

**Written Testimony of Alexandra N. Veitch
Director, YouTube Government Affairs & Public Policy, Americas**

**Senate Judiciary Committee
Subcommittee on Privacy, Technology & the Law**

**“Algorithms and Amplification: How Social Media Platforms’ Design Choices Shape Our
Discourse and Our Minds”**

April 27, 2021

Introduction

Chairman Coons, Ranking Member Sasse, and distinguished members of the subcommittee:

Thank you for the opportunity to appear before you today. My name is Alexandra Veitch, and I am the Director of Public Policy for the Americas at YouTube. As part of my role, I lead a team that advises the company on public policy issues around online, user-generated content.

At YouTube, we believe that the Internet has been a force for creativity, learning, and access to information. Supporting this free flow of ideas is at the heart of our mission to give everyone a voice and show them the world. We have built and continue to improve YouTube to empower users to access, create, and share information like never before; this has enabled billions to benefit from a bigger, broader understanding of the world. In addition, our platform has created economic opportunities for small businesses across the country and around the world, and we have provided artists, creators, and journalists a platform to share their work. Over the last three years, we’ve paid more than \$30 billion to creators, artists, and media companies around the world¹. And according to an Oxford Economics report², YouTube's creative ecosystem supported the equivalent of 345,000 full time jobs in 2019 in the United States.

¹ <https://blog.youtube/inside-youtube/letter-from-susan-our-2021-priorities/>

²

<https://blog.oxfordeconomics.com/content/from-opportunity-to-impact-assessing-the-economic-societal-and-cultural-benefits-of-youtube-in-the-us>

Over the years, we have seen more and more people come to YouTube to share their experiences and understand their world more deeply. This is especially true when it comes to learning new skills, participating in civic engagement, and developing informed opinions about current events. With so many users around the world looking to YouTube for information, we have a responsibility to provide a quality experience and support an informed citizenry. Over the past several years, responsibility has been a critical area of investment across our company. We have focused extensively on developing policies and building product solutions to live up to this responsibility while preserving the opportunity of an open platform. And we work continuously to identify areas where we can do more. In my testimony today, I will (1) explain how we think about algorithms, (2) discuss our approach to responsibility and how technology supports this work, (3) highlight our efforts to provide more transparency and visibility into how YouTube works, and (4) illustrate how our products protect users.

How YouTube thinks about algorithms

YouTube is a multi-faceted video-sharing platform enjoyed by billions of consumers and creators. The popularity of our service unlocks business opportunities for creators and helps businesses grow their reach. With this volume of users and economic engine for creators comes a significant responsibility to protect those users. With more than 500 hours of video uploaded to YouTube per minute, enabling users to easily find content they are looking for and will enjoy while protecting them from harmful content simply would not be possible without the help of technology.

Because of the importance of algorithms in the YouTube user experience, we welcome the opportunity to clarify our approach to this topic. In computer science terms, an algorithm is a set of instructions that direct a computer to carry out a specific task. An algorithm can be simple—asking a computer to calculate the sum of two numbers—or extremely complex, such as machine learning algorithms that consistently refine their ability to accomplish the goal for which they were programmed. An algorithm can manage a few inputs or nearly limitless inputs, and they can do one thing or perform a number of functions at once. Nearly everything that people do today on their devices is made possible by algorithms.

YouTube uses machine learning techniques to manage and moderate content on YouTube. YouTube's machine learning systems sort through the massive volume of content to find the most relevant and useful results for a user's search query, to identify opportunities to elevate authoritative news, and to provide a user with additional context via an information panel if appropriate. We also rely on machine learning technology to help identify patterns in content that may violate our Community Guidelines or videos that may contain borderline

content—content that comes close to violating our Community Guideline but doesn't quite cross the line. These systems scan content on our platform 24/7, enabling us to review hundreds of thousands of hours of video in a fraction of the time it would take a person to do the same. For example, more than 94% of the content we removed between October and December of 2020 was first flagged by our technology. This underscores just how critical machine learning is for content moderation.

