

**GOOGLE'S SUBMISSION IN RESPONSE TO
SUBCOMMITTEE QUESTIONS FOR THE RECORD
FOLLOWING March 25, 2021 HEARING**

The Honorable Jan Schakowsky (D-IL)

1. Please provide our committee with a list within 30 days of all the Universities, think tanks, NGOs, publications, journalists, activists, and activist groups that Google or YouTube currently supports financially or has provided more than \$50,000 in funding to in the past five years. This should include all grants or other payments paid through a third party that ultimately produced reports, studies, editorials, or other publicly available content.

Google actively donates and engages with a wide range of organizations every year. For example, through Google.org, we grant more than \$200M to nonprofits and social enterprises across the globe every year. Our goal is to use our philanthropic capital to help stimulate innovative approaches to solving problems, and provide comprehensive support for marginalized communities.

Through Google.org Impact Challenges, we award community-driven nonprofits and social enterprises with support to make their community, and beyond, a better place. Each challenge is an invitation for ideas that tackle a specific problem or support a particular community, whether that is a city, continent, or the entire world. We look for nonprofits and social enterprises with ideas that can create change at the pace and scale the world needs today. The best and boldest ideas are given a strategic package of funding, mentorship, and technical support. For information on our past challenges and winners, please see <https://www.google.org/opportunities/>.

Since 2003, Google Ad Grants has provided nonprofits with up to \$10,000 per month in free Search ads to help them attract donors, recruit volunteers, and promote their missions. In June 2020, we pledged an additional \$200M in Ad Grants for a total of \$1B in 2020 (<https://blog.google/outreach-initiatives/nonprofits/1-billion-ad-grants/>).

Throughout the pandemic, we have been actively working on ways to support communities and small businesses in the U.S. and around the world. In March 2020, we announced a \$125M Grow with Google Small Business Fund (<https://bit.ly/3tBVJFw>). In partnership with Opportunity Finance Network (OFN), the fund provides low-interest loans to community development financial institutions (CDFI), who in turn provide loans to small businesses in underserved communities in the U.S. Google.org also made a \$5M grant to OFN to further support CDFIs as they grow their capital and build their capacity during this time of crisis. In June 2020, we expanded these donations, adding \$45M in loans to the fund and \$5M in Google.org grants to OFN, with a specific focus on Black communities (<https://blog>.

[google/outreach-initiatives/grow-with-google/small-business-fund-cdfi](https://www.google.com/outreach-initiatives/grow-with-google/small-business-fund-cdfi)). Information on awardees and grantees can be found on OFN's website here: <https://ofn.org/googlesmallbizfund#Awardees> and <https://ofn.org/google-org-grant-program#Grantees>.

We have also supported health organizations, governments, and health workers throughout the pandemic. That has included more than \$550M in Ad Grants to help more than 100 government agencies globally provide critical information on how to prevent the spread of COVID-19 and other measures to help local communities, as well as \$100M in different COVID-19 relief projects through Google.org. For more information on those initiatives please see <https://www.blog.google/inside-google/company-announcements/commitment-support-small-businesses-and-crisis-response-covid-19/> and <https://blog.google/outreach-initiatives/google-org/100-million-dollar-contribution-covid-19-relief/>.

We have also invested heavily in organizations addressing racial inequity. In June 2020, we announced various commitments and programs dedicated to racial equity, including a more than \$175M economic opportunity package to support Black business owners, startup founders, job seekers, and developers, in addition to YouTube's \$100M fund to amplify Black creators and artists (<https://www.blog.google/inside-google/company-announcements/commitments-racial-equity/>). Part of that commitment included a \$12M pledge to support organizations working to advance criminal justice reform, including The Leadership Conference Education Fund, the NAACP Legal Defense and Educational Fund's Policing Reform Campaign, and the Movement for Black Lives. For information on our commitment to advancing racial justice, please see <https://www.google.org/inclusion/racial-justice/>.

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 available at <https://www.google.com/publicpolicy/transparency/>.

2. Does Google or YouTube share data on disinformation with independent scholars or researchers? If so, what information is shared, with whom is it shared, and what is the process for sharing it?

We believe transparency is essential to earning and sustaining the trust of our users and our business partners. And 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 launched the first Threat Analysis Group (TAG) Bulletin (<https://blog.google/threat-analysis-group>). 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 such groups, including researchers working in this space, and that our information sharing can help confirm findings from security firms and other industry experts.

