

Responses to Questions for the Record
U.S. Senate Committee on the Judiciary
Subcommittee on Competition Policy, Antitrust, and Consumer Rights
“Algorithms and Amplification: How Social Media Platforms’ Design Choices Shape Our
Discourse and Our Minds”
April 27, 2021

Alexandra Veitch
Director Of Government Affairs And Public Policy For The Americas And Emerging
Markets, YouTube

Written Questions Submitted by Sen. Sasse to Alexandra Veitch:

Question 1. How does your firm define the ideal user experience? Does your answer differ between adult and children?

Because the answers to these questions are related, we have grouped together the responses to Question Numbers 1 and 4.

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. 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.

YouTube is committed to supporting users of all ages to develop healthy relationships with technology and use our services in ways that support their mental health and wellbeing. We work to understand both the opportunities and the challenges our users face, including through research and engagement with experts. And we build on this by developing features in our products and design policies that help our users — particularly children and families — form healthier digital habits.

YouTube has always had policies that govern how people may use the service, including restrictions on the types of content that they may post. These policies are designed and regularly updated to make YouTube a safer and more welcoming place for users, creators, and advertisers. They reflect years of experience and investment, and an ongoing conversation between YouTube and its users, creators, and advertisers.

For many years, we've worked with child development specialists to provide guidance to YouTube creators on how to create high-quality content for children. As a result of extensive consultations on children's media, digital learning, and the study of good citizenship, we established a set of quality principles (available at <https://support.google.com/youtube/answer/10774223>) to help guide our kids and family creator ecosystem. These principles include content that demonstrates or encourages respect and healthy habits, such as being a good friend, content that is thought-provoking or imaginative, such as arts and crafts activities, and content that celebrates and encourages diversity and inclusion. We also use these principles to determine which high-quality content we raise up in our recommendations on YouTube. This means that when you're watching content for kids and families on YouTube, we aim to recommend videos that are age-appropriate, educational, and inspire creativity and imagination. Additionally, these principles supplement our Community Guidelines (available at <https://support.google.com/youtube/answer/9288567>), which help provide a safe viewing experience for our users.

In 2015, we created YouTube Kids as a way for kids to more safely explore their interests and curiosity while providing parents more tools to control and customize the experience for their families. Videos on YouTube Kids include popular children's videos, as well as collections of channels selected by the YouTube Kids team and trusted kids content partners like Sesame Workshop, PBS Kids, and UNICEF. Our approach to determining eligibility in YouTube Kids is to first identify a very small subset of channels that have a low likelihood of uploading inappropriate content from the full YouTube corpus of content that is available.

We also provide parents the ability to fully customize what their kids watch. For example, parents can block a video or channel, handpick which videos to make available for their kids, and select an option to only allow content selected by trusted partners. We are constantly working to improve our safeguards and offer more features to help parents create the right experience for their families. For more information on the content policies for YouTube Kids, please see <https://support.google.com/youtube/answer/10938174?hl=en>.

After we launched YouTube Kids, we heard from parents of older children that tweens and teens have different needs which were not fully met by our products. That is why we worked with parents and experts across the globe in areas related to child safety, child development, and digital literacy to develop a solution for parents of tweens — which we call Supervised Experience — that launched in March 2021 on the main YouTube platform. Supervised Experience allows parents to choose between three different content choices: content generally suitable for viewers aged 9+; content generally suitable for viewers 13+; and the 'Most of YouTube' option, which excludes all age-restricted content (18+).

To be clear, kids under 13 who are not in a Supervised Experience are not allowed on YouTube. When a parent opens a supervised Google account for their child, the experience feels much

like regular YouTube but certain features are disabled for younger audiences. Underlying the ads and creator policies applicable to both YouTube Kids and Supervised Experiences are core principles aimed at maintaining an appropriate viewing environment for children and families and protecting young users' privacy. For example, on both YouTube Kids and Supervised Experiences we prohibit personalized advertising and age-sensitive ad categories (such as those available at <https://support.google.com/youtube/answer/9713557?hl=en>). We do not allow ads that contain adult, dangerous, sexualized or violent content; ads that mislead or make deceptive claims; and sensitive product categories such as ads for pharma/healthcare, dating, politics, weight loss, beauty and fitness.

To help parents understand more about the YouTube Supervised Experience, we have developed guides (such as <https://www.youtube.com/myfamily/>), videos (such as <https://www.youtube.com/watch?v=oVOa6nDU7HQ>), and support pages (available at <https://support.google.com/youtube/answer/10314940?hl=en>), building on Google's successful Be Internet Awesome digital literacy resources, and in partnership with the National PTA, Parent Zone UK and other leading experts. We will continue to partner with these and other groups to provide easy to use resources specifically for parents to help them keep their kids safe online.

Question 2. From your firm's perspective, what is the ideal amount of time for a user to spend on the platform over a specified period of time? How much is too much?

Our business relies on the trust of our users. We want to build these relationships for the long term. That is why we bake user choice and user control right into the product with things like timers and the opportunity to turn auto-play off.

We also employ systems that enable our users to make decisions on screen time that are best for them and their families. For example, Family Link is a downloadable app (available at <https://families.google.com/familylink/>) now available by default in the latest Android operating system and Chromebooks that helps parents guide their child's experience as they explore online. The app lets parents set digital ground rules for their family, such as managing the apps their child can use, keeping an eye on screen time, or setting a bedtime and daily limits for their child's device. SafeSearch, a default Family Link supervision feature, helps filter out explicit search results on Google Search. And, as discussed in the response above, YouTube Kids provides a separate YouTube experience designed especially for children that parents can customize. Parents can decide what content to make available for their child to watch, set a timer to control screen time, block videos or channels, and more. For more information on the tools we provide on YouTube Kids to help parents manage screen time, please see <https://support.google.com/youtubekids/answer/6130558>.

Question 3. If an individual pulled up her account on your platforms, could somebody at your company walk her through the exact calculation behind why the first post ended up where it is? And if not, what are the limitations, and how close could you get to a full explanation?

Recommendations on YouTube help users discover videos they may enjoy and creators reach new viewers to grow their audience across the platform. Recommendations occur in two main places: a user's homepage and the "Up Next" panel. Our recommendation system takes into account many signals, including: a user's viewing habits, such as clicks, watchtime, survey responses, sharing, likes, and dislikes; and 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.

We also empower our users by giving them control over personalized recommendations. In June 2019, we launched new features that give users more control over what recommendations appear on the homepage and in their "Up Next" suggestions. Three specific changes were made in these new features to make it easier for users to block channels from recommendations, give users the option to filter recommendations on Home and on "Up Next," and give users more information about why we are suggesting a video:

- **Explore topics and related videos on your Homepage and in Up Next videos:** Users can now more easily explore topics and related videos on their Homepage and in Up Next videos to see more of what they love on YouTube. Users now see options based on existing personalized suggestions that are meant to help users find what they are looking for faster.
- **Remove suggestions from channels you don't want to watch:** Users can easily let us know to stop suggesting videos from a particular channel by tapping the three-dot menu next to a video on the homepage or Up Next, then "Don't recommend channel." After that, users should no longer see videos from that channel suggested to them on YouTube.
- **Learn more about why a video may be suggested to you:** Sometimes we recommend videos from channels users have not seen before based on what other viewers with similar interests have liked and watched in the past. When we're suggesting videos based on this, users now see more information underneath the video in a small box. Our goal is to explain why these videos surface on your homepage in order to help you find videos from new channels you might like.

