even where it may lead to more incorrect content removals (thus infringing on legitimate expression) or more aggressive detection regimes (such as in the case of using artificial intelligence to detect threats of self harm, thus potentially affecting privacy interests).

There are significant trade-offs regulators must consider when identifying metrics and thresholds. For example, a requirement that companies “remove all hate speech within 24 hours of receiving a report from a user or government” may incentivize platforms to cease any proactive searches for such content, and to instead use those resources to more quickly review user and government reports on a first-in-first-out basis. In terms of preventing harm, this shift would have serious costs. The biggest internet companies have developed technology that allows them to detect certain types of content violations with much greater speed and accuracy than human reporting. For instance, from July through September 2019, the vast majority of content Facebook removed for violating its hate speech, self-harm, child exploitation, graphic violence, and terrorism policies was detected by the company’s technology before anyone reported it. A regulatory focus on response to user or government reports must take into account the cost it would pose to these company-led detection efforts.

Even setting aside proactive detection, a model that focuses on average time-to-action of reported content would discourage companies from focusing on removal of the most harmful content. Companies may be able to predict the harmfulness of posts by assessing the likely reach of content (through distribution trends and likely virality), assessing the likelihood that a reported post violates (through review with artificial intelligence), or assessing the likely severity of reported content (such as if the content was reported by a safety organization with a proven record of reporting only serious violations). Put differently, companies focused on average speed of assessment would end up prioritizing review of posts unlikely to violate or unlikely to reach many viewers, simply because those posts are closer to the 24-hour deadline, even while other posts are going viral and reaching millions.

A regulation requiring companies to “remove all hate speech within 24 hours of upload” would create still more perverse incentives. Companies would have a strong incentive to turn a blind eye to content that is older than 24 hours (and unlikely to be seen by a government), even though that content could be causing harm. Companies would be disincentivized from developing technology that can identify violating *private* content on the site, and from conducting prevalence studies of the existing content on their site.

The potential costs of this sort of regulation can be illustrated through a review of Facebook’s improvements in countering terror propaganda. When Facebook first

began measuring its average time to remove terror propaganda, it had technical means to automatically identify some terror propaganda at the time of upload. After considerable time and effort, Facebook engineers developed better detection tools, including a means for identifying propaganda that had been on the site for years. Facebook used this tool to find and remove propaganda, some of which had long existed on the site without having been reported by users or governments. This effort meant that Facebook's measurement of average time on site for terror propaganda removed went up considerably. If Facebook would have faced consequences for an increase in its "time-to-action," it would have been discouraged from ever creating such a tool.

Whichever metrics regulators decide to pursue, they would be well served to first identify the values that are at stake (such as safety, freedom of expression, and privacy) and the practices they wish to incentivize, and then structure their regulations in a way that minimizes the risk of perverse incentives. This would require close coordination among various regulators and policymaking bodies focused on sometimes competing interests, such as privacy, law enforcement, consumer protection, and child safety. A regulation that serves the regulatory interests of one body well may in fact harm the interests of another.

Question 4: Should regulation define which “harmful content” should be prohibited on internet platforms?

Many governments already have in place laws and systems to define illegal speech and notify internet platforms when illegal speech is present on their services. Governments are also considering whether to develop regulations that specifically define the “harmful content” internet platforms should remove—above and beyond speech that is already illegal.

International human rights instruments provide the best starting point for analysis of government efforts to restrict speech. Article 19 of the [International Covenant on Civil and Political Rights \(ICCPR\)](#)⁵ holds:

1. Everyone shall have the right to hold opinions without interference.
2. Everyone shall have the right to freedom of expression; this right shall include freedom to seek, receive and impart information and ideas of all kinds, regardless of frontiers, either orally, in writing or in print, in the form of art, or through any other media of his choice.
3. The exercise of the rights provided for in paragraph 2 of this article carries with it special duties and responsibilities. It may therefore be subject to certain restrictions, but these shall only be such as are provided by law and are necessary:
 - (a) For respect of the rights or reputations of others;
 - (b) For the protection of national security or of public order (ordre public), or of public health or morals.

Similarly, Article 10 of the [European Convention on Human Rights](#) carves out room for governments to pass speech laws that are necessary for:

...[T]he interests of national security, territorial integrity or public safety, for the prevention of disorder or crime, for the protection of health or morals, for the protection of the reputation or rights of others, for preventing the disclosure of information received in confidence, or for maintaining the authority and impartiality of the judiciary.⁶

In countries with strong democratic institutions, laws governing speech are passed as part of a transparent process that weighs different societal interests, with due respect for human rights and with the check of an impartial judiciary. Drawing a clear line between legal and illegal speech helps citizens to know what is and is not permissible. Judicial checks and balances help to reduce the likelihood of capricious restrictions on legitimate speech.

Despite these international human rights norms, laws affecting speech and even criminalizing it vary widely both by jurisdiction and by the type of harm contemplated, even among liberal democratic nations.

Among types of content most likely to be outlawed, such as content supporting terrorism and content related to the sexual exploitation of children, there is significant variance in laws and enforcement. In some jurisdictions, for instance, images of child sexual exploitation are lawful if they are computer generated. In others, praise of terror groups is permitted, and lists of proscribed terror groups vary widely by country and region. Laws vary even more in areas such as hate speech, harassment, or misinformation. Many governments are grappling with how to approach the spread of misinformation, but few have outlawed it. As the UN and others have noted, the general criminalization of sharing misinformation would be “incompatible with international standards for restrictions on freedom of expression.”⁷ Thresholds also vary for defamatory speech and offensive material.

