



September 1, 2020

The Honorable Mike Lee
United States Senate
361A Russell Senate Office Building
Washington, DC 20510

Dear Senator Lee:

Thank you for your letter dated July 30, 2020, regarding Google's content moderation standards. We welcome the opportunity to provide more information about our content moderation policies.

Our mission at Google is to organize the world's information and make it universally accessible and useful. We are dedicated to providing access to information and freedom of expression, and we empower people to engage a diversity of sources and opinions. That said, we do have content rules that we enforce in an effort to make sure our platforms are safe, such as the [Community Guidelines](#) that set the rules of the road for what we do not allow on YouTube. Our guidelines are public and we rely on a combination of technology and people to enforce them.

To be clear, our content moderation standards are apolitical, unbiased and do not preference one point of view over another. We apply our policies to all content creators across the board and will not allow any form of political bias. Our platforms empower a wide range of people and organizations from across the political spectrum, giving them a voice and new ways to reach their audiences. Some of our biggest critics on the right and left have gotten millions of views and subscribers through our platforms.

We agree that Section 230 is a vital part of the internet ecosystem. Section 230 helps internet companies address harmful content and allows us to enforce content rules that protect our users. It also safeguards a platform's ability to promote open access to information and protect free expression online. Instead of overblocking speech, the law supports platforms to responsibly manage content.

Please find answers to your specific questions below.

1. What content-moderation standards do you employ when you remove content from your platform, where the content does not violate state or federal laws? Specifically, please explain your standards for removing content related

to:

Removal of content may occur for two reasons: it violates the law or it violates the 'rules of the road' for that product or service. We comply with the law in each country in which we operate and remove illegal content on our platforms in that country. In every country in which we operate, the unique cultures, histories, and forms of government have produced different laws governing what is considered permissible expression.

In addition, we develop and maintain 'rules of the road' which outline what types of content and behaviors are acceptable for each product or service. Known as '[content policies](#)' or '[Community Guidelines](#)' for YouTube, we aim to make them clear and easily accessible to all users and content creators – whether those are video creators, webmasters, app developers, commenters, or advertisers. These 'rules of the road' articulate the purpose and intended use of a given product or service and represent a crucial part of what makes that product unique. They also explain what types of content and behaviors are not allowed, and the process by which a piece of content, or its creator, may be removed from the service. We also have [ads policies](#) governing the use of our advertising products that all advertisers need to follow in order to be able to advertise on our platforms. These policies are designed not only to abide by laws but to ensure a safe and positive experience for our users. This means that our policies prohibit some content that we believe to be harmful to users and the overall advertising ecosystem.

Each of the products and services we offer has a different purpose, and people have different expectations of what kind of content they will interact with on each. So, we tailor our approach to the content that should be available on each product and service carefully.

Specific standards and practices pertinent to your question include:

a. COVID-19

- i. We are committed to providing timely and helpful information to people around the world during the COVID-19 pandemic. In response, we are elevating authoritative information and removing harmful misinformation that risks individual harm and undermines Government efforts to reduce infection rates. We've removed 200,000 coronavirus videos globally with dangerous or misleading coronavirus information on YouTube. We've also removed over 200 million coronavirus related ads globally for policy violations including price-gouging, capitalizing on global medical supply

- shortages, and making misleading claims about cures.
- ii. We [recently introduced](#) additional Google Ads protections by expanding our dangerous or derogatory content policies for both publishers and advertisers to include content about a current, major health crisis that contradicts authoritative scientific consensus. As a result, content contradicted by scientific consensus during COVID-19 such as origin theories, claims the virus was created as a bioweapon, as well as content that claims the virus is a hoax or government-funded will not be permitted on our ads platform. We have developed new resources to connect users to helpful information and resources, including a dedicated experience for COVID-19 on Google Search that provides easy access to authoritative information from government health authorities along with data, news, and locally-relevant information from trustworthy sources, YouTube Information Panels, homepage promotions around the world, and a website with resources dedicated to COVID-19 education and prevention.
 - iii. YouTube's Community Guidelines outline what content is not allowed on YouTube. On YouTube, our Community Guidelines prohibit content that encourages dangerous or illegal activities that risk serious physical harm or death, including certain types of medical misinformation. As the COVID-19 situation has evolved, we have partnered closely with global and local health authorities to ensure our policy definition and enforcement is effective in removing violative content where there is a serious risk of egregious harm. This work has evolved into a comprehensive [COVID-19 medical misinformation policy](#), which prohibits, for example, content that denies the existence of the coronavirus or encourages the use of home remedies in place of medical treatment. At a high level, we do not allow content that spreads medical misinformation that contradicts the World Health Organization (WHO) or local health authorities' medical information about COVID-19. Content relating to COVID-19 may be found to violate other policies as well, including our hate speech and harassment policies, depending on the content and context.