Question 4. If a family member asked you for advice on whether and under what conditions the family member should allow her child to use your firm's social media platform, what advice would you offer? How would it change based on age?

Please refer to the answer to Question No. 1, above.

Question 5. Do you think that there is anything fundamentally wrong with the business model of your firm along the lines of the criticisms outlined by Dr. Donovan and Mr. Harris?

At YouTube, we take a comprehensive approach to combating online extremism, abuse and other harmful content. Over the last few years, we have significantly increased our investments in the systems and processes that allow us to remove violative content, raise up authoritative content, reward trusted creators and artists, and reduce the spread of borderline content. Our teams work hard every day to ensure that we are allowing for a broad range of political speech while adhering to our responsibility of making sure that our platform is not abused to incite real-world harm or spread harmful misinformation.

Managing misinformation and harmful conspiracy theories is challenging because the content is always evolving, but we take this issue very seriously. Due to the shifting tactics of groups promoting conspiracy theories, we continuously invest in the policies, resources, and products needed to protect our users from harmful content. We have clear and public policies that we apply consistently. Across all of Google platforms, we have dedicated significant resources to content moderation efforts: we spent more than one billion dollars on content moderation efforts in 2020 and have nearly twenty thousand employees dedicated to ensuring the protection of our platforms.

In addition to removing content that violates our policies, we also work to ensure that we do not proactively expose users to content that is potentially harmful. We reduce recommendations of the “borderline content”, or content which comes close to but does not violate our policies, that represents less than 1% of the content watched on YouTube. In January 2019, we announced that we would begin reducing recommendations of borderline content or videos that could misinform users in harmful ways, and we continue to extend these efforts to more countries outside of the United States and into non-English-language markets. We have launched more than 30 different changes to our recommendations systems on YouTube in order to reduce recommendations of borderline content and harmful misinformation. In 2019, we saw a more than 70% average drop in “watch time” of this content coming from non-subscribed recommendations in the United States.

Determining what is harmful misinformation or borderline content is challenging, especially given the wide variety of videos uploaded to YouTube. To do it, we rely on external evaluators from around the world to provide input on the quality of a set of videos. These evaluators use the same rater guidelines as Google Search to guide their work (these guidelines are available at <https://support.google.com/youtube/answer/9230586>). Each evaluated video receives up to nine different ratings, with some content requiring ratings from certified experts in the field. For example, medical doctors provide guidance on the validity of videos about specific medical treatments to limit the spread of medical misinformation. Based on consensus input from these raters, we use well-tested machine learning systems to build models that help

review hundreds of thousands of hours of videos every day to identify and limit the spread of borderline content. The accuracy of these systems continues to improve over time.

We also raise up authoritative voices by providing users with more information about the content they are seeing to allow them to make educated choices. For more information on information panels, please see <https://support.google.com/youtube/answer/9004474>, and on fact-checks, please see <https://support.google.com/youtube/answer/9229632>.

The openness of our platforms has helped creativity and access to information thrive. It is our responsibility to protect that, and to prevent our platforms from being used to incite hatred, harassment, discrimination, and violence. We are committed to taking the steps needed to live up to this responsibility today, tomorrow, and in the years to come.

Question 6. If somebody like Dr. Donovan who spends her time researching online extremists has her feed filled with extremist content, should we consider that to be a job well done in optimizing her user experience or something that should worry us?

The rise of misinformation in recent years led us to further expand the ways we use our recommendation system to address problematic misinformation and borderline content. This includes conspiracy theory videos (“the moon landing was faked”) or other content that spreads misinformation (“orange juice can cure cancer”).

We are able to do this by using classifiers to identify whether a video is “authoritative” or “borderline”. To determine authoritativeness, we look at a few key questions: Does the content deliver on its promise or achieve its goal? What kind of expertise is needed to achieve the video goal? What is the reputation of the speaker in the video and the channel it is on? What is the main topic of the video (e.g., News, Sports, History, Science, etc)? Is the content primarily meant to be satire? These answers and more determine how authoritative a video is. The higher the score, the more the video is promoted when it comes to news and information content.

To determine borderline content, we assess factors that include — but aren’t limited to — whether the content is: inaccurate, misleading or deceptive; insensitive or intolerant; or harmful or with the potential to cause harm. The results are combined to give a score for how likely the video contains harmful misinformation or is borderline. Any video classified borderline is demoted in recommendations. These human evaluations then train our system to model their decisions, and we now scale their assessments to all videos across YouTube.

To create a safe environment for our users at YouTube, we approach our content recommendations with our four “Rs” of responsibility: we remove content that violates our policies as quickly as possible, we raise up authoritative voices, we reduce the spread of content that brushes up against our policy (i.e., borderline content), and we reward trusted creators.

- **We remove content that violates our policies.** In Q2 2022 alone, we removed more than 4.4 million videos for violating our Community Guidelines, more than 93% of which were first flagged by our automated systems. More than 72% of those first flagged by our systems received 10 or fewer views. Most recently, we began disclosing our violative view rate (VVR), or the number of views on content that violates our Community Guidelines as a portion of total viewership. More information on our VVR is available at <https://transparencyreport.google.com/youtube-policy/views>. In Q2 2022 our VVR was 0.09-0.11%, meaning that out of every 10,000 views on YouTube, 9-11 came from violative content.
- **We raise up authoritative voices when people are looking for news and information.** For example, we raise content from authoritative channels for newsworthy events in YouTube search results and “Up Next” recommendations panels. In addition to current events, authoritativeness is also important when it comes to topics prone to misinformation, such as vaccines. In these cases, we aim to surface videos from experts, like public health institutions, in search results. We also provide informational panels to provide contextual information on specific content in both the search results page and the video watch page. In October 2020, we updated YouTube’s hate and harassment policy to prohibit content targeting individuals or protected groups based on conspiracy theories used to justify real world violence (e.g., QAnon and Pizzagate).
- **We reduce recommendations of borderline content.** As referenced above, in January 2019, we announced improvements to our recommendation system to greatly reduce recommendations of borderline content and content that could misinform users in harmful ways. Globally, borderline content that comes from our recommendations represents less than 1% of the content watched on YouTube.
- **We reward trusted creators by allowing those channels to make money on our site.** We have strict policies on the kind of videos we allow ads to appear on that are strictly enforced, and creators must meet specific eligibility requirements to join our monetization program.