Second, in June 2020, YouTube launched a website called How YouTube Works (<https://www.youtube.com/howyoutubeworks>), 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, we publicly share significant data on actions we take to enforce our policies. For example, YouTube publishes quarterly data in our Community Guidelines enforcement report (<https://transparencyreport.google.com/youtube-policy/removals>). 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. In fact, we recently added a new metric to this report known as Violative View Rate (VVR) (<https://transparencyreport.google.com/youtube-policy>). This metric is an estimate of the proportion of video views that violate our Community Guidelines in a given quarter (again, excluding spam). Our data science teams 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.

In addition to sharing information externally, we also consult with a diverse set of external and internal stakeholders during our policy development, which can include expert input, user feedback, and regulatory guidance. This collaborative approach taps into multiple areas of expertise within and beyond our company and is typically driven by our Trust & Safety teams, whose mission includes tackling online abuse by developing and enforcing the policies that keep our products safe and reliable. These teams include product specialists, engineers,

lawyers, data scientists, and others who work together around the world and with a network of in-house and external safety and subject matter experts.

Where appropriate, these teams consult in-depth studies or research by a mix of organizations, academics, universities, or think tanks who have topical expertise in specific matters. These analysts study the evolving tactics deployed by bad actors, trends observed on other platforms, and emerging cultural issues that require further observation. We also engage in conversations with regulators around the world, and their perspectives and concerns directly inform our policy process.

3. Clarify who is responsible for Spanish-language content moderation in the U.S. Please indicate which specific team is responsible for that work and share the title of the manager who most directly carries that responsibility. In parallel, please also indicate the title of the manager who is responsible for efforts to counter Spanish-language disinformation in the U.S. and the name of the relevant team working on this issue. In both instances, please identify where these teams and managers sit in the organizational chart and to whom they report.

a. Furthermore, clarify whether Spanish-language content in the U.S. is a specifically identified and assigned responsibility. Do the teams and managers involved have specific responsibility for Spanish-language content in the U.S. or are they responsible for all Spanish-language content on the platform globally?

We have clear, public-facing policies that outline what is allowed on our services, and what is not. Our policies are not language-specific, and we do not differentiate when it comes to content on our platforms. We apply our policies equally and protect our users no matter where they are — in the United States or in a foreign country (including where Spanish is the predominant language). We use both automated and manual detection and enforcement efforts to protect our users. Content moderation at Google is primarily managed by Trust & 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. These teams also work with our legal and public policy teams, and oversee the vendors we hire to help us scale our content moderation efforts, as well as provide the native language expertise and the 24-hour coverage required of a global platform.

Our business model depends on us being a useful and trustworthy source of information for everyone. We will continue to be vigilant to prevent deceptive content, no matter what language it is in.

4. Phrases and symbols associated with minority communities are systematically more likely to be flagged as hate speech on social media platforms, regardless of how benign the content is. The Sikh community has noticed the targeted removal of benign posts that reference Sikh religious artifacts, political speech, and community organizing. How

will YouTube create channels of recourse that will prevent the structural and systematic suppression of content from minority communities? (Ex. YouTube’s removal of all content related to the Indian Farmers’ Protest from its platform.)

To enforce our content moderation policies at the scale of the web, we use a combination of human reviewers and cutting-edge machine learning to combat violent and extremist content. We’ve dedicated significant resources to content moderation efforts: we spent nearly 1.2B dollars on content moderation efforts in 2020. We also have nearly 22,000 employees dedicated to ensuring the protection of our platforms.

We’re constantly innovating to improve our machine learning and algorithms to spot content in violation of our policies. Those innovations, of course, also improve our ability to distinguish between violations and benign content. Among other AI principles, we believe AI should be socially beneficial, avoid creating or reinforcing unfair bias, and be accountable to people. We continue to be one of the leading companies investing heavily in responsible AI and implementation. This includes hundreds of people working in culture, education, and participation; technical progress; internal review processes; and, community outreach and exchange. More information on our AI Principles and progress can be found in our 2020 progress update report here: <https://ai.google/static/documents/ai-principles-2020-progress-update.pdf>.

We have taken multiple approaches to mitigating discriminatory bias in machine learning systems (MLs), including by promoting and supporting research in this area and developing and releasing tools and resources to mitigate discriminatory bias in MLs directed at both Google’s own products and more generally. For example, the What-If Tool lets users analyze an ML model without writing code, giving them the ability to test how the code performs for different groups, including historically marginalized people. For more information, please see <https://ai.googleblog.com/2018/09/the-what-if-tool-code-free-probing-of.html>.

