



PROTECTING THE OPEN INTERNET

Regulatory principles for policy makers



Introduction

The pressures driving internet regulation around the world are wide-ranging, highlighting the diverse challenges, competing policy priorities, and the implications of widespread technological adoption that societies face. How these challenges are addressed will affect services of all sizes, the ability of billions of people to share information with the world and hear from others across borders, the future of the digital economy, and the survival of a free and secure Open Internet.

The potential consequences are vast, go far beyond today's headlines, and are bigger than any one company. There are no easy answers, and there are a lot of trade-offs. By designing regulation around the largest services today, or by only responding to the challenges faced in certain regions, the future of the internet will be defined by these choices, and the innovation needed to solve challenges and widen participation will fall short. The Open Internet is not something to be taken for granted, and in the coming years, decisions will be made that define its future. The risk that the rhetoric of policy and language of law will be co-opted and weaponised by those seeking to usher in an age of techno-nationalism is real.

Regulatory approaches to new industries are often shaped by the policy responses designed in the aftermath of the industrial revolution, oriented towards frameworks that specify standards for outcomes of mechanical processes. This approach struggles to adapt to the unpredictable and rapidly evolving nature of human use of technology and expression. More broadly, the policy issues faced are often rooted in societal challenges. They demand a whole of society response and will not be solved by the removal of content online alone. Bad actors seeking to exploit online services to undermine elections, spread disinformation, and harm others will not be deterred by their accounts being removed.

This paper explores a range of public policy challenges, how they intersect with issues of competition, content moderation, and the role and responsibilities of services like Twitter. We offer these principles to inform the policy debate, recognising the need to balance tackling harm with protecting the global free and secure Open Internet.



Guiding principles for regulation

1 The Open Internet is global, should be available to all, and should be built on open standards and the protection of human rights.

2 Trust is essential and can be built with transparency, procedural fairness, and privacy protections.

3 Recommendation and ranking algorithms should be subject to human choice and control.

4 Competition, choice, and innovation are foundations of the Open Internet and should be protected and expanded, ensuring incumbents are not entrenched by laws and regulations.

5 Content moderation is more than just leave up or take down. Regulation should allow for a range of interventions, while setting clear definitions for categories of content.



The Open Internet is global, should be available to all, and should be built on open standards and the protection of human rights.

The Open Internet has been an unprecedented engine for economic growth, cultural development, and self-expression. But to continue this impact, it must be available to all.¹ A foundational goal of all digital policy should be to protect the global, free, and secure Open Internet.

The infrastructure of the internet is itself now a geopolitical space. Governments should prioritize policies, partnerships, and investments at home and abroad that support and defend the Open Internet, both through regulatory and standards bodies, as well as ensuring domestic regulation does not undermine global norms or set dangerous precedents. Open standards championed by these bodies will provide for greater interoperability, connection, and competition.

Access is a critical issue. Throttling or blocking of the internet must be resisted, and the principle that information should be able to safely and securely move across borders freely as part of a global internet should be core to democratic regulation. Enacting and enforcing these rules without considering the global nature of the internet runs the risk of isolating citizens from the global conversation that the Open Internet serves, with a social and economic cost.

Rhetoric and policies pursuing national data sovereignty should be avoided and scrutinised. Some actors seek to exploit this concept to strengthen control of and access to data, weakening the Open Internet through forced data localisation and limits on the free flow of data. The principle that data belongs to a person does not mean that all people's data belongs to the state.

Policy makers should avoid the use of extra-territorial application of national content standards as this further undermines the global internet and encourages a race to the bottom, with the entire world's open communications imperiled by those actors least committed to freedom of expression.

Both governments and industry should ensure their approach to addressing online harm is consistent with universally recognized human rights norms, including proportionality and the protection of privacy and freedom of expression.

¹ contractfortheweb.org/principle-theme/access/



Trust is essential and can be built with transparency, procedural fairness, and privacy protections.

There's a deficit in trust in many online services and government functions alike. It's essential every sector works to rebuild trust, beginning with greater transparency. People should understand the rules of online services and the way that governmental legal powers are used. Transparency enables accountability for companies and Governments. Without transparency, there can be no accountability.

One of the critical areas where policymakers and regulators can enhance transparency is ensuring that laws governing information provide suitable flexibility for valuable disclosures, for example, the provision of data to academics and researchers. While Twitter has taken the decision to publish archives of removed content attributed to state-linked information operations, there's a rich spectrum of work that could be enabled through smart regulation of such disclosures.