We are proud of these efforts, and believe they work to help curb the spread of extremist content. A growing number of independent researchers have been looking into how tech platforms impact the consumption of borderline content, and, while ongoing study continues, certain recently published papers conclude YouTube recommendations do not steer viewers towards extreme content. Instead, consumption of news and political content on YouTube more generally reflects personal preferences that can be seen across their online habits. To encourage further research, we recently launched the YouTube Research Program (more information is available at <https://research.youtube/>), which gives academic researchers access to our raw data in a user privacy sensitive way, allowing them to derive metrics or develop insights of their own based on that data.

Question 7. If a user’s behavior indicates that he or she consistently chooses to focus on content that makes them angry or anxious or depressed, what is the responsibility of the platform in terms of curating the content that it places in front of that user?

Please refer to the answer to Question No. 8, below.

Question 8. Do you measure whether content that stimulates a reaction of anger elicits more engagement from users? If so, what are the results, and how are those results used to shape the platform?

Because the answers to these questions are related, we have grouped together our response to Question Numbers 7 and 8.

Across our different platforms, we strive to connect people with “high-quality information” — the most useful, trustworthy, and helpful content — at the moment a person needs it. At the same time, we work to prevent user and societal harm and limit the reach of “low-quality information” — content that strays furthest from those qualities. We build products that help users access high-quality information that supports their mental health and wellbeing.

At YouTube, we take the health and well-being of all our creators and viewers seriously. Awareness and understanding of mental health is important and we support creators sharing their stories, such as posting content discussing their experiences with depression, self-harm, or other mental health issues. We also have a clear set of publicly available Community Guidelines (available at <https://support.google.com/youtube/answer/9288567>) that outline content that is not allowed on the platform.

Per our suicide and self-injury policy, we prohibit content that promotes self-harm or is intended to shock or disgust viewers. We remove content promoting or glorifying suicide, content providing instructions on how to self-harm or die by suicide and content containing graphic images of self-harm posted to shock or disgust viewers. The Stanford Internet Observatory highlighted how we handle searches related to self-harm and suicidal ideation, writing that they were “impressed that YouTube’s Community Guidelines on suicide and self-injury provide resources, including hotlines and websites, for those having thoughts of suicide or self-harm, for 27 countries,” and that creator platforms like YouTube “offer meaningful support in their policies both for people who are recovering from mental health issues and those who may be considering self-harm.” The report detailing these findings is available at <https://cyber.fsi.stanford.edu/io/self-harm-policies-report>.

YouTube uses a combination of machine learning and human review to enforce our policies and we regularly report on the content removed for violating our policies in our quarterly Community Guidelines Enforcement Report (available at <https://transparencyreport.google.com/youtube-policy/removals>). In Q2 2022, we removed over 900,000 videos for violating our violent or graphic content policy and over 1.3 million videos for violating our child safety policy, the majority of which were removed before they

had received 10 views. Most recently, we began disclosing our violative view rate (VVR), or the number of views on content that violates our Community Guidelines as a portion of total viewership. More information on our VVR is available at <https://transparencyreport.google.com/youtube-policy/views>. In Q2 2022, our VVR was 0.09-0.11%, meaning that out of every 10,000 views on YouTube, 9-11 came from violative content.

Question 9. In the hearing, I noted that political debate about regulating social media platforms often fails to respect the distinctions between the First Amendment limitations of government actors and the greater freedom enjoyed by private actors. While your firm is not a government entity, we should not overlook the history of cooperation between social media platforms and federal, state, and local governments in some content decisions. Many of these efforts, such as limiting access to terrorist propaganda, have long earned bipartisan support. Nevertheless, learning more about the role of government actors in your content decisions would be a significant and helpful step in informing the ongoing debate. Please give an overview of your partnerships with federal, state, and local governments in making all aspects of content decisions (e.g., developing standards, removing content that violates these standards, deprioritizing or demonetizing content, flagging content as potentially inaccurate or misleading or providing links to authoritative resources, promoting certain kinds of content).

Please refer to the answer to Question Number 10, below.

Question 10. At the federal level, to what extent is this liaising with the federal government done by political appointees instead of career officials? Does contact occur at the agency level or the White House level? Please characterize the nature of this contact.

Because the answers to these questions are related, we have grouped together our responses to Question Numbers 9, 10, and 12.

At YouTube, we develop and maintain Community Guidelines that outline what type of content is not allowed on our platform. These policies (available at <https://support.google.com/youtube/answer/9288567>) act as our ‘rules of the road’ and we aim to make them clear and easily accessible to all users and content creators. Our Community Guidelines also define what content and behaviors are not allowed, and the process by which a piece of content, or its creator, may be removed from service. These Community Guidelines are one of the ways we balance allowing diverse voices on the platform while limiting harmful content. We enforce our content policies at scale and take tens of millions of actions every day against content that does not abide by the policies for one or more of our products.

Courts of law and government agencies around the world regularly request that we remove content from YouTube. Some requests may allege infringement of intellectual property rights, while others claim violation of local laws prohibiting certain types of user generated content,

such as defamation or Neo-Nazi content in certain members states of the European Union. We review these requests closely to determine if content should be removed or restricted because it violates a local law or is contrary to our terms of service and content policies. In order for us to evaluate a request from a government entity, it must be in writing and provide a clear explanation of how the content is violative of law or policies.

Legal standards vary by country, and content that violates a specific law in one country may be legal in others. Typically, we remove or restrict access to the content only in the country where it is deemed to be illegal. But when content is found to violate our Community Guidelines we remove it globally.

It is important to note that YouTube does not always remove content in response to a government request. Some requests may not be specific enough for us to know what the government wanted us to remove (for example, no URL is listed in the request) or lack sufficient explanation of why the government believes content to violate a particular law. For defamation claims, we generally rely on courts to decide if a statement is defamatory according to local law, and do not rely on informal letters from government agencies.

With regard to government requests for user information, a variety of laws allow government agencies around the world to request the disclosure of user information for civil, administrative, criminal, and national security purposes. Each request is carefully reviewed to make sure it satisfies applicable laws. For more information, see our policies for how Google handles government requests for user information (available at <https://policies.google.com/terms/information-requests>).

We also publicly share information on government requests to remove content to shed light on the scale and scope of government requests for content removals. Data on government removal requests can be located on the Google Transparency Report (available at <https://transparencyreport.google.com/government-removals/overview?hl=en>).

There are also a number of ways we proactively work with law enforcement agencies to assess threats and to counter attempts to deceive, harm, or take advantage of people using our platforms. Our interactions with law enforcement include the following:

- First, we maintain regular communication channels with law enforcement, as communication with law enforcement, industry partners, and the federal government is a key component of our efforts to keep users safe. We rely on information from industry peers and law enforcement to ensure the integrity of our platform and act swiftly in response to crises or when we detect abuse that may threaten public safety or the integrity of democratic processes.
- Second, we have dedicated teams of analysts and security experts who work around the clock to identify and investigate possible instances of coordinated influence

operations on our platform. Google's Threat Analysis Group works to identify bad actors and terminate their channels and accounts, works with other technology companies to share intelligence and best practices, and shares threat information with law enforcement. For more information regarding our Threat Analysis Group, please see <https://blog.google/threat-analysis-group/>.