We have similarly developed the Responsible AI Toolkit (https://www.tensorflow.org/responsible_ai), a suite of tools we use internally and have open sourced to, among other things, improve fairness and eliminate or mitigate bias in MLs. We have further mitigated potential bias in our own products by refining the Perspective API developed by Alphabet subsidiary Jigsaw to more accurately identify toxic language in online forums (<https://developers.google.com/machine-learning/practical/fairness-indicators>) and by making improvements to Google Translate to reduce gender bias (<https://ai.googleblog.com/>).

No system for filtering content is perfect, which is why we also provide users with an opportunity to appeal content decisions. For more information on the YouTube appeals process, please see <https://support.google.com/youtube/answer/185111?hl=en>. We are committed to constantly improving our systems to be more inclusive and will continue to spend significant time and resources on this important issue.

5. Children and young adults who are members of minority communities are plagued by uniquely impactful hate speech and cyberbullying on social media platforms. While more than 15% of youth social media users experiencing cyber bullying, people of color receive disproportionately more vitriolic messages that can have substantial impact on emotional and cognitive development. How has YouTube sought to provide recourse to individuals facing online harassment at the intersection of cyberbullying and racial/ethnic hate speech?

YouTube is committed to protecting our users from hate speech and cyberbullying. Our Community Guidelines set the rules of the road for what content is and is not allowed on YouTube, and we review these guidelines to adapt to the evolving nature of potentially violative content. Our Child Safety policy, for example, includes specific Cyberbullying and harassment involving minors provisions that prohibit content that targets individuals for abuse or humiliation; reveals personal information like email addresses or bank account numbers; records someone without their consent; sexually harasses; or encourages others to bully or harass (<https://support.google.com/youtube/answer/2801999?hl=en>). We have a zero tolerance policy for predatory behavior on YouTube. If we think a child is in danger based on reported content, we'll help law enforcement investigate the content.

More generally, we have policies that prohibit content that targets an individual with prolonged or malicious insults based on intrinsic attributes, as well as policies for users to report inappropriate content. For more information about our policies around harassment and cyberbullying, along with how to report it, please see <https://support.google.com/youtube/answer/2802268?hl=en> and <https://support.google.com/youtube/answer/2802027>. There are also other policies that are helpful in combating evolving threats and in providing as much coverage as possible.

- **Harassment & Cyberbullying:** YouTube prohibits content that threatens individuals. We also don't allow content that targets an individual with prolonged or malicious insults based on intrinsic attributes. These attributes include their protected group status, which include 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. For example, we don't allow content in which someone simulates violence toward an identifiable individual, such as by using a weapon on a dummy representing that person.
- **Hate Speech:** We remove 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.

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 (<https://transparencyreport.google.com/youtube-policy/removals?hl=en>). In the fourth quarter of 2020 alone, we removed more than 65,000 channels, more than 77,000 videos, and over 136M comments for violating our harassment and cyberbullying policies. For more information, please see <https://transparencyreport.google.com/youtube-policy/removals?hl=en>.

6. How often do policy compliance standards change on YouTube, and who is involved in creating these policies? What is the role of international policy teams in determining global policies? How do international policies impact the policies of other nations? What safeguards does YouTube have to prevent the limitations from one country, such as India, from impacting users based in the United States?

Drafting content policies is a time-consuming process. It can take months before we feel comfortable taking action on a new policy. To help us identify emerging harms and gaps in our existing policies, we consider expert input, user feedback, and regulatory guidance. We rely on research performed by analysts who study the evolving tactics deployed by bad actors, trends observed on other platforms, and emerging cultural issues that require further observation. We also engage in conversations with regulators around the world. Their perspectives and concerns directly inform our policy process. This collaborative approach taps into multiple areas of expertise within and beyond our company and is typically driven by our Trust & Safety teams. Their mission includes tackling online abuse by developing and enforcing the policies that keep our products safe and reliable. The teams include product specialists, engineers, lawyers, data scientists, and others who work together around the world and with a network of in-house and external safety and subject matter experts.