b. violent riots, and how you distinguish them from peaceful protests

- i. Violent or gory content intended to shock or disgust viewers, or content encouraging others to commit violent acts, are not allowed on YouTube. Footage, audio, or imagery involving protests or riots with the intent to shock or disgust viewers is prohibited under our Community Guidelines. However, we have carved out exceptions to the Community Guidelines for material that is Educational, Documentary, Scientific, and/or Artistic (EDSA). Videos and comments that fall under those exceptions are crucial to understanding the

world and to chronicling history, whether it is documenting wars and revolutions, or artistic expression that may include nudity.

- ii. The teams that review content at YouTube are able to see the surrounding context during their review of reported content, including the video description, other content uploaded to the channel, and metadata (titles, tags or captions). These contextual clues are important in evaluating the intent of the upload. In addition, our review tool captures the timestamp at which a video was flagged and our webforms ask submitters to include timestamps. This enables our reviewers to focus on the potentially problematic moments within a video.

c. hate speech

- i. YouTube has clear policies prohibiting hate speech, such as content that promotes or incites hatred against individuals or groups based on attributes like their race, religion, disability, age, gender, veteran status, or sexual orientation. We don't allow the promotion or incitement of violence against victims of a violent event, such as a school shooting, or their kin. In 2019, working in collaboration with creators and external organizations such as NGOs, we made significant updates to our [hate speech](#) and [harassment](#) policies, which dramatically improved our ability to protect the YouTube community. Last quarter alone, thanks to this update and our ongoing enforcement, we removed more than 80K videos and 150M comments for violating these policies.
- ii. We value diversity and respect for others, and we strive to avoid offending users, so we don't allow ads or destinations for ads that display shocking content or promote hatred, intolerance, discrimination, or violence. We have strict ad policies to protect users against [Dangerous or Derogatory](#) content. Under this policy we don't allow, for example, content that incites hatred against, promotes discrimination of, or disparages an individual or group on the basis of their race or ethnic origin, religion, disability, age, nationality, veteran status, sexual orientation, gender, gender identity, or other characteristic that is associated with systemic discrimination or marginalization. To give you a better understanding about the scale of our enforcement, [in 2019 alone](#), we blocked and removed 2.7 billion bad ads—including 670,000 ads for violating our Dangerous or Derogatory content policy. We also suspended nearly 1 million advertiser accounts for policy violations. On the publisher side, we terminated over 1.2 million accounts and removed ads from over 21 million web pages that are part of our publisher network for violating our policies. Terminating accounts—not just removing an individual ad or page—is an especially effective enforcement tool that we use if advertisers or publishers engage in egregious policy violations or have a history of violating policy.

d. protections for the unborn

- i. The Google Ads [Healthcare and medicines policy](#) requires all advertisers in the US who want to run search ads related to abortions to be certified as either providing abortions or not providing abortions. Based on the certification, Google will automatically generate an in-ad disclosure that the advertiser either “Provides abortions” or “Does not provide abortions.” This additional transparency will help users decide which abortion-related ads are most relevant to them.
- ii. We are committed to protecting users and ensuring that advertisers are using our platforms in a responsible manner, and we will continue to look for ways to strengthen our ads-related policies.

e. Misinformation

- i. On YouTube and Search, we continuously improve our ranking systems to elevate authoritative sources and decrease the propagation of low quality information. For example, Search ranks websites [based on hundreds of factors](#), from relevance to authoritativeness to recency. This ranking is informed by feedback from Search raters from all around the country and the world who assess each and every improvement to Search based on our publicly available [rater guidelines](#).
- ii. On YouTube, we take a holistic approach to disinformation through several policies in our Community Guidelines. These policies include prohibitions against spam, deceptive practices, scams, hate speech, harassment, and harmful manipulated media. For example, our deceptive practices policy prohibits content that deliberately seeks to spread disinformation that could suppress voting or otherwise interfere with democratic or civic processes, such as demonstrably false content making claims of different voting days for different demographics. In addition, we have mechanisms in place to reduce the recommendation of content that brushes right up against our policy line, including harmful misinformation. Since making changes to our recommendations systems, we've seen a substantial drop in borderline content and misinformation. We are proud of our efforts to tackle misinformation and fake news and will continue to work hard and do everything we can to be a useful and trustworthy source of information for everyone and curb misinformation in our products.