- Third, we have processes in place to proactively refer to law enforcement imminent threats and certain illegal activity occurring on our platform. These imminent threats are raised to us in a variety of ways. In addition to our ongoing threat analysis just noted, we also have processes in place for our various internal product teams, as well as external sources, to escalate potential threats and criminal activity if they see it on our platforms. Google has a team - the CyberCrime Investigation Group (CCIG) - in place to assess escalated threats, and that team refers matters to law enforcement whenever appropriate.

Additionally, we developed the YouTube Trusted Flagger program to help provide more sophisticated reporting processes for government agencies and non-governmental organizations (NGOs) that are particularly effective at notifying YouTube of content that violates our Community Guidelines. The program provides these partners with dedicated reporting processes and a channel for ongoing discussion and feedback about YouTube's approach to various content areas. The program is part of a network of around 300 government partners and NGOs that bring valuable expertise to our enforcement systems. Participants in the Trusted Flagger program receive training in enforcing YouTube's Community Guidelines, and because their flags have a higher action rate than the average user, we prioritize them for review. Once flagged, our YouTube content reviewers make the call on what content is or isn't removed by applying our YouTube policies. We apply our policies transparently and consistently, without taking into account either the political leanings of the author, or their place in society.

Question 11. How often do you receive communications from legislators urging more aggressive content moderation?

Our mission at YouTube is to "give everyone a voice and show them the world." As such, our search and recommendation systems are not designed to filter or demote videos or channels based on specific political perspectives. We take extraordinary lengths to build our products and enforce our policies in a consistent manner. Our business model is dependent on being a useful and trustworthy source of information for everyone, so we have a natural, long-term business incentive to prevent anyone from interfering with the integrity of our products.

When developing and updating our policies, we solicit perspectives from a range of voices, including creators, subject-area experts, free speech proponents, and policy organizations from across the political spectrum. Once a policy has been developed, we invest significant time making sure new policies are consistently enforced by our global team of reviewers who

follow objective guidelines Our work is also guided by our four pillars of responsibility: removing violative content; raising up authoritative content; reducing the spread of borderline content; and rewarding trusted creators — what we refer to as the 4Rs of responsibility. YouTube’s Community Guidelines outline what content is permissible and what content is violative on our platform. We enforce these Community Guidelines using a combination of human reviewers and machine learning, and apply them to everyone equally — regardless of the subject or the creator’s background, political viewpoint, position, or affiliation.

We are committed to ensuring that our participation and associations with politics, trade associations, or third-party organizations is open, transparent and clear to our users, shareholders, and the public. Information on our public policy engagement, political contributions, lobbying efforts, and memberships — including our biannual list of politically engaged trade associations, independent third-party organizations, and other tax-exempt groups that receive the most substantial contributions from our U.S. Government Affairs and Public Policy Team — is published in our US Public Policy Transparency Reports (available at <https://www.google.com/publicpolicy/transparency/>).

Question 12. Has your firm ever been the subject of threats of negative consequences by federal officials for refusing to take certain content moderation steps?

Please refer to the answer to Question Number 10, above.

Question 13. Why is YouTube’s autoplay feature enabled by default?

The Autoplay feature on YouTube makes it easier to decide what to watch next. When Autoplay is on, another related video will automatically play after a video ends. Users can have different Autoplay settings for different devices. For example, users can set Autoplay to be "On" in the YouTube app on their mobile device, but set it to "Off" when watching YouTube on their computer. If users are connected to a mobile network, Autoplay will stop automatically if they have been inactive for 30 minutes. If users are on Wi-Fi, Autoplay will stop automatically after four hours. For more information on YouTube Autoplay, please see <https://support.google.com/youtube/answer/6327615>

For teens and for those in our YouTube Supervised Experience and in YouTube Kids, autoplay is now off by default. We have also launched additional parental controls in the YouTube Kids app, including the ability for a parent to choose a “locked” default autoplay setting. Our “Take a Break” reminders and bedtime settings are also on by default in these experiences. In addition, we have worked with creators to develop a series of PSAs to help children reflect on the time they spend online, and build empathy for others.

Question 14. What percentage of your users subscribe to YouTube Premium?

As of September 2021, we have 50 million Music and Premium subscribers, including trial members. It is an honor to build a membership that allows people to more deeply immerse

themselves in music, learning, fashion, gaming, and more, all the while supporting the creators and artists that make it possible. Music and Premium subscriptions are key pillars of YouTube's monetization, enabling unique content and communities to flourish (more information on monetization is available at

<https://blog.youtube/news-and-events/10-ways-monetize-youtube/>).

Question 15. Aside from disabling ads and the other advertised features of the service, does subscribing to YouTube Premium otherwise alter how content is presented to a user?

We offer YouTube Premium to amplify the user experience on YouTube. With YouTube Premium, subscribers can watch millions of videos without interruptions by ads before and during a video, including video overlay ads. Subscribers also do not see third-party banner ads and search ads, but may still see branding or promotions embedded in the content by the creator and promotional links, shelves and features in and around the content that are added or enabled by the creator. These links, shelves, and features could be for their website, merchandise, a membership to their channel, event tickets, or other related destinations they are promoting.

Written Questions Submitted by Sen. Hirono to Alexandra Veitch:

Question 1. Last fall, Jack Dorsey testified before the full Senate Judiciary Committee that “people should have choices about the key algorithms that affect their experience online,” including the ability to use algorithms created by third parties to rank and filter the content they see.

- a. Do you agree that “people should have choices about the key algorithms that affect their experience online?”**
- b. Does YouTube commit to allowing people to choose the algorithm that recommends videos on the platform, including a choice of algorithms developed by third parties?**

We give our users a number of choices that help form their experience on our platforms. First, recommendations on YouTube help users discover videos they may enjoy, and they help creators reach new viewers and grow their audience across the platform. Recommendations work in two main places: a user’s homepage and the “Up Next” panel. And users have control over their data that is stored and how it is saved and used on our platform. For example, users can access and delete their search and watch history, and any videos removed from their history will no longer influence their recommendations, and queries they have removed will not appear in their search bar as suggestions (more information on user choice concerning their data in YouTube is available at

<https://www.youtube.com/howyoutubeworks/user-settings/privacy/#your-data-in-youtube>).

Our recommendation system constantly evolves, learning every day from over 80 billion pieces of information we call signals. A number of signals build on each other to help inform our dynamic system about what a particular user finds satisfying: clicks, watchtime, survey responses, sharing, likes, and dislikes (more information about these signals is available at <https://blog.youtube/inside-youtube/on-youtubes-recommendation-system/>).