Another area where we use machine learning is for recommendations. Recommendations on YouTube help users discover videos they may enjoy, and they help creator content reach new viewers and grow their audience across the platform. We share recommendations on YouTube's homepage and in the "Up next" section to suggest videos a user may want to watch after they finish their current video. Our recommendation systems take into account many signals, including a user's YouTube watch and search history (subject to a user's privacy settings) and channels to which a user has subscribed. We also consider a user's context—such as country and time of day—which, for example, helps our systems show locally relevant news, consistent with our effort to raise authoritative voices. Our systems also take into account engagement signals about the video itself—for example, whether others who clicked on the same video watched it to completion or clicked away shortly after starting to view the video. It is important to note that, where applicable, these signals are overruled by the other signals relating to our efforts to raise up content from authoritative sources and reduce recommendations of borderline content and harmful misinformation—even if it decreases engagement.

We also empower our users by giving them significant control over personalized recommendations, both in terms of individual videos as well as the way that watch and search history may inform recommendations. Users control what data is used to personalize recommendations by deleting or pausing activity history controls. Signed out users can pause and clear their watch history, while signed in users can also view, pause, and edit watch history at any time through the YouTube history settings. Clearing watch history means that a user will not be recommended videos based on content they previously viewed. Users can also clear their search history, remove individual search entries from search suggestions, or pause search history using the YouTube History settings.

In-product controls enable users to remove recommended content—including videos and channels—from their Home pages and Watch Next. Signed in users can also delete YouTube search and watch history through the Google My Account settings, set parameters to automatically delete activity data in specified time intervals, and stop saving activity data

entirely. We also ask users directly about their experiences with videos using surveys that appear on the YouTube homepage and elsewhere throughout the app, and we use this direct feedback to fine-tune and improve our systems for all users.

YouTube's Pillars of Responsibility: the 4Rs

Responsibility is our number one priority at YouTube. Some speculate that we hesitate to address problematic content because it benefits our business; this is simply false. Failure to consistently take sufficient action to address harmful content not only threatens the safety of our users and creators, it also threatens the safety of our advertising partners' brands. Our business depends on the trust of our users and our advertisers. This is why we have made significant investments over the past few years in teams and systems that protect YouTube's users, partners, and business. Our approach towards responsibility involves 4 "Rs" of responsibility, described in detail below.

REMOVE VIOLATIVE CONTENT: Our Community Guidelines provide clear, public-facing, guidance on content that is not allowed on the platform. These include policies against spam, deceptive practices, scams, hate, harassment, and identity misrepresentation and impersonation. We remove content that violates our policies as quickly as possible, and removed videos represent a fraction of a percent of total views on YouTube. We work continuously to shrink this even further through improved detection and enforcement, relying on a combination of technology and people.

We are dedicated to providing access to information and freedom of expression, but YouTube has always had clear and robust content policies. For example, we have never allowed pornography, incitement to violence, or content that would harm children. Harmful content on our platform makes YouTube less open, not more, by creating a space where creators and users may not feel safe to share. That's why our policy development team systematically reviews and updates all of our policies to ensure that they are current, keep our community safe, and preserve openness. They frequently consult outside experts and YouTube creators during the process, and consider regional differences to ensure proposed changes can be applied fairly and consistently around the world.

Our COVID-19 Medical Misinformation Policy³ represents one such example of YouTube working closely with experts. Over the course of the last year, we have worked with and relied on information from health authorities from around the world to develop a robust policy

³ <https://support.google.com/youtube/answer/9891785>

anchored in verifiably false claims tied to real world harm. Our policy addresses false and harmful claims about certain treatments and public health measures, as well as misinformation about COVID-19 vaccines. Each claim we prohibit has been vetted as verifiably false by the consensus of global health authorities, including the CDC. This policy has evolved alongside misinformation trends about the pandemic, and we have invested significant time and resources into carefully developing the policy and creating the necessary training and tools required to enforce it.

We adapted this policy over time to address the challenges of this pandemic. We began to remove content for COVID-19 misinformation in March 2020, under provisions of our policy prohibiting Harmful and Dangerous content. But as the pandemic progressed, we developed a fulsome and separate COVID-19 misinformation policy. In October 2020, we further expanded the policy to include vaccine misinformation about COVID vaccines. Since March 2020, we have vigorously enforced our COVID-19 misinformation policy to protect our users, removing 900,000 videos worldwide. And in the fourth quarter of 2020, we removed more than 30,000 videos for violating the vaccine provisions of our COVID-19 misinformation policy.