f. Terrorist influence

- i. We have strict prohibitions on terrorist content. We do not allow any content produced by or in support of a foreign terrorist organization on YouTube or other Google hosted platforms. If a violent group is not a designated terrorist organization, our policies still prohibit hate speech, gratuitous violence, incitement to violence, and other forms of intimidation.
- ii. Whilst we remove terrorist content that violates our policies, such as content

depicting terrorist acts, we do allow for contextualized coverage from authoritative sources on YouTube, like news outlets. The context in which a piece of content is created or shared is an important factor in any assessment about its quality or its purpose. We are attentive to educational, documentary, scientific or artistic contexts including journalistic intent, where the content might otherwise violate our policies.

2. How did you formulate standards for the above categories of content? What sources did you look to informing your content policies?

As discussed above, we design the 'rules of the road' across all our products and services to protect users from harm while supporting the purpose of the product. For each product and service, we tailor these policies to strike the appropriate balance between providing access to a diversity of voices and limiting harmful content and behaviors. This balance can differ from one product to the next, in part because harm manifests differently in each service and context. While a universally recognized harm may be prohibited across all our products and services, it can appear on each product and service differently. So, we must evaluate the potential for harm specific to each product and design our policies accordingly. This includes harm to an individual and harm that may affect an entire society, such as an attempt to interfere with elections or civic processes.

To help us identify emerging harms and gaps in our existing policies, we consider expert input, user feedback, and regulatory guidance. We rely on research performed by analysts who study the evolving tactics deployed by bad actors, trends observed on other platforms, and emerging cultural issues that require further observation. We also engage in conversations with regulators around the world. Their perspectives and concerns directly inform our policy process.

At YouTube, we're committed to maintaining an open platform where many different viewpoints and types of speech are welcome, including speech that some people might find problematic. We consult with experts as we look to draw the line in the right place, and review our policies and systems on an ongoing basis.

3. What are the prerequisites for a content-moderator position at your company? Do you inquire about the political or other beliefs of a candidate before making a hiring decision? Where you use contractors to serve in these roles, how do you ensure that they follow your internal guidelines and

standards?

While we rely heavily on machines and technology to reach our goals on information quality and content moderation, human reviewers also play a critical role. Content moderators help us assess context and nuance, to evaluate content we've never seen before, and make distinctions and decisions

We work with third-party vendors and contractors to help us scale our content moderation efforts, and provide the native language expertise and the 24-hour coverage required of a global platform. When we work with these providers, we engage in regular site visits and audits to ensure that our guidelines and Supplier Code of Conduct are respected.

As discussed further below, we have robust systems to ensure that employees' and contractors' personal views do not impact our products and that our policies are enforced without regard to political viewpoints. These include policies that prohibit unethical behavior and altering or compromising Google's systems to achieve some personal goal or benefit. If we find instances of a Googler or contractor violating our policies, we'd take appropriate action.

4. What is the internal process that your content moderators follow to remove content that violates your standards?

To enforce our policies at the scale of the web, we rely on a mix of automated and human efforts to spot problematic content. In addition to flags by individual users, sophisticated automated technology helps us detect problematic content at scale. Our automated systems are carefully trained to quickly identify and take action against spam and violative content.

This includes flagging potentially problematic content for human reviewers, whose judgement is needed for the many decisions that require a more nuanced determination. The context in which a piece of content is created or shared is an important factor in any assessment about its quality or its purpose. We are attentive to educational, scientific, artistic, or documentary contexts, including journalistic intent, where the content might otherwise violate our policies.

YouTube maintains a more detailed and living set of enforcement guidelines that provide guidance on the enforcement of the public Community Guidelines. These enforcement guidelines are extensive and dynamic to ensure that the policies apply to changing trends and new patterns of controversial content online.

In addition, our expert teams around the world handle the investigations of more sophisticated threat actors that are adept at circumventing the automated defenses we build into our products. New forms of abuse and threats are constantly emerging that require human ingenuity to assess and plan for action before an automated system can address them at scale. So, we operate dedicated threat intelligence and monitoring teams (e.g. Google's Threat Analysis Group), which provide insights and intelligence to our policy development and enforcement teams so they can stay ahead of bad actors.