- **Clicks:** Clicking on a video provides a strong indication a user will also find it satisfying. After all, a user would not click on content they do not want to watch.
- **Watchtime:** Watchtime — which videos a user watched and for how long — provides personalized signals to our system about what users most likely want to watch.
- **Survey Responses:** To ensure viewers are satisfied with the content they watch, we measure what we call “valued watchtime” — the time spent watching a video that viewers consider valuable. We measure valued watchtime through user surveys that ask viewers to rate the video watched from one to five stars, giving us a metric to determine how satisfying viewers found the content. If a video is rated one to two stars, we ask the viewer why they gave such a low rating. Similarly, if the viewer gives the video four to five stars, we ask why — was it inspirational or meaningful? Only videos that are rated highly with four or five stars are counted as valued watchtime.

- **Sharing, Likes, Dislikes:** On average, people are more likely to be satisfied by videos that they share or like. Our system uses this information to try to predict the likelihood that viewers will share or like further videos.

In June 2019, we launched new features that give users more control over what recommendations appear on the homepage and in their “Up Next” suggestions (more information is available at <https://blog.youtube/news-and-events/giving-you-more-control-over-homepage/>). Three specific changes were made:

- **Explore topics and related videos on your Homepage and in Up Next videos:** Users can now more easily explore topics and related videos on their Homepage and in Up Next videos to see more of what they love on YouTube. Users now see options based on existing personalized suggestions that are meant to help users find what they are looking for faster.
- **Remove suggestions from channels you don’t want to watch:** Users can easily let us know to stop suggesting videos from a particular channel by tapping the three-dot menu next to a video on the homepage or Up Next, then “Don’t recommend channel.” After that, users should no longer see videos from that channel suggested to them on YouTube.
- **Learn more about why a video may be suggested to you:** Sometimes we recommend videos from channels users have not seen before based on what other viewers with similar interests have liked and watched in the past. When we’re suggesting videos based on this, users now see more information underneath the video in a small box. Our goal is to explain why these videos surface on your homepage in order to help you find videos from new channels you might like.

These features make it easier for users to block channels from recommendations, give users the option to filter recommendations on Home and on “Up Next,” and give users more information about why we are suggesting a video.

In addition to providing choice, we believe that transparency around our processes is important. As an indication of this commitment, we recently launched the YouTube Research Program. The program gives academic researchers access to our raw data in a user privacy sensitive way, allowing them to derive metrics or develop insights of their own based on that data. This increased access includes expanded data quotas, ability to derive analytics, and in-house YouTube support. As we expand this access to our API, we’ll work closely with researchers to gain feedback and improve collaborations. Finally, we provide users with the ability to influence the ads they see online. As explained in the Google Privacy Policy, if you use a Google Account, Google may use information you provide to show you relevant ads. Users can always learn more information concerning personalized advertising and controlling the ads they see on Google’s support pages (see

<https://support.google.com/google-ads/answer/2549116> and <https://support.google.com/accounts/answer/2662856>). In addition, users can turn off ad personalization in their Google Account settings, available at <https://adssettings.google.com/>.

Question 2. Former-WDBJ reporter Alison Parker was murdered live on-the-air in 2015. Not long after, videos of her murder were uploaded to YouTube. Despite the videos violating YouTube’s Community Guidelines regarding “violent or dangerous content” and efforts by Alison’s father Andy Parker, the HONR Network, and others (including my office) to identify these videos to YouTube, many copies of the video remain easily accessible on the platform.

- a. What specific steps has YouTube taken to ensure that videos of Ms. Parker’s murder are either removed from YouTube or not uploaded in the first place?**
- b. Please explain YouTube’s process for proactively identifying content that may violate YouTube’s Community Guidelines, including any automated processes employed by YouTube such as hashing or machine learning.**
- c. Once content has been flagged as possibly violating YouTube’s Community Guidelines—whether through YouTube’s internal processes or through the efforts of users—please explain the process for determining whether the content should be removed or have other action taken.**
- d. When YouTube becomes aware of its platform being used to spread violent and graphic footage of a particular event, does YouTube take additional steps beyond its normal content moderation practices to remove such footage? If so, please describe those steps in detail.**

Violent footage filmed by perpetrators such as Alison Parker's murderer are not allowed on YouTube, and we are committed to removing any content that violates our policies. Our policies specifically prohibit videos that aim to shock with violence, or accuse victims of public violent events of being part of a hoax. For certain egregious content, like videos containing footage of Alison Parker’s murder, we use digital fingerprinting to detect and remove it from YouTube. We have automatically detected and removed over 400 reuploaded versions of the murderer’s video before they were ever flagged by a user.

YouTube maintains detailed policies against violent or dangerous content

At YouTube, we develop and maintain Community Guidelines that outline what type of content is not allowed on our platform (available at <https://support.google.com/youtube/answer/9288567?hl=en>). These policies act as our ‘rules of the road’ and we aim to make them clear and easily accessible to all users and content creators. Our Community Guidelines define what content and behaviors are not allowed, and the process by which a piece of content, or its creator, may be removed from service. These

Community Guidelines are one of the ways we balance allowing diverse voices on the platform while limiting harmful content.

Our Community Guidelines are designed to protect the YouTube Community and include a number of policies that protect users from violent or dangerous content. These include policies that address violent extremism, violent or graphic content, harmful or dangerous content, election misinformation and hate speech. They also include policies that prohibit content that praises, incites, or promotes violence, or that which is misleading and deceptive with serious risk of egregious harm (including content that interferes with democratic processes). For example:

- **Harmful or Dangerous Content:** YouTube prohibits content that encourages dangerous or illegal activities that risk serious physical harm or death. This policy applies to various types of conduct, including content that provides instructions to build a bomb meant to injure or kill others, and content that promotes or glorifies violent tragedies, such as school shootings.
- **Firearms:** YouTube prohibits content intended to sell firearms, instruct viewers on how to make firearms, ammunition, and certain accessories, or instruct viewers on how to install those accessories. We also do not allow live streams that show someone holding, handling, or transporting a firearm.
- **Harassment & Cyberbullying:** YouTube prohibits content that threatens individuals or targets an individual with prolonged or malicious insults based on intrinsic attributes. This policy includes, for example, targeting an individual and making claims that they are involved in human trafficking in the context of a harmful conspiracy theory where the conspiracy is linked to direct threats or violent acts.
- **Hate Speech:** YouTube removes content promoting violence or hatred against individuals or groups based on age, caste, disability, ethnicity, gender identity and expression, nationality, race, immigration status, religion, gender, sexual orientation, veteran status, or victims of a major violent event and their kin. This policy also prohibits videos alleging that a group is superior in order to justify discrimination, segregation or exclusion based on qualities like age, gender, race, caste, religion, sexual orientation or veteran status. This would include, for example, videos that promote or glorify Nazi ideology, which is inherently discriminatory. We also prohibit content denying that well-documented violent events, like the Holocaust or the shooting at Sandy Hook Elementary, took place.
- **Violent Criminal Organizations:** YouTube prohibits content intended to praise, promote, or aid violent criminal organizations is not allowed on YouTube. An example of content that violates this policy would be videos or comments directing users to sites hosting manifestos from the perpetrators of violent incidents, such as the Christchurch attacker. This policy also applies to channels and content produced by violent criminal or terrorist organizations, including that:

- praising or memorializing prominent terrorist or criminal figures to encourage others to carry out acts of violence;
 - praising or justifying violent acts carried out by violent criminal or terrorist organizations;
 - aimed at recruiting new members to violent criminal or terrorist organizations designated by the U.S. government;
 - depicting hostages or posted with the intent to solicit, threaten, or intimidate on behalf of a violent criminal or terrorist organization; and
 - depicting the insignia, logos, or symbols of violent criminal or terrorist organizations to praise or promote them.
- **Violent or Graphic Content:** YouTube prohibits violent or gory content intended to shock or disgust viewers, or content encouraging others to commit violent acts. This includes, but is not limited to, content that incites others to commit violent acts against individuals or a defined group of people; fights involving minors; and content where animals are encouraged or coerced to fight by humans. It also applies to footage, audio, or imagery involving road accidents, natural disasters, war aftermath, terrorist attack aftermath, street fights, physical attacks, sexual assaults, immolation, torture, corpses, protests or riots, robberies, medical procedures, or other such scenarios with the intent to shock or disgust viewers. Further, footage or imagery showing bodily fluids, such as blood or vomit, with the intent to shock or disgust viewers is prohibited by this policy.

Content that violates our policies against violent extremism, detailed above, includes material produced by government-listed foreign terrorist organizations. We do not permit these terrorist organizations to use YouTube for any purpose, including recruitment. YouTube also strictly prohibits content that promotes terrorism, such as content that glorifies terrorist acts or incites violence. Content produced by violent extremist groups that are not government-listed foreign terrorist organizations is subject to our policies, including those described above, regardless of the uploader or the group to which they belong.

Finally, we regularly review and update our policies to address new and emerging threats. In 2017, we introduced a tougher stance toward videos with supremacist content, including limiting recommendations and features like comments and the ability to share the video. In June 2019, we clarified our hate speech policy to specifically prohibit videos alleging that a group is superior to justify discrimination, segregation or exclusion based on qualities like age, gender, race, caste, religion, sexual orientation or veteran status. In October 2020, we expanded our hate and harassment policies to prohibit content that targets an individual or group with conspiracy theories that have been used to justify real-world violence, such as QAnon.

Due to the evolving nature and shifting tactics of groups promoting these conspiracy theories, we will continue to adapt our policies to stay current and remain committed to taking the steps needed to live up to this responsibility.

YouTube has robust methods to enforce its policies regarding violent or dangerous content

We rely on a combination of people and technology to enforce our policies. YouTube has invested significantly in our automated detection systems and our engineering teams continue to update and improve them month by month. In 2017, we significantly expanded our use of machine learning technology to help detect potentially violative content and send it for human review at scale. Machine learning is well-suited to detect patterns and find content similar (but not identical) to other content we've already removed, even before it's ever viewed. Automated flagging also allows us to identify and act more quickly and accurately to remove violative content, lessening both the burden on human reviewers and the time it takes to remove such content. For example, in Q2 2022, more than 93% of videos that we removed as violative were first flagged through our automated systems.

Our machine learning systems help our human review teams remove content at scale, with the speed and volume that could not be achieved with people alone. Recently, we removed nearly 73,000 videos or promotion of violence and violent extremism between April and June 2022. We also use hashes (or "digital fingerprints") to automatically identify copies of known violative content before they are ever made available for viewing. These systems automatically remove content only where there is high confidence of a policy violation and the other content is flagged for human review. YouTube has invested significantly in these human reviewers, who remain essential to both removing content and training machine learning systems to help make contextualized decisions on content.

Additionally, we release a quarterly Community Guidelines Enforcement Report (available at <https://transparencyreport.google.com/youtube-policy/removals>) that shows the progress we are making in removing violative content from our platform. In Q2 2022 alone, we removed more than 4.4 million videos for violating our Community Guidelines. More than 72% of those first flagged by our systems received 10 or fewer views and more than 93% were detected by our automated flagging system.

Written Questions Submitted by Sen. Coons to Alexandra Veitch:

Question 1. I appreciate your testimony that YouTube is focused on long-term usage of its products rather than prioritizing short-term engagement. Some have raised concern, however, about whether the goals and metrics that social media platforms use to measure performance and pay for product teams may create structural incentives within the company potentially at odds with this long-term interest.

- a. Please enumerate the metrics that YouTube uses to set goals for its product teams—particularly the team responsible for recommending videos—for purposes of performance and pay.**
- b. Please explain why YouTube chose those metrics.**
- c. Does YouTube take any integrity-related metrics into account for its performance and pay assessments for any of its teams, including those focused on trust, safety, and integrity?**

Please refer to the answer to Question Number 2, below.

Question 2. Some have raised concern about whether social media companies are internally organized to ensure that trust and safety teams have adequate voice, staffing, and resources in product development.

- a. What teams are responsible for ensuring that the recommendation algorithm does not promote violating content? Please elaborate on how this process is structured and coordinated internally.**
- b. To the extent there are tensions between topline goal metrics and integrity metrics in assessing business units or product features, how do you resolve them, and who makes those decisions?**

Because the answers to these questions are related, we have grouped together our responses to Question Numbers 1 and 2.

Our first priority is to protect our users. Content moderation at Google and YouTube is primarily managed by Trust and Safety teams across the company. 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 also work with our legal and public policy teams, oversee the people we hire to help us scale our content moderation efforts, and provide the native language expertise and the 24-hour coverage required of a global platform. Google employs review teams across many offices globally and across the U.S. to ensure we have a diverse set of reviewers who are reviewing publisher sites, apps, and content.

We do not have to choose between safety and growing our business. In fact, our business depends on our providing a safe environment for users and advertisers. How YouTube lives up

to its publicly declared goals is a core part of the product, and much of what we do on the product team is to deliver on Trust and Safety objectives and live up to YouTube's responsibilities.

As we consider the YouTube product experience, we must take into account how our content policies impact our viewers and our creators. The Community Guidelines act as the "rules of the road" for what is and is not allowed, and we apply them consistently across YouTube. This means the policies and how they work are a core part of the overall product experience for our entire ecosystem — including our viewers, creators, and advertisers. Our responsibility relies on using our systems to connect viewers to high-quality information and to minimize the chances they'll see problematic content. It also involves assessing how creators—the vast majority of whom are looking to do the right thing — are rewarded. These levers work together to form a comprehensive approach to responsibility.

We prioritize responsibility even when it comes at a cost. Our advertiser-friendly content guidelines and the YouTube Partner Program set out YouTube's revenue sharing program, and if a creator fails to comply with those guidelines, we cut them off from making money on the platform. When we do that, YouTube foregoes revenue, as well.