Just as due process is a core feature of robust judicial systems, procedural fairness should be a core function of online services. These concepts should be at the core of regulation, particularly where governments seek to require services to remove content and companies take action under their terms of service.

Regulation by proxy, where governments use broad standards to push the burden of defining types of content onto service providers to avoid having to do so in legislation, is a dangerous trend, particularly when set alongside seemingly contradictory obligations to protect certain types of content. This is fundamentally a constitutional issue as much as a trust issue. Both individuals and companies need notice of what is prohibited by law so that they can act accordingly.

Technology will continue to accelerate and change far faster than laws will be passed, with decentralised services and blockchain technology already upending traditional regulatory approaches. Some governments will seek to control these new services or undermine the global adoption of them while their proxies seek to influence domestic policy debates.



Legislators should ensure clear harmonized standards for safeguarding and processing personal data, supplemented by regulatory guidance as new issues emerge, recognising that it's neither feasible nor desirable to legislate for every potential scenario of how personal data is used in primary legislation. Fragmented and inconsistent frameworks, both within countries and internationally, weaken consumer protection and the ability of well-understood norms to develop. While many services do collect data to enable them to provide advertising, granular privacy controls balance the functionality of online services with consumer control while serving a desire to allow people who use services to make informed decisions about the data they share. Individuals should know, and have meaningful control over, what data is being collected about them, how it's used, and when it's shared. In the long run, regulation should protect and encourage services based on a range of business models, not just those built on advertising.

Policymakers should protect the ability to use the internet without having to disclose your real identity, legal ID, or detailed personal information. This is foundational to a universally accessible internet for all, and it's essential to recognise that not all services require the same amount of personal information to be disclosed or verified and nor should they be required to.

Recommendation and ranking algorithms should be subject to human choice and control.

As algorithms and machine learning increasingly shape our online experiences, the decisions people make online have long-lasting consequences, some of which we may not be able to foresee. Recognising that content moderation and content organisation are two different spheres of work, particularly when content is recommended without a positive signal to seek it out, policymakers should prioritise empowering people to have control over algorithms they interact with and ultimately drive an ability to make our own choices between algorithms. Choice can also help foster greater understanding and awareness of how algorithms impact people's online experiences, leading to greater digital literacy.

While algorithmic transparency is an important part of deepening understanding of how these systems work, both in terms of process and training data, the focus on source code for algorithms, a literal interpretation of the phrase "algorithmic transparency" offers flawed and unclear benefits. While in a limited context this may provide a small, highly technical audience with insights, it does little to change the experience of people online.



The first step is the ability to control whether an algorithm is shaping your experience. For example, in 2018, Twitter introduced the ability to turn off our Home timeline ranking algorithm, returning people to a reverse-chronological order of Tweets. This control enables transparency — people can see how the content appears in the two environments. In the long term, as we envision through our @bluesky project, this control will extend to the choice between ranking algorithms, built on an open standard for social media that we hope Twitter will ultimately become a client of. The idea of “Protocols not platforms”² is instructive not only for the technological potential for standardization of ranking algorithms but also the underlying impact this would have on protecting free expression and driving competition.

Competition, choice, and innovation are foundations of the Open Internet and should be protected and expanded, ensuring incumbents are not entrenched by laws and regulations.

A less competitive internet trends towards a less open internet. There’s a risk that some regulatory interventions will undermine competition and entrench incumbent services, reducing consumer choice. It’s not unique to the technology sector that incumbents will often seek to use new regulations to protect their own market dominance, and just because services are online does not mean they depend on the Open Internet. Indeed, in some cases, a less open internet may suit certain businesses more.

Strong net neutrality protections, which protect against a two-tier internet that treats data according to an ability to pay for prioritisation, are needed to protect new entrants and innovators from well-resourced incumbents and infrastructure gatekeepers.

Competition in the online service space depends on a number of pillars, which sometimes are portrayed as only benefiting large providers. This framing is often misleading, given that these protections currently — and should continue to — benefit services of all sizes and are of most importance to those with fewer resources.

² knightcolumbia.org/content/protocols-not-platforms-a-technological-approach-to-free-speech



Firstly, intermediary liability protection is a foundation of the global, Open Internet and critical to the competition online. Without this foundation, the internet as we know it — allowing the speech, interaction, and discovery of billions of people — would cease to exist. Intermediary liability protections enable two crucial functions. They stop people from silencing others by litigating against a service rather than the person responsible for posting the content. Without these protections, services would be forced to choose between expensive litigation or removing content on their service. Secondly, it enables companies to take proactive action on content that may be legal but violates their terms of service without fear of litigation.