We comply with the law in each country in which we operate and remove illegal content on our platforms in that country. In every country where we operate, the unique cultures, histories, and forms of government have produced different laws governing what is considered permissible expression. Of course, legal standards vary greatly 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/region where it is deemed to be illegal. We also report on such removals via our Transparency Report (<https://transparencyreport.google.com/government-removals/overview?hl=en>).

In addition, we develop and maintain rules of the road, which outline what types of content and behaviors are acceptable for each product or service. Known as Content Policies or Community Guidelines, we aim to make them clear and easily accessible to all users and content creators—whether those are video creators, webmasters, app developers, or advertisers. These rules of the road articulate the purpose and intended use of a given product or service and represent a crucial part of what makes that product unique. They also explain what types of content and behaviors are not allowed, and the process by which a piece of content, or its creator, may be removed from the service.

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

We also make our policy information and process clear and easily available to all. We develop comprehensive help centers, Community Guidelines websites, and blog posts that detail the specific provisions of our policies. In addition, we regularly release reports that detail how we enforce those policies or review content reported to be in violation of local law. For more information, please see our YouTube Community Guidelines enforcement report (<https://transparencyreport.google.com/youtube-policy/removals>), our Threat Analysis Group's Quarterly Coordinated Influence Operations Bulletin (<https://blog.google/threat-analysis-group>), and reports made available on the Google Transparency Report Website (<https://transparencyreport.google.com/>). Finally, as noted in our response to Question No. 2, in June 2020, YouTube launched a website called How YouTube Works (<https://www.youtube.com/howyoutubeworks>), 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.

7. Can you explain what exactly you have done to ensure YouTube's tools —algorithms, recommendation engines, and targeting tools — are not amplifying conspiracy theories and disinformation, connecting people to dangerous content, or recommending hate groups or purveyors of disinformation and conspiracy theories to people? For example, can you provide detailed answers on some of the Capitol riot suspects' use history, to include the following:

a. Of the videos the individuals watched with Stop the Steal content, calls to question the election, white supremacy content and other hate and conspiracy content, how many were recommended by YouTube to the viewer?

To create a safe environment for our users at YouTube, we approach our content moderation practices 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 lines (i.e., borderline content), and we reward trusted creators.

- We remove content that violates our policies. Our policies are updated regularly to meet new challenges and include comprehensive guidelines on hate speech, terrorism, and harassment, along with other harmful content. In Q4 2020 alone, we removed 9.3M videos for violating our community guidelines. Approximately 74% of those first flagged by our systems received 10 or fewer views. In Q4, we also removed over 72,000 videos for violating our policies against Violent Extremism. We have removed over 850,000 YouTube videos with dangerous or misleading COVID-19 information.
- We raise up authoritative voices when people are looking for breaking 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. For example, since 2018, we have seen 25M impressions on our QAnon info panel. On October 15, 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). And, 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, 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.
- We also reduce recommendations of borderline content. In January 2019 we announced this improvement to our recommendation system to greatly reduce recommendations of borderline content and content that could misinform users in harmful ways. This is content which comes close to but doesn’t quite violate our policies and represents less than 1% of the content that is watched on YouTube. The result was a 70% average drop in watch time of this content coming from non-subscribed recommendations in the U.S. that year. While algorithmic changes take time to ramp up and you might see consumption of borderline content go up and

down, our goal is to have views of non-subscribed, recommended borderline content below 0.5%. We have now rolled out changes to our recommendation system to reduce recommendations of borderline content in every market where we operate.

- Finally, 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.

In addition, our expert teams around the world handle the investigations of more sophisticated threat actors that attempt to circumvent the automated defenses we build into our products. New forms of abuse and threats are constantly emerging that require human ingenuity to assess and plan for action before an automated system can address them at scale. So we work with dedicated threat intelligence and monitoring teams—Google’s Threat Analysis Group—to provide insights and intelligence to our policy development and enforcement teams so they can stay ahead of bad actors. We also enable users and trusted organizations to flag content that may be problematic. Flagged content is reviewed by a member of our Trust & Safety teams who will decide whether the content warrants action. We have reviewers evaluating user flags 24/7 from teams based in offices around the world.

We’re constantly innovating to improve our machine learning and algorithms to spot content in violation of our policies. And we partner with a network of academics, industry groups, and subject matter experts to help us better understand emerging issues. We will continue to be vigilant in evaluating our responses to removing harmful content from our platforms.

