

Attachment—Additional Questions for the Record

**Subcommittee on Communications and Technology and
Subcommittee on Consumer Protection and Commerce
Joint Hearing on
“Fostering a Healthier Internet to Protect Consumers”
October 16, 2019**

Ms. Katherine Oyama, Global Head of Intellectual Property Policy, Google, Inc.

The Honorable Anna G. Eshoo (D-CA)

- 1. Many of the issues raised today are serious ones that we, as a society, need to work to resolve. How is Google working on these issues today? One example I’m especially worried about is how you deal with the sales of illegal opioids on your platforms today. How does Section 230 play a role in those practices?**

Response: The opioid epidemic is complex and tragically historic, and many organizations and families are working to find solutions. Google has been doing its part specifically to help support efforts relating to prevention, treatment, and recovery. Google is supportive of legislative approaches to increase user awareness about the dangers of opioids, and provide resources to people with substance use disorders, and their families, with help for treating addiction and remaining in recovery. Google actively promotes prescription drug take back programs, like the DEA’s Rx Take Back Day in April and October, as well as permanent disposal efforts run by pharmacies (e.g., CVS, Riteaid, Walgreens, and hospitals) or municipalities (law enforcement), to lessen the risk that legitimate, but unused prescriptions are diverted for misuse.

Google also has a number of efforts to combat the sales of illicit drugs, including illegal opioids, across our platforms. Such content is strictly against many of our products’ policies, such as Ads and YouTube. Google regularly removes policy-violating content relating to controlled substances, including opioids, from its ads platforms and YouTube (with the help of a third party, Legitscript, and our internal manual/automated filters). In March 2020, our Ads platforms will begin blocking marketing of opioid painkillers generally.

- On Web Search, our approach is two-fold. Most links to purported offers to sell opioids without an Rx are actually non-delivery scams. Google has undertaken extensive efforts to combat those listings (e.g., over 2 billion opioid-related spam listings were actioned, delisted or heavily demoted).
- In Spring 2018 we began to work with the FDA to delist--in other words, remove from search results--websites that are the target of an FDA Warning Letter for purporting to sell opioids without a prescription. We have removed hundreds of sites under this process.

Queries seeking information about opioids and treating addiction are much more common than "buy [opioid]" queries by many orders of magnitude. On our Ads platform, we have also taken efforts to prevent abusive addiction treatment advertisements, particularly by patient brokers.

We are able to take these actions precisely because of CDA 230's incentives for responsible content moderation efforts. The law's "good samaritan" provisions were designed to incentivize self-monitoring and facilitate content moderation. Without it, we could not have the types of rigorous policies and programs that we have in place to defend against misuse of our platforms.

The Honorable Kathy Castor (D-FL)

- 1. On June 19, 2019, The Verge published an investigation into one of Facebook's content moderation sites in Tampa, FL, which is operated by the firm Cognizant. The article details allegations of appalling working conditions including sexual harassment, verbal and physical fights, theft, and general filthiness in addition to adverse mental health effects associated with the nature of their work.**

- a. Operationally, how should tech platforms moderate their content? What role should human content moderators play? What role should technology play?**

Response: Our strategy for tackling illegal and potentially harmful content is tailored to each of our platforms. For each of our products, we have a specific set of rules and guidelines that are suitable for the type of platform, how it is used, and the risk of harm associated with it.

We use a mix of machines and people to enforce our policies at scale. Machine learning is allowing us to identify and remove violative content faster than ever before, and our investment in technology enables us to address enforcement of our content policies at scale. For instance, on YouTube, over 93% of videos removed for violating our policies in the third quarter of 2019 were first detected and flagged by our automated detection systems.

However, while machine learning is good at examining large sets of data and looking for content similar to previously removed material, the technology is still evolving. It cannot perfectly identify all violations, and, in general, people are better at making the nuanced, context specific judgments that are often necessary when it comes to evaluating speech and other expression. When videos are flagged by our machines, they are then sent to our trained teams of reviewers, who can analyze the content and take quick action. In many cases, we are able to take action before content has been viewed or accessed by anyone; for instance, during the third quarter of this year, nearly 68% of videos on YouTube that were first flagged by machines were removed before a single view.

Together, this work represents a significant investment. We have over 10,000 people across

Google working on content moderation and removal on our platforms and have invested hundreds of millions of dollars in these efforts.

b. What standard should a private company use to evaluate content? “Quasi constitutional”, a “community standard” established by the company along the lines of other private media, other?