Critically, they enable services to set their terms of service to serve their audience best. Whether prohibiting profanity on a children’s service or allowing discussion of controversial subjects, this diversity is essential to the competition between services while also enabling the greatest range of choice and vehicles for expression.

Policymakers should avoid mandating technical means of implementation that have the effect of further entrenching services based on those tools and technologies, or by benefiting those that have the financial and technical means to deploy the particular implementation proposed, not to mention the vendors promising a simple solution. Opportunities to expand interoperability and the adoption of open standards will empower people with greater choice and flexibility about how they interact with online services and drive competition.

Finally, the technologies that underpin the ability to address and remove the most harmful content and respond to further harms remain in proprietary silos, becoming exponentially more effective as businesses scale, further enhancing dominance and undermining competition. Content moderation technology is one of the most significant barriers to entry, particularly as regulators set ever stricter requirements on the time taken to remove harmful content. Policymakers should encourage and facilitate a fundamental change in the availability of proactive technologies and the data that underpin them to enable new services and tools to be made more accessible to a greater range of services, including providing a robust legal framework for information sharing.



Content moderation is more than just leave up or take down. Regulation should allow for a range of interventions while setting clear definitions for categories of content.

Legislation and regulation should set clear standards for the types of content they seek to address, with substantive definitions and boundaries and consistent with human rights standards. Where the content at issue is lawful, but a government believes there's a need to intervene, the regulatory framework should clearly distinguish between these types of content. Government requests for the removal of specific pieces of content based on illegality should be based upon legal process and provide for transparency about how these powers are used. It's a fundamental question of due process that a government agency, not a private actor, is responsible for determining criminality. Companies should be free to provide notice to people that this was the basis for action being taken.

Secondly, we believe the regulatory debate needs to reflect how content moderation is now more than just leaving content up or taking it down. Providing users with context, whether concerning an account, piece of content, or form of engagement, is more informative to the broader public conversation than removing content while providing controls to people and communities to control their own experience is empowering and impactful. Equally, de-amplification allows a more nuanced approach to types of speech that may be considered problematic, better striking a balance between freedom of speech and freedom of reach. Long term, how attention is directed is a critical question.

Thirdly, regulatory frameworks that address system-wide processes, as opposed to individual pieces of content, will be able to better reflect the challenges of scale that all modern communications services involve, in addition to the way that challenges change depending on if you are trying to protect a certain group, like young people, or a particular type of behavior, such as platform manipulation.

³ blog.twitter.com/en_us/topics/company/2020/new-approach-to-synthetic-and-manipulated-media.html



As has been noted by a range of voices, the combination of significant administrative penalties for individual pieces of content and expected removal in short time periods — whether one hour or 24 hours — creates a significant corporate incentive to over-remove content, particularly in edge cases, and one that more acutely impacts small companies and new services who have more limited resources to litigate or pay fines. These frameworks must be underpinned with strong, independent processes and free from political interference while allowing for civil society participation.

Mistakes will happen, as they do in all large processes involving human decision-making. To avoid incentivising over-removal, regulation that assesses the system-wide performance of how services enforce their terms of service will provide essential flexibility and reduce incentives to over-moderate content while incentivising investment in technological solutions despite the inevitable errors that come from imperfect tools and robust appeal mechanisms.

Conclusion

This paper sets out high-level principles to inform debates about content moderation and regulation happening around the world. There are clear areas where the continued lack of regulation puts the onus on technology companies to fill the vacuum with their own standards, for example, political advertising. At the same time, there's a desire to deal with the companies and issues most commonly in the headlines today, without sufficient consideration of how this will impact the future of the Internet or where policy objectives might be contradictory and need resolving by Governments directly.

The Open Internet is more at risk now than ever before. Governments who seek to defend and expand online freedom cannot stand by while other countries seek to silence critics, censor journalists, and block access to information. The harassment of employees of service providers is a worrying norm, accelerated by proposals to require local staff to be liable for decisions rather than the corporate entity. Similarly, the targeting of independent journalists and activists highlights the willingness of some states and actors to use digital policy and manipulation to control political debate. As the control of digital infrastructure is increasingly a focus of geopolitical action, these issues cannot be viewed in isolation. It is essential that there is a coordinated, multi-stakeholder strategy to respond to these threats and defend the free, secure, and global Open Internet.