Once we have implemented a policy, we rely on a combination of people and technology to enforce it. Machine learning plays a critical role in content moderation on YouTube, and we deploy it in two key ways: proactively identify and flag harmful content, and automatically remove content that is very similar to what has been previously removed. In both cases, data inputs are used to train the systems to identify patterns in content—both the rich media content in videos, as well as textual content like metadata and comments—and our systems then use those patterns to make predictions about new examples to match. Machine learning is well-suited to detect patterns, which also helps us to find content similar to other content we have already removed, even before it is ever viewed. We sometimes use hashes (or “digital fingerprints”) to catch copies of known violative content before they are even made available to view. The systems then automatically remove content only where there is high confidence of policy violation—e.g., spam—and flag the rest for human review.

Machine learning is critical to keeping our users safe. In the fourth quarter of 2020, of the 9.3 million videos removed for violating our Community Guidelines, 94% of those videos were first flagged by machine detection. But the human review piece here is critical as well: machines are effective for scale and volume but are not able to analyze and evaluate context, whereas human reviewers allow us to evaluate context and consider nuance when enforcing our policies. Once our machine learning systems flag a potentially violative video, reviewers then remove videos that are violative while non-violative videos remain live. These decisions are in

turn used as inputs to improve the accuracy of our technology so that we are constantly updating and improving the system's ability to identify potentially violative content. In addition, when we introduce a new policy or alter an existing one, it takes our systems a bit of time to catch up and begin to detect relevant content. As we explained when we updated our hate speech policy, our enforcement of new policies improves quarter over quarter.

But as with any system, particularly operating at scale like we do, we sometimes make mistakes, which is why creators can appeal removal decisions. Creators are notified when their video is removed, and we provide a link with simple steps to appeal the decision. If a creator chooses to submit an appeal, it goes to human review, and the decision is either upheld or reversed. And we are transparent about our appeals process. As reported in our most recent Transparency Report, in Q4 2020, creators appealed a total of just over 223,000 videos. Of those, more than 83,000 were reinstated.

We also recently added a new metric to the YouTube Community Guidelines Enforcement report⁴ known as Violative View Rate (VVR). This metric is an estimate of the proportion of video views that violate our Community Guidelines in a given quarter (excluding spam). Our data science teams have spent more than two years refining this metric, which we consider to be our North Star in measuring the effectiveness of our efforts to fight and reduce abuse on YouTube. In Q4 of 2020, YouTube's VVR was 0.16-0.18%, meaning that out of every 10,000 views on YouTube, 16-18 come from violative content. We have added historical data for this metric to our Transparency Report, showing that, since Q4 of 2017, we have seen a 70% drop in VVR. This reduction is due in large part to our investments in machine learning to identify potentially violative content. Going forward, we will update this metric quarterly alongside our regular data updates.

RAISE UP AUTHORITATIVE VOICES: YouTube is a source for news and information for people around the world—whether about events unfolding in local communities or more existential global issues like climate change. Not all queries are the same. For topics like music or entertainment, relevance, newness, and popularity are most helpful to understand what people are interested in. But for subjects such as news, science, and historical events, where accuracy and authoritativeness are key, the quality of information and context are paramount.

Our search and recommendations systems are designed to raise up authoritative voices in response to user queries, especially those that are “news-y” or related to topics prone to

⁴ <https://transparencyreport.google.com/youtube-policy>

misinformation. In 2017, we started to prioritize authoritative voices such as local and national news sources for information queries in search results and “watch next” recommendation panels. This work continued with the addition of a short preview of text-based news articles in search results on YouTube, along with a reminder that breaking and developing news can rapidly change. And in 2018, we introduced Top News and Breaking News sections to highlight quality journalism⁵. Our work here is far from done, but we have seen significant progress in our efforts to raise authoritative voices on YouTube. Globally, authoritative news watchtime grew by more than 85% from the first half of 2019 to the first half of 2020, with a 75% increase in watchtime of news in the first 3 months of 2020 alone. And during the 2020 U.S. elections, the most popular videos about the election came from authoritative news organizations. On average 88% of the videos in top 10 search results related to elections came from authoritative news sources.

Authoritativeness is also important for topics prone to misinformation, such as videos about COVID vaccines. In these cases, we aim to surface videos from authoritative news publishers, public health authorities, and medical experts. Millions of search queries are getting this treatment today and we continue to expand to more topics and countries. In addition, in April 2020, we expanded our fact-checking panels in YouTube search results to the U.S., providing fresh context in situations where a news cycle has faded but where unfounded claims and uncertainty about facts are common. These panels highlight relevant, third-party fact-checked articles above search results for relevant queries, so that our viewers can make their own informed decision about claims made in the news.