5. How do you ensure that a content-moderation decision is not influenced by the personal beliefs or political views of the moderator?

Google's mission – to make the world's information accessible and useful for everyone – is not Democratic or Republican and is dependent on being a useful and trustworthy source of information. We design products that are for everyone and enforce our policies in an apolitical way.

On YouTube, we aim to enforce our policies consistently, with a focus on the content rather than the speaker or poster. This allows us to apply our policies consistently. For example, we will remove violative content if posted by elected officials or other public figures, unless there educational, documentary, scientific, or artistic considerations (e.g., a clip of a political figure included in a documentary or a news report).

For Ads, we have a wide range of ads policies that all advertisers need to follow in order to be able to advertise on our platforms. We build our ads products and enforce our policies in a neutral way. Our enforcement involves both automation and human review. Our automated system is able to detect and remove many infringing ads, for instance where ads are duplicates of previously reviewed ads. Ads that warrant human review are scrutinized by members of our Trust and Safety team, who use a variety of internal and external resources to assess the ad. Our policies consist of technical requirements (such as that your ad must lead to a functioning landing page) and content requirements.

We build our products for everyone. While our more than 100,000 employees around the world hold a wide variety of views, we have safeguards in place to ensure that we design and enforce our policies in a way that is free from improper bias.

As mentioned earlier, to ensure Search algorithms meet high standards of relevance and quality, we have a [rigorous process](#) that involves both live tests and feedback from thousands of trained external Search Quality Raters from around the world.¹

The [Search Quality Rater Guidelines](#) that define the goals of our ranking systems include the criteria that our raters use to assess the expertise, authority, and trustworthiness of pages.² These criteria do not include political ideology and specifically provide guidance for raters that *“ratings should be based on the instructions and examples given in these guidelines. Ratings should not be based on your personal opinions, preferences, religious beliefs, or political views.”* Furthermore, whether a business, individual, or organization buys ads is not a factor in our search algorithms. We never provide special treatment to advertisers in how our search algorithms rank their websites, nor how our policies are enforced, and nobody can pay us to do so.

In addition, we conduct live traffic experiments to measure how users interact with a new feature before releasing it more widely. Results from these experiments are reviewed by experienced engineers and search analysts. They collectively determine whether the change is approved to launch. In 2019, we conducted over 460,000 experiments with trained external Search Quality Raters and live tests, which resulted in more than 3,600 improvements to Google Search.

This commitment goes beyond ranking. A diverse set of external and internal stakeholders are consulted during policy development. Our process involves multiple Google teams, and leaders are involved in finalizing a new or updated policy. We outline our product policies and guidelines in help centers and other fora so that our users can understand the rules that apply to our products.

For some Search features such as [Discover](#), [news surfaces](#), or [featured snippets](#), to name some examples, we have policies specifying what is eligible to appear. These are intended to especially ensure we are not surfacing shocking, offensive, hateful, violent, dangerous, harmful or similarly problematic material in places where users may have higher expectations given the special formatting or presentation of such features. While we design our algorithms to prevent potentially policy-violating content from

¹ <https://www.google.com/search/howsearchworks/mission/users/>

² For more on Search Quality Raters Guidelines, see www.google.com/search/howsearchworks

appearing in these features, our systems are not perfect. If content violates our policies for these Search features, we will remove it and improve our systems to avoid such an occurrence. Failure to meet these eligibility requirements can lead to pages and sites being blocked from these special features; should pages and sites later meet eligibility requirements, they can reappear. None of these processes affect how pages or sites appear outside these features. Neither our policies nor our enforcement of those policies takes political viewpoints into account.

Content in Search might also be blocked or removed based on our legal, webmaster guidelines or voluntary removals policies, which are limited to things like copyright, spam and sensitive personal information like government IDs. In these cases, content that is reported to us or that we identify to be in violation of our [webmaster guidelines](#) is filtered from our results to adhere to the law and our policies.

In addition, we enforce our policies consistently, regardless of who or what is involved. “Gray area” cases – those that approach a policy boundary – are reviewed by multiple people to ensure that an appropriate decision is made, and we have a rigorous quality assurance process for all cases across our products. We approach with similar caution the development and use of the safety lists that help us ensure, for instance, that a website we demonetised for the most severe infringement of our advertising policies is not inadvertently offered the possibility to monetize again via another of our services.