Response: Across our products, we develop ‘rules of the road’ – known as ‘content policies’ or ‘community guidelines’ – which outline what types of content and behaviors are acceptable on each product or service. We aim to make them abundantly clear to all users and content creators and ensure they are easily accessible.

They articulate the purpose and intended use of a given product or service and represent 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

To that end, YouTube’s Community Guidelines prohibit certain categories of material, including sexually explicit content, spam (such as videos trying to trick people to ‘click through’ to another site), hate speech, harassment and incitement to violence. In evaluating videos, we consider purpose and context, including allowing content in some cases if the purpose is educational, documentary, scientific, or artistic in nature. If users are posting content related for this purpose, for instance, we encourage them to be mindful to provide enough information so viewers understand the context, such as through an introduction, voiceover commentary, or text overlays, as well as through a clear title and description. Providing documentary or educational context can help the viewer, and our reviewers, understand why potentially disturbing content sometimes remains live on YouTube.

Sometimes, we make mistakes in our decisions to enforce our policies, which may result in the unwarranted removal of content from our services. To address that risk, wherever possible, we make it clear to creators that we have taken action on their content and provide them the opportunity to appeal that decision. The decision will then be evaluated by a different member of our trust and safety team.

c. Given that private companies are not governed by standards that government would be when it decides not to post content, why do content moderators have to spend so much time reviewing and in such great detail evaluating explicit, violent, or hateful content? What value is there to society and the site owner to work to ensure that such explicit, violent, or hateful content is given every opportunity to be posted?

Response: YouTube is built on the premise of openness. Based on this open platform, millions of creators around the world have connected with global audiences and many of them have built thriving businesses in the process. But openness comes with its challenges, which is why we also have Community Guidelines that we update on an ongoing basis. Most recently, this includes our hate speech policy and our upcoming harassment policy. When you create a place designed

to welcome many different voices, some will cross the line. Bad actors will try to exploit platforms for their own gain, even as we invest in the systems to stop them. As discussed above, we rely on a combination of people and technology to flag inappropriate content and enforce our Guidelines. We continue to improve not only our enforcement processes, but also our policies over time; for instance, in the last year, we made significant improvements to our policies around hate speech and harassment. Problematic content represents a fraction of one percent of the content on YouTube and we're constantly working to reduce this even further.

- d. This explicit, violent, or hateful content often is known to be inconsistent with the tech platform's content bylaws. Why do tech platforms, like Facebook, force content moderators to not only look at but also evaluate in great detail explicit, violent, or hateful content that is often inconsistent with the tech platform's bylaws?**

Response: At YouTube, our Community Guidelines prohibit incitement to violence, hate speech, and graphic content, among other categories of content. And we continue to tighten our policies on what content can appear on our platform, or earn revenue for creators.

We've increased our enforcement teams and invested in powerful new machine learning technology to scale the efforts of our human moderators to take down videos and comments that violate our policies. Human reviewers remain essential to both removing content and training machine learning systems because human judgment is critical to making contextualized decisions on content.

Context is very important for all videos, but it's particularly important when evaluating these categories of content -- and we must be careful to assess potential educational, scientific, newsworthy, or a documentary content that is allowed on YouTube if it includes context. Human moderators have an important role in recognizing contextual nuances.

This context-sensitive approach is important to maintaining the benefits of YouTube as an open platform. Based on this open platform, millions of creators around the world have connected with global audiences and many of them have built thriving businesses in the process. If we were to take the approach of merely relying on automated approaches to relying on content, we would risk over-removing and limiting access to legitimate, non-violative speech.

- e. Should content moderators have more leeway to ban harmful content so they don't have to look at it over such lengthy time periods and evaluate the content in such detail?**

Response: 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, where the content might otherwise violate our policies. This work can be emotionally challenging, especially when the people who review the content against our policies are exposed to some of the most shocking or abhorrent types of content

that exist online. Google is determined to support the wellness of these workers through high wellness standards, verification of vendors' compliance with those standards, and research & technological innovation to promote wellness and mitigate trauma caused by content moderation.

f. What should industry best practices be for treating content moderators? Should Congress play a role in ensuring worker rights in this unique industry? If so, how?