8. Do you commit to working with the Alphabet Workers Union to combat extremism and other challenges?

At Google, we strongly believe in the freedom of expression, and we also feel a great responsibility to keep our users safe. We’ll continue engaging directly with all our employees on how best to combat violent and extremist content and protect our users.

9. On a recent Atlantic Council webinar, YouTube CEO Susan Wojcicki explained that YouTube did not implement a policy about election misinformation until after the states certified the election, on December 9. She said that starting then, a person could no longer allege the election was due to widespread fraud. First, this raises the obvious question: why did you wait until December 9?

a. She then continued to explain that due to a “grace period” after the policy was finally made, Donald Trump’s numerous violations did not count, and he only has one actual strike against him and will be reinstated when YouTube deems there is no longer a threat of violence. How will you make that assessment? How will YouTube decide that there is no longer a threat of violence? And does that mean you will allow Donald Trump, or

others with strikes against them, to reinstate their accounts and be allowed to continue spreading mis- and disinformation and conspiracy theories?

December 9, 2020 was not the beginning of our enforcement efforts against election-related content. Beginning long before election day, we worked hard to ensure the integrity of our elections. Our Threat Analysis team, for example, worked to combat influence operations before the election. And between September and December 2020, YouTube terminated over eight thousand channels and removed thousands of harmful and misleading election-related videos. We also raised authoritative content so that our users could readily find accurate information -- like where and how to vote. Our Info Panels with the election results have been shown more than 8B times.

On December 8, 2020, states across the country certified their election results for the first time. We relied on that certification as an independent third-party verification that the election results were accurate, and thus content challenging the accuracy of the results was removed pursuant to a new policy. Starting December 9, 2020, if we identified content on YouTube stating that the presidential election result was invalid as a result of fraud or irregularities, we removed it. This is in line with how we treat other U.S. presidential elections and we've consistently applied that policy to the 2020 U.S. presidential election since that time.

That approach was in place long before the events of January 6, 2021, and was applied to the videos that were uploaded by President Trump's channel. When President Trump uploaded a video on January 6, 2021 that violated our policies, that video was taken down. And when, on January 12, 2021, the same channel uploaded videos that we determined violated our incitement of violence policies, they were removed and the channel was suspended from uploading. We have always approached these issues carefully—balancing the importance of free speech, particularly in the political context, with the need to protect our users and our democracy.

We understand the need for intense scrutiny on our election-related work. We will keep engaging with experts, researchers, and organizations to ensure that our policies and products strike a balance between allowing for a broad range of political speech and making sure our platform isn't abused to incite real-world harm or to broadly spread harmful misinformation. And, as always, we'll apply learnings from this election to our ongoing efforts to protect the integrity of elections around the world. For more information on YouTube's efforts to protect the integrity of the 2020 elections, see our YouTube 2020 Election Retrospective, <https://www.youtube.com/howyoutubeworks/our-commitments/supporting-political-integrity/#election-news-and-information>.

10. According to numerous studies YouTube often performs worst of all major platforms at moderating and removing election misinformation. YouTube's Community Guidelines state that channels receiving three "strikes" in a 90-day period will be banned. However, after 90 days the strikes reset. Has YouTube evaluated whether the three "strikes"

policy allows repeat offenders to continue to post violative content by manipulating the 90-day period? How many channels have received three or more total strikes but were not removed because the strikes did not occur within a 90-day period? Please provide a list of these channels.

Our YouTube Community Guidelines prohibit spam, scams, deceptive practices and any content that seeks to incite violence. Since September 2020, we've terminated more than 8,000 channels and thousands of harmful and misleading election-related videos for violating our existing policies. More than 77% of those removed videos were taken down before they had 100 views. And since election day, relevant fact check information panels from third-party fact checkers were triggered more than 200,000 times above relevant election-related search results, including for voter fraud narratives such as "Dominion voting machines" and "Michigan recount." For additional information, please see <https://blog.youtube/news-and-events/supporting-the-2020-us-election/>.

To measure our effectiveness, and to create transparency around our enforcement, YouTube publishes quarterly data in our Community Guidelines enforcement report (<https://transparencyreport.google.com/youtube-policy/removals>), a section of the Google Transparency Report. The YouTube portion of the report provides public data about the number of videos we remove from our platform for each of our policy violation categories, as well as additional information about channel removals, views before removals, appeals and reinstatements, and human and machine flagging.