This of course does not mean that our products are ‘neutral.’ Any ranking inherently involves classification by reference to a specific set of goals or factors. However, we make sure that we publicly document the kinds of goals and factors our products optimize for, and we welcome feedback. For instance, our Search Quality Rater Guidelines outline how we characterize expertise, authoritativeness, or trust for Google Search.

6. If your content-moderation standards rely on guidance from a government entity, please explain your policy for allowing on your platform speech that disagrees with the government. If CDC guidance is the basis for removing content regarding COVID-19, how is that standard applied consistently? For example, since the CDC says that it is safe for schools to open, do you remove content from your platform that claims it is unsafe to reopen schools?

YouTube’s [COVID-19 misinformation policy](#), an extension of our harmful and dangerous content policy, prohibits content relating to COVID-19 that poses a serious risk of egregious harm. Our policy prohibits content containing certain types of medical misinformation that contradicts the World Health Organization (WHO) or

local health authorities' medical information about COVID-19. We prohibit content that contains misinformation on COVID-19 treatment, prevention, diagnostic methods, or transmission that runs counter to WHO or local health authority guidance and where there is a serious risk of egregious harm. The example content provided in the question would not be removed unless there were other statements in the same video that violated the COVID-19 misinformation policy or any other one of our Community Guidelines. This policy has evolved since March, and we continue to monitor the situation to make adjustments where appropriate.

7. Where do you clearly articulate your content-moderation standards? How do you convey your moderation standards to consumers? Do you regularly update your users on changes made to your policies?

Our policies work best when users are aware of the rules and understand how we enforce them. That is why we work to make this information clear and easily available to all. We develop comprehensive help centers, websites outlining our policies, and blog posts that detail the specific provisions of our policies as well as updates to these policies.

For YouTube, we recently launched [How YouTube Works](#), which includes a section on our Community Guidelines and enforcement thereof. In addition, we regularly release reports that detail how we enforce our policies or review content reported to be in violation of local law. The [YouTube Community Guidelines Enforcement Transparency Report](#) provides quarterly updates on the number of videos, channels, and comments removed from YouTube, including a breakdown of the policies under which this content was removed. It also details how we detect infringing videos (e.g., with automated systems, via user flags) and how many offending videos were removed without any user viewing them.

Our [How Search Works](#) site provides extensive information to anyone interested in learning more about how Google Search works. The site includes information about how we improve search quality, our approach to algorithmic [ranking](#), including publication of our [Search Quality Rater Guidelines](#) which define our goals for Search algorithms. It also includes an overview of our [approach](#) to removing content from search results, which links to our help center and provides a more detailed overview of our policies for [website removals](#), our policies for features like [autocomplete](#), and our [webmaster guidelines](#).

Our annual ['Bad Ads' report](#) outlines the scale of our work to enforce our [advertising policies](#), including the number of ads that were removed, the number of pages that we

stopped showing ads on, the number of advertiser and publisher accounts that were terminated throughout the year, and the number of updates we made to our policies over the course of the year.

Our Threat Analysis Group's [Quarterly Coordinated Influence Operations Bulletin](#) provides information about actions we take against accounts that we attribute to coordinated influence campaigns (foreign and domestic). Additionally, reports made available on the [Google Transparency Report](#) website provide information regarding government requests to remove content from our services, and how the actions of governments and corporations affect privacy, security, and access to information online. We also provide a publicly accessible, searchable, and downloadable [Google Transparency Report of election ad content](#) and spending on our platforms. Given recent concerns and debates about political advertising, and the importance of shared trust in the democratic process, we hope to improve voters' confidence in the political ads they may see on our ad platforms.

We will continue building upon these transparency efforts in the future, as they are an important component of ensuring an informed public dialogue about the role that our services play in society.

8. Are your users required to provide knowing consent to the standard before giving you their personal information, data, and content, which gives value to your platform?

By using our services, users consent to our [Terms of Service](#). We also require users to affirm consent in other circumstances, such as when they setup a Google Account. These Terms of Service help define Google's relationship with users as they interact with our services. For example, these terms include the following topic headings:

- [What users can expect from us](#), which describes how we provide and develop our services
- [What we expect from users](#), which establishes certain rules for using our services
- [Content in Google services](#), which describes the intellectual property rights to the content a user finds in our services — whether that content belongs to the user, Google, or others
- [In case of problems or disagreements](#), which describes other legal rights users have, and what to expect in case someone violates these terms

Besides these terms, we also publish a Privacy Policy which can be found [here](#).