Response: YouTube works closely with our vendor partners to ensure a standardized, comprehensive wellness program is delivered to all agents reviewing sensitive content. All agents have access to onsite counseling services, wellness breaks, resilience training, and organized wellness activities to ameliorate the mental health impact of reviewing sensitive content. We are competitive in the space of compensation and continue to evaluate pay relative to industry standards, and ensure those who are reviewing more sensitive content are compensated accordingly.

g. Is it common practice among tech platforms to use contractors to conduct content moderation for their sites? Why do some tech platforms use contractors to conduct content moderation for their sites? Should tech platforms do this?

Response: While we cannot speak to the practices of other companies, content moderation at Google and YouTube is primarily managed by Trust & Safety teams that sit across the company. These teams work with our in-house legal and policy departments on escalations and also oversee the vendors we hire to help us scale our operations. These teams are made up of engineers, content reviewers, and others who work across Google to address content that violates any of our policies. They are made up of a mix of full-time employees and contractors.

The Honorable Lisa Blunt Rochester (D-DE)

1. At the October 16, 2019, joint hearing, you provided commitments that Google will disclose information on diversity of your content moderators and issues with hiring diverse content moderator teams. Please provide that information to the Energy and Commerce Committee and my office.

Response: As you may know, Google was the first large technology company to publish workforce diversity data in 2014 (all of this information is available at [google.com/diversity](https://www.google.com/diversity)). We are committed to sharing our numbers every year, and 2019 was no different. We provide one of the most transparent data sets in our industry and our report allows the public to view a demographic representation of our employees, toggle between different job functions & leadership roles, and even view this data intersectionality. Beyond diversity of race and gender, we also included veterans, people with disabilities, and LGBTQ people in this data for the first time last year. To your question, the Google moderators are a mix of tech and non-tech employees who are representative of Google as a whole. Our 2019 Diversity Annual

Report can be found at <https://diversity.google/annual-report/> .

We believe it is important to be transparent about our challenges and key learnings in this arena. Our original decision to release our workforce (diversity) numbers led to other companies following suit. Google stands firm in its commitment to foster dialogue and to drive impact on this important issue.

2. What can the federal government do to improve the capacity and ability to effectively moderate online content, including technological research?

Response: Our strategy for tackling illegal and harmful content is tailored to each of our platforms. Across our products, our teams tackle a huge spectrum of online abuse, from scams, like the email from a ‘relative’ stranded abroad needing a bank transfer to get home safely, to abhorrent content, including child sexual abuse material (CSAM) online.

It is important to note, however, that rogue, off-shore sites that promote illegal and harmful activity are often commercial in nature -- they’re running scams or selling illegal goods to make a profit. A critical way to stop these sites is to cut off their money supply (i.e. payments processors, advertising services, etc.). Removing the site from Search, for example, doesn’t remove it from the Web, but money is the oxygen that many bad actors need to survive. Cut that off, and many of them will go away.

Furthermore, CDA 230’s civil law framework does nothing to alter existing law enforcement tools and liability framework for violations of federal criminal laws, which are expressly exempted from the scope of the Communications Decency Act. This is why combating difficult problems of illegal content requires a response from government, individuals, and organizations, often working in partnership. We stand ready to work in concert on these issues.

The Honorable Tom O’Halleran (D-AZ)

- 1. Ms. Oyama, as written in statute, Section 230 has “good Samaritan” language to incentivize online platforms to take actions “*in good faith to restrict access to or the availability of*” harmful content.**

Many platforms have established content or use of service policies to specify what behavior is allowed by the service, while others employ artificial intelligence formulas to automatically filter user-generated content. Some platforms also hire human content moderators to review and remove content posted by users on its platforms that is considered harmful, violent, or graphic. These content reviewers often suffer from Post-Traumatic Stress Disorder (PTSD).

- a. **What more can be done by the government and industry to ensure sufficient mental health services are made available to human content reviewers?**

Response: YouTube works closely with our vendor partners to ensure a standardized, comprehensive wellness program is delivered to all agents reviewing sensitive content. All agents have access to onsite counseling services, wellness breaks, resilience training, and organized wellness activities to ameliorate the mental health impact of reviewing sensitive content. We are competitive in the space of compensation and continue to evaluate pay relative to industry standards, and ensure those who are reviewing more sensitive content are compensated accordingly.