We also recently published a Retrospective on the 2020 U.S. Elections on YouTube (<https://www.youtube.com/howyoutubeworks/our-commitments/supporting-political-integrity/#election-news-and-information>), to provide a thorough overview of YouTube's work to support election integrity in the 2020 U.S. presidential and congressional elections.

11. Even before COVID-19 and as recently as January 2020, anti-vaccination groups and profiteers ran ads on YouTube leading to resurgences in diseases like measles that our country had all-but-eradicated.

a. What was Google's advertising revenue from anti-vaccination videos pre-COVID-19?

b. How have these numbers changed during the pandemic?

c. How many clicks for anti-vaccination conspiracy advertisements has YouTube generated?

A number of the policies and product features that we have used to address the COVID-19 crisis were already in place before the crisis began, and others were underway. For example, our ranking systems on Google Search and YouTube have been designed to elevate authoritative information in response to health-related searches for years. Before 2020, YouTube's advertiser content guidelines (<https://support.google.com/youtube/answer/6162278>)

already prohibited “harmful health or medical claims or practices,” and our work to update our YouTube recommendation systems to decrease the spread of misinformation, including, but not limited to, health-related misinformation, was announced in January 2019. For more information, please see <https://youtube.googleblog.com/2019/01/continuing-our-work-to-improve.html>.

Since the outbreak of COVID-19, our efforts have focused on keeping people informed with trusted and authoritative information, supporting people as they adapt to the current situation, and contributing to recovery efforts. To help ensure that people are well informed, we have taken multiple steps to organize and provide accurate and verifiable information on the pandemic. These efforts to fight misinformation across our platforms include our Homepage “Do the Five” promotion, amplifying authoritative voices through ad grants (<https://support.google.com/google-ads/answer/9803410>), and launching our COVID-19 site (https://www.google.com/intl/en_us/covid19/), which includes COVID-19 information, insights, and resources.

We have continued to adapt our policies over time to address the challenges of the pandemic. For instance, YouTube 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 (<https://support.google.com/youtube/answer/9891785?hl=en>). In October 2020, we further expanded the policy to include COVID-19 vaccine misinformation.

We also raise up content from authoritative sources such as the CDC and other authorities to help users get the latest COVID-19 information. With anti-vaccination content, for example, we elevate reliable information across both Google and YouTube regarding medical topics (including vaccination) from trustworthy sources, such as health authorities.

Meanwhile, with respect to COVID-related ads, our Ads policies (<https://support.google.com/adspolicy/answer/6008942>) are designed not only to abide by laws but to ensure a safe and positive experience for our users. This means that our policies prohibit some content that we believe to be harmful to users and the overall advertising ecosystem. This includes policies that prohibit ads for counterfeit products, dangerous products or services, or dishonest behavior, and any content that potentially seeks to profit from a sensitive event such as a public health emergency like the COVID-19 global health crisis. For more information on our policies, such as those addressing sensitive events, please see <https://support.google.com/adspolicy/answer/6015406?hl=en>. In addition, our dangerous or derogatory content policy (<https://support.google.com/adspolicy/answer/6015406>) prohibits content in ads that would advocate for physical or mental harm, such as content that denies the efficacy of vaccines, as well as content that relates to a current, major health crisis and contradicts authoritative scientific consensus. As a result, content contradicted by scientific consensus during COVID-19 such as origin theories, claims the virus was created as a bioweapon, as well as claims the virus is a hoax or government-funded are not permitted on our platform.

And these efforts to limit the spread of COVID-19 misinformation are working. There have been over 400B impressions on our information panels for COVID-19 related videos and searches, and we've removed more than 850,000 videos for violating our COVID-19 misinformation policies. And in 2020, we blocked or removed over 275M COVID-related ads (globally), including Shopping ads, for policy violations including price-gouging, capitalizing on global medical supply shortages, making misleading claims about cures, and most recently, fake vaccine doses.

We are proud of our efforts to combat health misinformation and address this unprecedented public health crisis. We will continue to work hard and do everything we can to help our communities in addressing this global pandemic.

12. Your testimony cited YouTube's success at removing COVID-19 and vaccine-related disinformation. At the same time, white supremacist content, racist, and other hate content is running rampant on the platform. Is there anything inherently more difficult about policing white supremacist, racist, and other hate content compared to health misinformation and disinformation? Why?

We believe strongly in the freedom of expression and access to information—we know that the overwhelming majority of creators follow our guidelines and understand that they are part of a large, influential, and interconnected community. But 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.