Question 3. In a response to questions for the record I posed last November, Facebook said that it has been building a "parallel viral content review system to flag posts that may be going viral—no matter what type of content it is" and prioritize these posts for review. The idea is that potential virality should be a signal that content would be particularly dangerous if it violates platform policies, and thus warrants a trigger for human review.

- a. Does YouTube incorporate potential virality in reviewing content?**
- b. Please explain the factors YouTube considers when determining which content to review or fact-check.**
- c. How long does it take on average for YouTube videos to be fact-checked?**
- d. What is YouTube's process for providing corrections when it learns someone watched a video that later proved to be misinformation or violative?**

YouTube's Community Guidelines outline what content is permissible and what content is violative on our platform. As noted above in our responses to Questions 1 and 2, these policies act as our "rules of the road" of what is and is not allowed, and we enforce our policies globally, independent of viewpoints or audience size.

For a number of years, the misinformation landscape online was dominated by a few main narratives — such as 9/11 "truthers," moon landing conspiracy theorists, and so-called "flat earthers." We were able to train our machine learning systems to reduce recommendations of those videos and other similar ones based on patterns in that type of content. But increasingly, a completely new narrative can quickly crop up and gain views. Or, narratives can slide from

one topic to another — for example, some general wellness content can lead to vaccine hesitancy. Each narrative can also look and propagate differently, and at times, even be hyperlocal.

We faced these challenges early on in the COVID-19 pandemic, such as when a conspiracy theory that 5G towers caused the spread of coronavirus led to people burning down cell towers in the UK. Due to the clear risk of real-world harm, we responded by updating our guidelines and making this type of content violative. In this case we could move quickly because we already had policies in place for COVID-19 misinformation based on local and global health authority guidance as well as our longstanding policies against inciting violence.

But not every fast-moving narrative in the future will have expert guidance that can inform our policies. And the fresher the misinformation, the fewer examples we have to train our systems. To address this, we continuously train our systems on new data.

In addition to reducing the spread of some content, our systems connect viewers to authoritative videos in search results and recommendations. For major news events that are developing, like a natural disaster, we surface developing news panels to point viewers to text articles. We also surface, in certain countries, fact checks when you search for a specific claim, or contextual information panels when you search for or watch videos related to topics prone to misinformation. We give viewers this context to help them make their own informed decisions about the videos they are watching. When viewers search YouTube for something related to a specific claim, sometimes we provide an information panel that includes a fact check from an independent third-party publisher. These information panels indicate whether claims related to your search are true, false, or something else like "partly true," according to the publisher's fact check. More information about our fact checking is available at <https://support.google.com/youtube/answer/9229632>, and more information about our information panels is available at <https://support.google.com/youtube/answer/9229632>.

We strive to remove violative content as quickly as possible. Further, as detailed in the YouTube Community Guidelines Enforcement Transparency Report (available at <https://transparencyreport.google.com/youtube-policy/removals>), in the second quarter of 2022, more than 4.4 million videos were removed from YouTube for violating our Community Guidelines — 93% of which were first flagged by machines rather than humans. Of those detected by machines, over 34% never received a single view, and just over 72% received fewer than 10 views.

Question 4. I appreciate YouTube’s new reporting of a “violative view rate,” which is an important estimation of what is not being caught and thus being viewed by users. To what extent can YouTube calculate this metric at a level of granularity that allows it to drive decision making about product features? For instance, can YouTube estimate this rate as to certain regions, or use it in potential A/B testing of products?

As discussed in our answers to Questions 1 and 2, in April 2021 we released a new metric called Violative View Rate (VVR) as part of our Community Guidelines Enforcement Report. It is now the primary metric used to measure our responsibility work across YouTube. We will continue to publish this metric so we can be held accountable to reducing it over time.

As noted in the responses above, we calculate VVR by taking a sample of videos on YouTube and sending it to our content reviewers who tell us which videos violate our policies and which do not. By sampling, we gain a more comprehensive view of the violative content we might not be catching with our systems.

Question 5. I appreciate your testimony that YouTube is committed to increasing transparency about your algorithmic recommendation engines. Many independent researchers have said it is essential to know not only what “signals” are considered but also how those signals are weighted; otherwise, it is not apparent which signals are actually driving the ranking.

- a. Please provide more information about the inputs, signals and weightings used in YouTube’s recommendation algorithm.**
- b. Will YouTube provide greater transparency going forward into its weightings of different inputs for its recommendation algorithm, whether through reporting, granting access to independent researchers, or conducting an audit?**

Our approach to recommendations is anchored in transparency and user choice and control. We have a number of resources available and are working on ways to make our systems easier to understand. We want to offer meaningful transparency to policymakers and users. We have an informative “How YouTube Works” site (available at <https://www.youtube.com/howyoutubeworks/>) with information about how we raise results and make improvements. We also provide significant resources to creators and work hard to inform creators about the ways our systems work and how to be successful on YouTube. This includes interactive websites, detailed Help Centers, videos, frequent blog posts, and user forums (examples of which are available at <https://www.youtube.com/channel/UChRfArvrzheW2E7b6SVT7vQ> and <https://support.google.com/youtube/answer/9002587>).

It is important to note that our recommendation system does not operate off of a “recipe book” of what to do. It constantly evolves, learning every day from over 80 billion pieces of

information we call signals. Providing more transparency is not as simple as listing a formula for recommendations, but involves understanding all the data that feeds into our system. A number of signals build on each other to help inform our dynamic system about what you find satisfying: clicks, watchtime, survey responses, sharing, likes, and dislikes (more information about these signals is available at

<https://blog.youtube/inside-youtube/on-youtubes-recommendation-system/>).

- Clicks: Clicking on a video provides a strong indication that users will also find it satisfying. After all, a user would not click on content they do not want to watch.
- Watchtime: Watchtime — which videos a user watched and for how long — provides personalized signals to our system about what users most likely want to watch.
- Survey Responses: To ensure viewers are satisfied with the content they watch, we measure what we call “valued watchtime” — the time spent watching a video that viewers consider valuable. We measure valued watchtime through user surveys that ask viewers to rate the video watched from one to five stars, giving us a metric to determine how satisfying viewers found the content. If a video is rated one to two stars, we ask the viewer why they gave such a low rating. Similarly, if the viewer gives the video four to five stars, we ask why — was it inspirational or meaningful? Only videos that are rated highly with four or five stars are counted as valued watchtime.
- Sharing, Likes, Dislikes: On average, people are more likely to be satisfied by videos that they share or like. Our system uses this information to try to predict the likelihood that viewers will share or like further videos.

We also empower our users by giving them control over personalized recommendations. As described in our response to Question 1, in June 2019, we launched new features that give users more control over what recommendations appear on the homepage and in their “Up Next” suggestions. These features make it easier for users to block channels from recommendations and provide users more information about why we are suggesting a video.

Question 6. In general, YouTube does not disclose data to the public regarding how many times any video has been recommended by its recommendation algorithm. This means that when YouTube removes videos from its platform for violating its content policies, the public has no way of knowing how many times those videos were recommended before their removal.