The Honorable Greg Walden (R-OR)

1. **At the hearing, Rep. Bilirakis asked EFF whether they have argued for including language mirroring legislation in trade deals explicitly for the purpose of “baking” language into an agreement to protect the statute domestically.**

For the record, Yes or No: Is including such 230-like language in trade agreements an attempt to preclude us – the committee of jurisdiction – from revisiting the statute? Do you see the intent of including such 230-like language in trade agreements is to ensure that we may not revisit the statute?

Response: No.

2. **During the hearing, you mentioned that Google has a tool to tag copyrighted works so that upon re-upload they can be stopped before spreading, in the case of illegal content like violent extremism. Can that tool be more widely deployed, both across different platforms, like Reddit, and to cover more content besides copyrighted works? If not, why not?**

Response: Google offers a wide range of services on many platforms. Our strategy and technologies for tackling illegal and harmful content can differ based on the nature of the services and the technology of the platforms. There is no one size fits all approach in these matters. Generally speaking, in addition to content flagged by our users, we develop and deploy cutting-edge technology to proactively identify, remove, and block offending content. In addition, across our products, our human review teams tackle a huge spectrum of online abuse, from scams, like the email from a ‘relative’ stranded abroad needing a bank transfer to get home safely, to abhorrent content, including child sexual abuse material (CSAM) online.

That said, we work with other companies in the industry where we can. For example, in the CSAM space, for over a decade we have been using PhotoDNA, along with other complementary systems, to identify, remove, and report copies of CSAM present in still images and share the digital fingerprints, or ‘hashes.’ We contribute new hashes to, and receive hashes from other platforms via a hash database maintained by NCMEC. In 2015,

YouTube engineers created CSAI Match, which can be used to identify uploaded videos that contain known CSAM. It is used for videos uploaded to YouTube and on Livestreams. CSAI Match is used by companies and organizations like Adobe, Reddit, Tumblr, among others. Building on our previous work to develop and extend machine learning systems, and share technology with industry in 2018, Google engineers launched the Content Safety API. This tool helps us to find and report new CSAM that was not possible using hash matching alone. It also allows us to identify CSAM at scale, by prioritizing for manual review content most likely to constitute CSAM. We make this technology available for free to industry and NGO partners.

- a. **In a response to Rep. Walberg, you pointed to Content ID as one tool Google uses for piracy on its platform. Understanding Section 230 immunity already has an exception for intellectual property in statute, we are interested in better understanding how these tools could be applied to other types of content on Googles' platforms in situations where such content is either criminally illegal (and therefore not covered by section 230) or "obscene, lewd, lascivious, filthy, excessively violent, harassing, or otherwise objectionable."**

Are there any other tools for copyrighted content that Google uses besides Content ID, Content Verification Program, or Content Match? If so, please provide the name and a short description of how the tool operates.

Response: Google's approach to combating piracy and illegal content on its platform goes beyond just the use of individual tools and policies. Rather, we believe that successfully decreasing incidents of copyright infringement involves a multi-pronged approach:

1. Create More and Better Legitimate Alternatives

Piracy often arises when consumer demand goes unmet by legitimate supply. The best way to battle piracy is with better, more convenient, legitimate alternatives to piracy, which can do far more than attempts at enforcement can. By developing products with compelling user experiences like Google Play Music and YouTube, Google helps drive revenue for creative industries and steer people toward legitimate alternatives. Google also supports the larger ecosystem by providing the cloud infrastructure that other legitimate services depend on to deliver fast, reliable streaming to their customers.

2. Follow the Money

Rogue sites that specialize in online piracy are commercial ventures, which means that one effective way to combat them is to cut off their money supply. Google is a leader in rooting out and ejecting rogue sites from our advertising and payment services, and we help establish best practices across the industry.

3. Be Efficient, Effective, and Scalable

Google strives to implement anti-piracy solutions that work at scale. For example, as early as 2010, Google began making substantial investments in streamlining the copyright removal process for search results. As a result, these improved procedures allow Google to process copyright removal requests for search results at the rate of millions per week.

4. Guard Against Abuse

Fabricated copyright infringement allegations can be used as a pretext for censorship and to hinder competition. Google is committed to ensuring that it detects and rejects bogus infringement allegations, such as removals for political or competitive reasons, even as it battles online piracy

5. Provide Transparency

Google is committed to providing transparency. In our external Transparency Report, Google discloses the number of requests it receives from copyright owners and governments to remove information from its services to inform ongoing discussions about online content regulation.