The variation among laws is significant because it underscores the difficulty of weighing competing values like safety and expression. If governments do not agree on how to balance these interests when writing laws, it is likely they will have very different ideas of how to define online content standards, which require even more specificity than laws in order to ensure consistent application.

Laws restricting speech are generally implemented by law enforcement officials and the courts. Fact finders investigate and due process is afforded to the speaker. The penalty for illegal speech can include fines or imprisonment.

Internet content moderation is fundamentally different. Companies use a combination of technical systems and employees and often have only the text in a post to guide their assessment. The assessment must be made quickly, as people expect violations to be removed within hours or days rather than the months that a judicial process might take. The penalty is generally removal of the post or the person's account.

Given these complexities, governments that seek to define for internet companies what content they should allow on their platforms should seek to:

- Create a standard that can be enforced practically, at scale, with limited context about the speaker and content, without undermining the goal of promoting expression.
- Take account of different user preferences and expectations for different sorts of internet services. Different definitions may be appropriate based on the nature of the service (search engines versus social media) or the nature of the audience (private chat versus public post, ephemeral versus more permanent, etc).

- Take account of individual controls, preferences, settings, etc. that a company may provide its users for some kinds of potentially harmful content (e.g., nudity).
- Provide flexibility so that platforms can adapt policies to emerging language trends and adversarial efforts to avoid enforcement. For instance, hateful speech, bullying, and threats of self-harm are often expressed through a lexicon of words that fall in and out of favor or evolve over time. Similarly, proscribed terror groups and hate groups may rename themselves or fracture into different groups.

Principles for Future Regulators

03

In some countries, a regulatory system for online content may require a new type of regulator. Addressing online communication goes beyond simply adopting the core capabilities of traditional regulators and channeling them towards the task of effective oversight. Instead, any regulator in this space will need proficiency in data, operations, and online content. Governments will also need to ensure that regulatory authorities pay due regard to innovation and protect users' rights online, including consideration of the following:

INCENTIVES

Ensuring accountability in companies' content moderation systems and procedures will be the best way to create the incentives for companies to responsibly balance values like safety, privacy, and freedom of expression.

THE GLOBAL NATURE OF THE INTERNET

Any national regulatory approach to addressing harmful content should respect the global scale of the internet and the value of cross-border communications. It should aim to increase interoperability among regulators and regulations. However, governments should not impose their standards onto other countries' citizens through the courts or any other means.

FREEDOM OF EXPRESSION

In addition to complying with Article 19 of the ICCPR (and related guidance), regulators should consider the impacts of their decisions on freedom of expression.

TECHNOLOGY

Regulators should develop an understanding of the capabilities and limitations of technology in content moderation and allow internet companies the flexibility to innovate. An approach that works for one particular platform or type of content may be less effective (or even counterproductive) when applied elsewhere.

PROPORTIONALITY AND NECESSITY

Regulators should take into account the severity and prevalence of the harmful content in question, its status in law, and the efforts already underway to address the content.

The Beginning of a Conversation

04

We hope that this paper contributes to the ongoing conversations among policymakers and other stakeholders. These are challenging issues and sustainable solutions will require collaboration. As we learn more from interacting with experts and refining our approach to content moderation, our thinking about the best ways to regulate will continue to evolve. We welcome feedback on the issues discussed in this paper.

Online Content Regulation

1. Zuckerberg, Mark. "The Internet Needs New Rules. Let's Start in These Four Areas." Washington Post, 30 March 2019.
2. Bowers, John, and Jonathan Zittrain. "Answering Impossible Questions: Content Governance in an Age of Disinformation." Harvard Kennedy School Misinformation Review, 2020.
3. As quoted in: Reynolds, Matt. "The Strange Story of Section 230, the Obscure Law That Created Our Flawed, Broken Internet." Wired, 28 March 2019, www.wired.co.uk/article/section-230-communications-decency-act.
4. "The GNI Principles." Global Network Initiative, 10 January 2020. <https://globalnetworkinitiative.org/gni-principles/>.
5. UN General Assembly, International Covenant on Civil and Political Rights, 16 December 1966, United Nations, Treaty Series, vol. 999, p. 171.
6. Council of Europe, European Convention for the Protection of Human Rights and Fundamental Freedoms, as amended by Protocols Nos. 11 and 14, 4 November 1950.
7. "Joint Declaration on Freedom of Expression and 'Fake News', Disinformation and Propaganda." The United Nations (UN) Special Rapporteur on Freedom of Opinion and Expression, the Organization for Security and Co-operation in Europe (OSCE) Representative on Freedom of the Media, the Organization of American States (OAS) Special Rapporteur on Freedom of Expression and the African Commission on Human and Peoples' Rights (ACHPR) Special Rapporteur on Freedom of Expression and Access to Information, 3 March 2015. As noted in the Joint Declaration and explored infra, those International Standards are: "that they be provided for by law, serve one of the legitimate interests recognised under international law, and be necessary and proportionate to protect that interest."