We also recognize that there may be occasions when it is helpful to provide viewers with additional context about the content they are watching. To that end, we have a variety of information panels that provide context on content relating to topics and news prone to misinformation, as well as the publishers themselves. For example, a user viewing a video about climate change—regardless of the point of view presented in the content—will see an information panel providing more information about climate change with a link to the relevant Wikipedia article.

Information panels provide critical context as well as point users to reliable sources of authoritative information. For example, when a user in the U.S. watches a video about COVID-19, we display an information panel that points to the CDC’s official resource for information about COVID-19 and the Google search results page with health information from

⁵ <https://blog.youtube/news-and-events/building-better-news-experience-on>

the CDC and local statistics and guidance. Beginning last week, when a U.S. user watches a video about COVID-19 vaccines, we show a panel that points to the CDC's online resource for vaccine information, with an additional link to the Google search results page with local information about vaccination. To date, our COVID-19 information panels have received more than 400 billion views.

For the U.S. 2020 election, we provided a range of new information panels in addition to our existing panels to provide additional context around election-related search results and video content. For example, when viewers searched for specific queries related to voter registration on YouTube, they were shown an information panel at the top of the page that linked to Google's "How to register to vote" feature for their state. When a viewer searched for 2020 presidential or federal candidates on YouTube, we surfaced an information panel with information about that candidate—including party affiliation, office, and when available, the official YouTube channel of the candidate—above search results. We also provided an election results info panel, the content of which was flexible in order to keep pace with new developments and key milestones along the road to inauguration. These panels were collectively shown more than 8 billion times.⁶

REDUCE THE SPREAD OF BORDERLINE CONTENT:

While we have strong and comprehensive policies in place that set the rules for what we don't allow on YouTube, we also recognize that there's content that may be problematic but doesn't violate our policies. Content that comes close to violating our Community Guidelines but does not cross the line—what we call "borderline content"—is just a fraction of 1 percent of what is watched on YouTube in the United States. We use machine learning to reduce the recommendations of this type of content, including potentially harmful misinformation.

In January 2019, we announced changes to our recommendations systems to limit the spread of this type of content. These changes resulted in a 70 percent drop in watchtime on non-subscribed recommended content in the U.S. that year. We saw a drop in watchtime of borderline content coming from recommendations in other markets as well. While algorithmic changes take time to ramp up and consumption of borderline content can go up and down, our goal is to have views of non-subscribed, recommended borderline content below 0.5%. We seek to drive this number to zero, but no system is perfect; in fact, measures intended to take this number lower can have unintended, negative consequences, leading legitimate speech to

6

<https://kstatic.googleusercontent.com/files/2dfe4c0bde90ab01c566d50ec70c8d74bc60f34ac5a409c28c5d5b1bac8aae93ee1187562a0e12f672d32df8d5cd8fa8fe52c95ee96293bc92b0ca7f9301d0b3>

not be recommended. As such, our goal is to stay below the 0.5% threshold, and we strive to continually improve over time.

This content is ever-evolving and it is challenging to determine what content may fall into this category. This is why we rely on external evaluators located around the world to provide critical input on the quality of a video, and these evaluators are trained with public guidelines⁷. Each evaluated video receives up to nine different opinions and some critical areas require certified experts. For example, medical doctors provide guidance on the validity of videos about specific medical treatments to limit the spread of medical misinformation. Based on the consensus input from the evaluators, we use well-tested machine learning systems to build models, which in turn review hundreds of thousands of hours of videos each day to find and limit the spread of borderline content. These models continue to improve in order to more effectively identify and reduce recommendations of borderline content.

Our efforts here have been publicly validated in several ways. Researchers in the United States and around the world who have studied YouTube have acknowledged that YouTube's recommendation algorithm actively discourages viewers from visiting radicalizing or extremist content. According to a 2020 study⁸ conducted by an Australian data scientist and a researcher at the University of California, Berkeley's School of Information: *"...data suggest that YouTube's recommendation algorithm actively discourages viewers from visiting radicalizing or extremist content. Our study thus suggests that YouTube's recommendation algorithm fails to promote inflammatory or radicalized content, as previously claimed by several outlets."* In 2018, a study⁹ from Pew Research also found that, on average, our recommendations point users to popular videos, which videos tend to be gaming, vloggers, and music rather than conspiracy theories or other types of misinformation—which again account only for a fraction of the content on YouTube.