- a. During last Tuesday’s hearing, I asked whether YouTube could commit to disclosing this information publicly. You said this was an “interesting idea” and agreed to follow up. Can YouTube now commit to disclosing this information? If not, why not?**
- b. In general, does YouTube collect data regarding the number of times that videos on the platform have been recommended by its algorithm?**
- c. Does YouTube currently share that data with anyone outside the company?**

YouTube has been a leader in transparency concerning how we operate our platform and when we enforce our Community Guidelines. We publish a quarterly YouTube Community Guidelines Enforcement Transparency Report (available at <https://transparencyreport.google.com/youtube-policy/removals>). As noted in some of our other answers, we strive to remove violative content as quickly as possible.

Our commitment to openness means we may have “borderline” content available on the platform that comes close to violating our policies but does not quite cross the line. It is important to protect freedom of expression, but we set a high bar for what videos we display prominently in our recommendations on the YouTube homepage or through the “Up Next” panel. We constantly work to improve our systems to reduce recommendations on borderline content and videos that could misinform users in a potentially harmful way, particularly in areas that rely on veracity — such as news, medicine, historical events, or science.

To determine borderline content, evaluators assess factors that include, but are not limited to, whether the content is: inaccurate, misleading or deceptive; insensitive or intolerant; and harmful or with the potential to cause harm. The results are combined to give a score for how likely it is that the video contains harmful misinformation or is borderline. Based on the consensus input from the evaluators, we use well-tested machine learning systems to build models. These models help review hundreds of thousands of hours of videos every day in order to find and limit the spread of borderline content. And over time, the accuracy of these systems will continue to improve. Today, consumption of borderline content that comes from our recommendations represents less than 1% of the content watched on YouTube. We have now launched changes to our recommendation system to reduce recommendations of borderline content in every market where we operate.

To encourage further research, we recently launched the YouTube Research Program, which gives academic researchers access to our data in a user privacy sensitive way, allowing them to derive metrics or develop insights of their own based on that data (more information is available at <https://research.youtube/>).

Question 7. During last Tuesday’s hearing, I asked whether YouTube generally requires employees to sign non-disclosure agreements (NDAs). You said you believe that YouTube has “standard agreements to protect proprietary information.”

- a. Please briefly describe the scope of the NDAs that YouTube employees are required to sign, including the types of information that employees and former employees are barred from disclosing.**
- b. Do the NDAs cover more than trade secrets, private user data, and/or technical information that employees learn while employed by YouTube?**

- c. Do the NDAs or any other contractual agreements bar employees from criticizing or disparaging YouTube after they leave the company?
- d. If an employee disagrees with a policy decision made by YouTube during the employee's period of employment, would the employee's NDA prohibit him or her from discussing that disagreement publicly after leaving the company?
- e. Do YouTube's NDAs contain binding arbitration clauses—meaning employees must submit to arbitration in the event of any legal dispute under the terms of the agreement?
- f. Has YouTube ever threatened legal action against an employee (current or former) based on potential violations of that employee's NDA?

In 2020, Google and Alphabet updated their workplace codes of conduct consistent with the following five guiding principles:

- **Commitment:** We will double down on our commitment to a respectful, safe, and inclusive working environment for all employees and members of our extended workforce.
- **Care:** We will continue to create an environment that emphasizes respect for each individual at all levels of the organization.
- **Transparency:** We will be open and transparent as an organization regarding the frequency of complaints of misconduct and our approach to investigating and responding to those allegations.
- **Fairness & Consistency:** We will make sure that everyone is treated respectfully, fairly, and compassionately, and we'll apply our employment policies consistently.
- **Accountability:** We will hold everyone responsible for their actions.

Google's employment agreement, which includes a NDA, does not prohibit the disclosure of facts underlying claims of harassment or discrimination. Specifically, its employment agreement provides that "nothing in this Agreement limits any right I may have to discuss terms, wages, and working conditions of employment, as protected by applicable law."

In addition, Google has made arbitration optional for all employees, temporary staff, vendors, and independent contractors with respect to harassment, discrimination, and retaliation disputes with Alphabet. It has also committed to annually review its sexual harassment and retaliation policies and practices to ensure consistency with the guiding principles.

Question 8. You stated in your written testimony that YouTube “disincentivize[s] creators seeking to profit off of misinformation by blocking ads on their content and suspending repeat offenders from our monetization program.” How does YouTube determine whether a content creator is a “repeat offender,” particularly with regard to misinformation?

- a. Does YouTube ever take public perception into consideration when determining whether or how to enforce its repeat offender policies? If so, how?**
- b. Are potential political repercussions ever taken into consideration when determining whether or how to enforce YouTube’s repeat offender policies? If so, how?**
- c. Are potential business consequences—such as potential dissatisfaction from advertisers—ever taken into consideration when determining whether or how to enforce YouTube’s repeat offender policies? If so, how?**

We believe strongly in the freedom of expression and access to information, and we know that the overwhelming majority of creators follow our guidelines and understand they are part of a large, influential, and interconnected community. However, we also know that we have a responsibility to protect our users, which is why we have policies prohibiting hate speech, terrorist content, and other content that violates our policies, as well as stricter standards for who can monetize their content. Each of the products and services we offer has a different purpose, and we tailor our approach carefully to the content that should be available on each product and service. However, our search and recommendation systems are not designed to filter or demote videos or channels based on specific political perspectives. We take extraordinary lengths to build our products and enforce our policies impartially.

While YouTube creates a space for ideas and expression, it is not a free-for-all. For example, it is a violation of YouTube’s hate speech policy for users to post videos that promote violence against particular ethnic or religious groups (for more information, please see <https://support.google.com/youtube/answer/2801939>). As described in our response to Question No. 2, creators who violate those rules may have their content removed or their accounts terminated. When we detect a video that violates our Community Guidelines, we remove the video and apply a strike to the channel. The strike restricts a creator’s ability to post or create content on the platform for one week. If the creator’s behavior warrants another strike within 90 days from the first, a new two-week prohibition from posting or creating content is implemented. A third strike within 90 days results in permanent removal of a channel from YouTube. Creators can appeal those strikes if they believe we are mistaken. We also terminate entire channels if they are dedicated to posting content prohibited by our Community Guidelines or contain a single egregious violation, like child sexual abuse material.

As referenced in our response to Questions 1 and 2, we reward our trusted creators through monetization programs. Creators who earn money on YouTube must follow YouTube’s channel

monetization policies (available at <https://support.google.com/youtube/answer/1311392>), which include YouTube's Community Guidelines, Terms of Service, Copyright, and Google AdSense program policies. Violation of our YouTube channel monetization policies may result in monetization being suspended or permanently disabled. Additionally, creators who follow our policies can apply to join the YouTube Partner Program, a program that sets a higher bar for which channels can make money on our site (for more information on our Partner Program, please see <https://support.google.com/youtube/answer/72851>).