9. Do you coordinate your content moderation standards with other online platforms or competitors? Have you ever discussed or reached an agreement regarding these standards—or the removal of content generally—with any other online platform or competitor? If so, please explain each discussion or agreement, its subject, and the parties thereto.

We work with many talented experts and organizations across the technology industry, government, and civil society to ensure that we are doing everything we can to set the right policies, establish industry best practices, and get ahead of emerging challenges. We do this in part by relying on a community of partners to help us identify content that violates our policies, seeking the advice of subject-matter experts as we craft and update policies.

We also work with other technology companies and industry partners to address challenges that span multiple products and ecosystems by identifying where cooperation would be beneficial and where the resources of a company like Google can help increase the capacity of others. This type of collaboration is often the most effective mechanism for fighting bad actors at scale. Examples of this cross-industry collaboration include the [Global Internet Forum to Counter Terrorism](#) (GIFCT) and the [Technology \(Tech\) Coalition](#) to help eradicate the horrors of Child Sexual Abuse Material (CSAM).

We are optimistic about the progress we have made on our own services, and working together with other companies and governments. It's important that additional frameworks governing online speech be carefully balanced, clear, and fit for purpose. We continue to develop and learn from these collaborations over time and seek more opportunities to develop best practices jointly with partners in industry and government.

10. Do you coordinate the removal of specific content with other online platforms or competitors? If so, please explain the process and what content has been subject to coordinated removal

Through the Tech Coalition, we make cutting-edge technology available to qualifying industry and non-governmental organizations for free in order to help identify, remove, and report illegal CSAM more quickly and at a greater scale. In the last decade-plus, member companies have made progress with the development and roll-out of innovative technology to combat CSAM, and, in 2020, the Coalition announced 'Project Protect,' a renewed investment and

strategic plan to enhance our collective work. 20

Tools like CSAI Match and Content Safety API, which were developed by Google and YouTube engineers, help prioritize potentially illegal content for review while identifying known and never-before-seen CSAM. In addition to being used on our platforms, these tools are also being used by companies like Adobe, Tumblr, and Reddit to aid in the faster identification of potential victims of CSAM, while reducing the toll on content moderators.

Among other important initiatives, GIFCT allows participating companies and organizations to submit hashes, or 'digital fingerprints,' of identified terrorist and violent extremist content to a database so that it can be swiftly removed from all participating platforms. By sharing best practices and collaborating on cross-platform tools we have been able to increase our hash-sharing database to 300,000 hashes.

11. Some of you have removed or threatened to remove the ability of third parties to monetize their content through your advertising platform on the basis of content found in the third party's comments section. What is your policy or standard for such demonetization? Do you believe the same standard should be applied to platforms currently protected by Section 230?

It is no accident that the greatest internet companies in the world were born in the United States. Section 230 is one of the foundational laws that has enabled America's technology leadership and success in the internet sector – allowing freedom of expression to flourish online.

Section 230 helps internet companies address harmful content, including comments. Changes to Section 230 could jeopardize removals of, among other things: terrorist content, spam/malware, scams, misinformation, manipulated media and hate speech, and other objectionable content. The ability to remove harmful but not necessarily illegal content has been particularly important during COVID-19. In just one week, we saw 18M daily malware and phishing emails related to the coronavirus and more than 240M COVID-related daily spam messages. We've removed over 200,000 YouTube videos with dangerous or misleading coronavirus information and 200 million coronavirus ads.

Furthermore, before companies advertise, they want some assurance that publishers' sites are appropriate for their ads. That's where our longstanding content policies come in – they are business-driven policies to ensure ads do not

appear alongside offensive content. For example, a company marketing baby clothes wouldn't want its paid ads to appear alongside violent or mature content. Our content policies cover the entire site where ads are displayed, including [user generated](#) comments sections. In 2011 we published a [blog post](#) making clear that website publishers are responsible for ensuring that all content, "including user-generated content such as forum posts, blog comments or outside feeds," complies with our policies. And in 2017, we published another [blog post](#), specifically about the risks associated with user comments, and suggestions for how publishers can manage those risks.

Section 230 allows us to have content rules that we enforce in an effort to ensure that our platform is safe for our users. We've always had robust policies and are never going to be "neutral" about issues like child abuse, terrorism, harassment, etc.

Sincerely,

A handwritten signature in blue ink, appearing to read 'M. C. Erickson' with a horizontal line extending to the right.

Markham C. Erickson

Vice President, Government Affairs & Public Policy