REWARD TRUSTED CREATORS: In our mission to create and sustain an open, global platform, YouTube has also expanded economic opportunity for small businesses, artists, creators, journalists, rightsholders, and more. For many, sharing video content on YouTube is

⁷ <https://support.google.com/youtube/answer/9230586>

⁸ <https://journals.uic.edu/ojs/index.php/fm/article/view/10419/9404>

⁹

<https://www.pewresearch.org/internet/2018/11/07/many-turn-to-youtube-for-childrens-content-news-how-to-lessons/>

not just a hobby, but a business. Globally, the number of creators earning five figures annually increased more than 40% from December 2018 to December 2019.

Chris Bosio, owner of the Tampa barbershop Headlines, is one such creator¹⁰. When initial business was slow, Chris used his free time to teach the other barbers in his shop new techniques. Impressed by his down-to-earth, easy-to-understand lessons, his business partners convinced him to upload a tutorial to YouTube to bring attention to the shop. The video was a hit, so Chris kept creating videos, and before long he saw his subscribers turn into clients. Within a couple of months, Headlines went from five clients a week to 800, many of them mentioning that they had watched Chris's YouTube videos before coming into the shop.

As Chris's channel grew, his subscribers began to ask him to make his own shaving accessories—a revenue stream Chris hadn't considered yet. He learned how to make a shaving gel from YouTube tutorials and soon launched a line called Tomb45. YouTube is the main way Chris promotes the line, which creates a constant sales funnel for the company. With the support of his subscribers, Tomb45 sold over 70% of its inventory the first day. Today, Tomb45 sells 10,000 products a month in 15 countries. His YouTube business was a lifeline during the COVID-19 shutdown, when barbershops were forced to close. Without it, Chris isn't sure his business could have survived.

Today, millions of channels from over 90 different countries earn revenue from their videos by participating in our YouTube Partner Program (YPP). Through YPP, creators earn revenue generated from advertising that is shown to viewers before or during a video. This revenue from ads is shared between YouTube and the creator, with the creator receiving a majority share—thus empowering creators to directly profit from their work.

But generating revenue on YouTube is a privilege, reserved for creators who meet specific eligibility requirements. To be eligible for YPP, a creator must have more than 1,000 subscribers, 4,000 watch hours in the last 12 months, and a track record of adhering to our Community Guidelines. YPP creators must also adhere to Google Ads policies, and any videos that are monetized must also meet an even higher bar by adhering to YouTube's Advertising-Friendly Guidelines¹¹ for content. These guidelines outline what content may not be monetized and what content warrants limited monetization. If creator content violates any of our Community Guidelines, that content will be removed from YouTube. We also enforce our Advertising-Friendly Guidelines by limiting or blocking ads on videos in accordance with those

¹⁰ <https://blog.youtube/creator-and-artist-stories/chris-bossio-tomb45/>

¹¹ <https://support.google.com/youtube/answer/6162278>

guidelines. Creators who repeatedly violate any of our rules may be suspended from YPP for 90 days and need to apply again in order to rejoin YPP.

The 4Rs and Misinformation

We confront new challenges of balancing openness with responsibility every day. This is especially true when combating harmful misinformation on our platform. We invest a great deal of resources in research, policy development, technology, and experimentation to inform our approaches and improve our effectiveness in addressing misinformation on our platform. We continuously review our policies to evaluate whether the lines are in the right place; we prominently raise authoritative sources to the top of search results and make authoritative information readily available via a range of information panels; we regularly update our recommendation systems to hone our reduction of borderline content, including harmful misinformation; and we disincentivize creators seeking to profit off of misinformation by blocking ads on their content and suspending repeat offenders from our monetization program. We believe our 4R approach to responsibility provides a powerful and effective range of tools to combat harmful misinformation online, but we know there is more we can do. We commit to continuing to improve in our efforts to combat harmful misinformation on our platform.

Prioritizing transparency and accountability

At YouTube, we believe transparency is essential to earning and sustaining the trust of our users and our business partners. As a part of Google, we have led the way for the industry in terms of reporting on content removal at the request of governments and according to our own Community Guidelines, as well as information about government requests for information about users. We continue to expand our initiatives and the information we share, and we have rolled out three major resources over the last 12 months that underscore our commitment to transparency.

First, in May 2020, we collaborated with Google to launch the first Threat Analysis Group (TAG) Bulletin¹². The Bulletin—published on the TAG blog every quarter—discloses removal actions that Google and YouTube have taken to combat coordinated influence operations in a given quarter. Our hope is that this bulletin helps others who are also working to track these groups, including researchers working in this space, and that our information sharing can help confirm findings from security firms and other industry experts.