For more information, our 2018 "[How Google Fights Piracy](#)" report explains the programs, policies, and technology we put in place to combat piracy online and ensure continued opportunities for creators around the world.

b. I understand Google makes Content ID available to companies that, “own exclusive rights to a substantial body of original material that is frequently uploaded by the YouTube user community.” Can you please describe how Google determines what qualifies as a “substantial body?”

Response: Content ID access requires users to make a certain level of operational investment, without which other rights holders could have their rights impaired and lawful expression could be inappropriately impacted. We maintain [minimum standards](#) for Content ID access to preserve accuracy and quality and evaluate candidates on an individual basis. Some of the factors we consider when matching a rightsholder to one of the tools in our copyright management suite are:

- Whether the company or rightsholder holds exclusive copyrights in their content. Often different entities own broadcast rights, physical resale, and online distribution rights and we must take instructions from the entity that controls online distribution rights.
- Whether the company or rightsholder has a history of sending YouTube complete and valid copyright takedown requests to remove allegedly infringing content.
- Whether the company or rightsholder owns the rights to a variety of content that’s frequently uploaded to YouTube. There is no set requirement for the number of uploads of a rightsholder’s content that must have occurred to grant access, but we

have found that those with significant experience managing their content on the platform are most able to navigate the complexity of copyright licensing on Content ID.

- i. What is the threshold to qualify for Content ID? Is it measured in number of copyrighted rights, the value of those rights, the amount of ad revenue those rights bring to Google’s platform, or another metric? If “other,” please elaborate on what that metric is.**

Response: As discussed in the previous question, Content ID access requires users to make a certain level of operational investment, without which other rights holders could have their rights impaired and lawful expression could be inappropriately impacted. We maintain [minimum standards](#) for Content ID access to preserve accuracy and quality and evaluate each individually. However, neither “the value of those rights”, nor “the amount of ad revenue” play a factor in the evaluation.

That said, YouTube endeavors to make Content ID available to as many rightsholders as possible without sacrificing accuracy and quality. We have found that even major rightsholders can mismanage Content ID, however, resulting in unfair or inaccurate claims on legally-uploaded content. Thus, we try to ensure that we grant Content ID to responsible, enterprise-scale rights holders committed to training and management of the powerful tool. For creators and rightsholders who require more simplified copyright tools, we’ve developed the Copyright Match Tool and the Content Verification 1 Program.²

YouTube's recently launched a Copyright Match Tool, which uses the power of the Content ID matching system system to find re-uploads of creator videos on YouTube. Instead of uploading a reference file to YouTube, creators that upload a video to YouTube are shown subsequent uploads of their videos. Creators and content owners can then review the matching videos and file bulk takedowns for any they wish to remove. They can also choose to contact the uploader. In this way, creators remain in control of the works they create on YouTube. This new tool greatly simplifies copyright management so that creators can focus their time on making great videos.

1 Available at <https://support.google.com/youtube/answer/7648743> .

2 Available at <https://support.google.com/youtube/answer/6005923?hl=en> .

- ii. I understand that Google allows some aggregators access to Content ID, and encourages some smaller content creators that Google determines does not qualify for access to Content ID to hire an aggregator. Why?**

Response: We have found that creators often prefer to have an experienced aggregator that specializes in this work to manage it, which can be complex and time-consuming. There is also a significant investment in initially learning how to use copyright management systems appropriately. YouTube offers certification courses to train individuals on the use of Content

ID. We have found that those who are not certified are at much greater risk of making mistakes which can have a significant negative impact on other rightsholders and creators. Even after certification, Content ID requires daily engagement to resolve complex issues.

In addition, in many cases, creators, such as independent musicians, may have already assigned their rights to a music label or hired an aggregator in order to distribute their works on other platforms such as Spotify or iTunes. Despite this, we often learn that a small creator is unaware that an aggregator is already representing their rights on YouTube.

1. Must this aggregator have, “exclusive rights” to their clients’ content, or does Google make an exception for aggregators?

Response: There is no requirement that the aggregator own exclusive rights to a client’s works. The client must own the right to manage the work online and is free to designate the aggregator to represent those rights.

3. We get caught up talking a lot about scale and resources in the conversation about litigation risk from modifying Section 230. One of the benefits of scale, as Google has shown time and again, is that vast amounts of data is needed to create more sophisticated machine learning and algorithms. Do you see the hundreds of hours of video uploaded to YouTube every minute as an asset to be testing the types of technologies Dr. Farid mentions in his testimony?