¹² <https://blog.google/threat-analysis-group>

Second, in June 2020, we launched a website called How YouTube Works¹³, which was designed to answer the questions we most often receive about our responsibility efforts and to explain our products and policies in detail. How YouTube Works addresses some of the important questions we face every day about our platform, and provides information on topics such as child safety, harmful content, misinformation, and copyright. The site also covers timely issues as they arise, like our COVID-19 response, and our work to support election integrity. Within the site, we explain how we apply our responsibility principles—which work alongside our commitment to users’ security—to manage challenging content and business issues.

Third, YouTube publishes quarterly data in our Community Guidelines enforcement report¹⁴. This report provides public data about the number of videos we remove from our platform for each of our policy verticals (except spam) as well as additional information about channel removals, views before removals, appeals and reinstatements, and human and machine flagging. And as noted above, just this month, we updated our report to include the Violative View Rate to reflect how effectively we identify and take action on violative content.

These are important steps but we know we are being called to do more so that we can be held accountable for the decisions we make—algorithmic or otherwise. We appreciate the feedback that we have received from Members of Congress on our efforts to date and look forward to continuing to examine additional steps that could be taken to build upon our transparency efforts. We will continue to expand the information we share through our transparency report, cross-industry initiatives, blog posts, public disclosures, and other mechanisms like tools for researchers. Our goal is to achieve transparency and accountability by providing meaningful information while protecting our platform.

Child safety and digital wellbeing

As discussed above, responsibility is our number one priority at YouTube—and nowhere is this more important than when it comes to protecting kids. We continue to make significant investments in the policies, products and practices to help us do this. From our earliest days, YouTube has required our users to be at least 13 years old, and we terminate accounts belonging to people under 13 when they are discovered. In 2015, we created YouTube Kids, an app just for kids that we created to provide a safe destination to explore their interests while providing parental controls. With availability in over 100 countries, now over 35 million viewers

¹³ <https://www.youtube.com/howyoutubeworks/>

¹⁴ <https://transparencyreport.google.com/youtube-policy/removals>

use YouTube Kids every week. We continue to expand product availability, add new features, and offer several parental tools.

We have also heard from parents and older children that tweens have different needs that were not being fully met by our products. As children grow up, they have insatiable curiosity and a need to gain independence and find new ways to learn, create, and belong. Over the last year, we have worked with parents and child development experts across the globe in areas related to child safety, child development, and digital literacy. This collaboration informed a recently announced new supervised experience for parents on our main YouTube platform with three content settings for parents to choose from. The YouTube supervised experience looks much like YouTube's flagship app and website, but with adjustments to the features children can use and ads protections. For example, comments and live chat are disabled, as well as the ability to upload content and make purchases. Additionally, automatic reminders will appear for breaks and bedtime, which they can adjust to reinforce healthy screen-time habits.

In addition to these specially designed products, our YouTube main app treats personal information from anyone watching children's content on the platform as coming from a child, regardless of the age of the user. This means that on videos made for kids—whether explicitly designated as such by the creator or identified as child-directed by our content classification systems—we limit data collection and use, and as a result, we restrict or disable some product features. For example, we do not serve personalized ads on this content on our main YouTube platform or support features such as comments, live chat, notification bell, stories, and save to playlist. To be clear, we have never allowed personalized advertising on YouTube Kids.

Conclusion

Technology that uses algorithms is critical to our day-to-day operations, both in terms of basic user-facing functionality as well as content management and moderation at scale. But so, too, is input from people, whether by evaluating context, providing feedback on the quality of videos, or controlling how and when they choose to use YouTube. Just as machine learning systems are constantly taking new inputs to hone their pattern detection and efficacy, we work continuously to address new threats and identify ways to improve our systems and our processes. Responsibility is and will continue to be our number one priority—our business depends on it.

We look forward to continuing to engage and discuss areas where we share priorities, and how we can join together to support research to identify novel approaches to problems that threaten

YouTube, LLC
901 Cherry Avenue
San Bruno, CA 94066
USA



Phone +1 650 253.0000
www.youtube.com

both our users and your constituents, as well as thinking about how media literacy efforts can help users develop skills to build resiliency against misinformation.

Thank you for the opportunity to discuss our work with you today.