Response: As with all technologies, the real hard work is in turning raw machine learning models into great products and features that solve problems. Data is just one of many important factors in developing a smart and useful model. Quantity of data is not the main key to success. In fact, with worldwide reach, standardized technology and communications protocols, and rapid price decreases in things like cloud platforms and storage, data has become easy to obtain and create. In fact, large quantities of data can even be paralyzing if not properly understood. Success in machine learning, as in almost any other area of tech, requires selection and execution. Data quality often counts way more than data quantity.

Beyond the data that defines what machines learn, there is a great deal of work to be done in how machines learn. Lots of researchers are looking to find ways for machines to learn effectively with smaller amounts of data. Google has a big interest in addressing this open question of research - we would benefit a great deal if machines learned faster with less data. This is why we have made datasets available for others to train their own models upon -- for example our [Open Images Dataset](#) or our [AVA Video Dataset](#) for Human Action Understanding which we used in training some of our own machine learning models.

4. When law enforcement provides information of actual criminal activity, should platforms be required to act to remove it? How quickly should they respond, and should that vary by provider, for example should smaller providers have more time to respond?

Response: Google appreciates that law enforcement agencies face significant challenges in protecting the public against crime and terrorism. We engage in ongoing dialogue with law enforcement agencies to understand the threat landscape and respond to threats that affect the safety of our users and the broader public.

While we cannot speak for other companies, courts and government agencies around the world regularly request that we remove information from Google products. Government bodies ask us to remove or review content for many reasons. Some requests allege defamation, while others claim that content violates local laws. The laws surrounding these issues vary by country/region. Since each request can differ greatly, our teams evaluate and review the content in context in order to determine whether or not content should be removed due to violation of local law or our content policies. We release a transparency report on these requests which can be found [here](#).

5. Does the Google Play store need section 230 protections, or because you only allow developers/apps in the Play store that meet your Terms of Service, you could create accountability within that portion of Google without Section 230?

Response: The Play Developer Distribution Agreement and incorporated Developer Program Policies do give Google the right to determine what content we do and don't want on Play and to take action against apps and developers that don't meet those policies. And while we rely on the agreement and policies in taking action, CDA 230 remains an important part of enforcing our rigorous policies and programs. Section 230 helps ensure that when we take action against a developer or an app (or remove an app, which might otherwise lead to legal complaints from users), we don't end up in years of litigation with costly discovery. The protections that 230 affords makes it possible for us to have the types of rigorous policies and programs that we have in place to defend against misuse of our platforms.

6. Recently, Mark Zuckerberg said Facebook's size is the only reason it can effectively fight election interference, citing that, quote

"it's why Twitter can't do a good of job as we can...I mean, they face, qualitatively, the same types of issues. But they can't put in the investment. Our investment on safety is bigger than the whole revenue of their company."

He's right, size matters, and we expect a lot more from those who have a "bigger sword." In your estimation, is Google doing enough to earn it's 230 protection?

Response: We cannot speak for Mr. Zuckerberg, but at Google, we know that combating difficult problems of illegal content requires a response from government, individuals, and organizations, often working in partnership. We have not waited for legislation to act in tackling illegal or harmful content. We are committed to doing our part.

Our strategy for tackling illegal and harmful content is tailored to each of our platforms. Across our products, our teams tackle a huge spectrum of online abuse, from scams, like the email from a 'relative' stranded abroad needing a bank transfer to get home safely, to

abhorrent content, including child sexual abuse material (CSAM) online.

For each product, we have a specific set of rules and guidelines that are suitable for the type of platform, how it is used, and the risk of harm associated with it. These approaches range from clear community guidelines, with mechanisms to report content that violates them, to increasingly effective artificial intelligence (AI) and machine learning that can facilitate removal of harmful content before a single human user has been able to access it to blocking and removing content when we are notified that a video violates our guidelines. We now have over 10,000 people across Google working on content moderation and removal on our platforms and have invested hundreds of millions of dollars in these efforts.

The Honorable Richard Hudson (R-NC)

- 1. One of the best parts of my job is having the privilege of representing the brave men and women who are stationed at Fort Bragg. Additionally, my district represents one of the fastest growing veteran populations in the country. I take it as my responsibility to advocate for them in everything I do. As you all are aware, the opioid epidemic is something that has ravaged our country and disproportionately affected veterans. One of the underlying issues in this area is the availability of these drugs and how easy it can be for an individual to gain access to them. Unfortunately, we have seen that these drugs are often available through illegal online sales that help fuel this crisis.**

- a. Can you please explain how your company monitors content on your platform to ensure the illegal sale of opioids does not occur?**

Response: The opioid epidemic is complex and tragically historic, and many organizations and families are working to find solutions. Google has been doing its part specifically to help support efforts relating to prevention, treatment, and recovery. Google is supportive of legislative approaches to increase user awareness about the dangers of opioids, and provide resources to people with substance use disorders, and their families, with help for treating addiction and remaining in recovery. Google actively promotes prescription drug take back programs, like the DEA's Rx Take Back Day in April and October, as well as permanent disposal efforts run by pharmacies (e.g., CVS, Riteaid, Walgreens, and hospitals) or municipalities (law enforcement), to lessen the risk that legitimate, but unused prescriptions are diverted for misuse.

Google also has a number of efforts to combat the sales of illicit drugs, including illegal opioids, across our platforms. Such content is strictly against many of our products' policies, such as Ads and YouTube. Google regularly removes policy-violating content relating to controlled substances, including opioids, from its ads platforms and YouTube (with the help of a third party, Legitscript, and our internal manual/automated filters). In March 2020, our Ads platforms will begin blocking marketing of opioid painkillers generally.

- On Web Search, our approach is two-fold. Most links to purported offers to sell opioids without an Rx are actually non-delivery scams. Google has undertaken

extensive efforts to combat those listings (e.g., over 2 billion opioid-related spam listings were actioned, delisted or heavily demoted).

- In Spring 2018 we began to work with the FDA to delist--in other words, remove from search results--websites that are the target of an FDA Warning Letter for purporting to sell opioids without a prescription. We have removed hundreds of sites under this process.

Queries seeking information about opioids and treating addiction are much more common than "buy [opioid]" queries by many orders of magnitude. On our Ads platform, we have also taken efforts to prevent abusive addiction treatment advertisements, particularly by patient brokers.

b. When you find such content on your platform, do you engage law enforcement in addition to removing the content from your platform?

Response: Google is proud to work with the FDA, the DEA, and other regulatory and law enforcement agencies involved in enforcing laws and regulations concerning the sale of illicit drugs online. Over the past few years, Google has referred thousands of potential rogue pharmacies to law enforcement, including the FDA's Office of Criminal Investigations and, more recently, the DEA's Special Operations Division (to whom we have sent a number of referrals this year alone). We have also referred to the FDA pharmaceutical advertisers who have sought to evade Google's filters for prescription drug advertisements. Google also takes a proactive role in assisting the FDA and other law enforcement agencies in investigative efforts. For instance, in October 2012, Google voluntarily participated in the successful "Operation Pangea V", in which the FDA, in partnership with international regulatory and law enforcement agencies, as well as other companies around the world, took collective action against more than 4,100 internet pharmacies online.

The Honorable Tim Walberg (R-MI)

1. **In response to one of my questions during the hearing, you talked about Google's risk engine and how it is very, squarely in-line with Google's interest to not serve ads next to illegal activity. Indeed, you go on to say, "our advertisers only want to be serving good ads to good content," additionally mentioning that over 2 billion ads are stricken every year, "before they're able to ever hit any page...": making the point that you can do so at scale. This of course, makes sense, given this is where Google makes its money. But if Google has proven it can identify harmful or illegal content in the ads space to protect its bottom line, why then can't Google apply this same process to:**

a. Search

b. Images

- c. YouTube**
- d. Google News**
- e. Google Assistant**
- f. Groups**

Response: Our strategy for tackling illegal and potentially harmful content is tailored to each of our platforms. For each of our products, we have a specific set of rules and guidelines that are suitable for the type of platform, how it is used, and the risk of harm associated with it. These approaches range from clear policies and community guidelines, with mechanisms to report content that violates them, to increasingly effective artificial intelligence (AI) and machine learning that can facilitate removal of harmful content before a single human user has been able to access it. We also now have over 10,000 people across Google working on content moderation and removal on our platforms and have invested hundreds of millions of dollars in these efforts.

We use a mix of machines and people to enforce our policies at scale. Machine learning is allowing us to identify and remove violative content faster than ever before. And our investment in technology enables us to address enforcement of our content policies at scale. Machines flag suspect videos for review by trained teams, who can analyze the content and take quick action. This system has had a major impact on the way we tackle harmful content, and has helped our human reviewers remove content more quickly.