We also strive to be a safe and inclusive space for all of our users. Improvements are happening every day, and we will continue to adapt, invent, and react as hate and extremism evolve online. We're committed to this constant improvement, and the significant human and technological investments we're making demonstrate that we're in it for the long haul. We've dedicated significant resources to content moderation efforts: we spent nearly 1.2B dollars on content moderation efforts in 2020. We also have nearly 22,000 employees dedicated to ensuring the protection of our platforms.

One of the most complex and constantly evolving areas we deal with is hate speech. That is why we systematically review and re-review all of our policies to make sure we are drawing the line in the right place, often consulting with subject matter experts for insight on emerging trends. For our hate speech policy, we work with experts in subjects like violent extremism, supremacism, civil rights, and free speech from across the political spectrum.

Hate speech is a complex policy area to enforce at scale, as decisions require nuanced understanding of local languages and contexts. To help us consistently enforce our policies, we have expanded our review team's linguistic and subject matter expertise. We also deploy

machine learning to better detect potentially hateful content to send for human review, applying lessons from our enforcement against other types of content, like violent extremism.

We also have recently taken a tougher stance on removing hateful and supremacist content from YouTube. Since early 2019, we've increased by 46 times our daily hate speech comment removals on YouTube. Thanks to this update and our ongoing enforcement, in Q4 2020 alone, we removed 175,000 videos and more than 182M comments for violating these policies. For additional information regarding enforcement of, and improvements to, our hate speech policies, please see: <https://blog.youtube/news-and-events/make-youtube-more-inclusive-platform/>, <https://transparencyreport.google.com/youtube-policy/featured-policies/hate-speech>, and <https://blog.youtube/news-and-events/our-ongoing-work-to-tackle-hate>.

Additionally, in October 2020, we launched a Community Guidelines YouTube update on harmful conspiracy theories (<https://blog.youtube/news-and-events/harmful-conspiracy-theories-youtube/>), which expanded our hate speech and harassment policies to prohibit content that targets an individual or group with conspiracy theories that have been used to justify real-world violence. For example, content such as conspiracy theories saying individuals or groups are evil, corrupt, or malicious based on protected attributes (e.g., age, race, religion, etc.), or hateful supremacist propaganda, including the recruitment of new members or requests for financial support for their ideology, all violate our hate speech policy (<https://support.google.com/youtube/answer/2801939>) and are subject to removal as such.

The openness of our platforms has helped creativity and access to information thrive. It's our responsibility to protect that, and 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.

The Honorable Bobby L. Rush (D-IL)

1. The last time you testified before Congress in October 2020, the Leadership Conference on Human and Civil Rights submitted a letter in the record calling for all tech companies to submit to a civil rights audit.

a. Has Google or any of its subsidiaries completed a civil rights audit?

b. If so, when will the results of the audit be made public? If no, do you plan to conduct such an audit?

Google is deeply committed to civil rights across our company and in our products. Our work on civil rights is conducted through our Human Rights Program, a central function responsible for ensuring—across Google and all its products (such as hardware, Search, Cloud, and YouTube)—that we are meeting our commitment to the United Nations Guiding Principles on

Business and Human Rights, Global Network Initiative Principles, and other civil and human rights instruments. Our senior management oversees the implementation of the civil rights work and provides regular updates to the Audit and Compliance Committee of Alphabet's Board of Directors. In 2020, the Board amended the charter of the Audit and Compliance Committee to explicitly include oversight of civil rights issues.

We incorporate civil rights principles into our long-term strategies and day-to-day decision-making. Google has product-specific counsel advising product teams who are part of the development and deployment of new products and features. Product and regional counsel, in coordination with subject-matter and regional experts among Google's policy staff, assess legal and policy risks. We also have a Human Rights Program which manages human rights due diligence, including human rights impact assessments, work with internal stakeholders in areas such as data governance, content policies, and supply chain.

Engaging with external experts and affected stakeholders is essential to our civil rights work. Regular engagement and formal consultation with civil society and other stakeholders informs the development of our products and policies. These engagements help us identify, prioritize, and address existing and potential civil rights impacts. They also present important opportunities for Google to get feedback on how and where we should consider improvements to our policies, practices, and services.

For instance, we've improved our algorithms to serve all users and reduce stereotypical representations of people and other forms of offensive results—this work continues. We use filters to block offensive language from appearing and continually update and improve these filters to block as much offensive content as possible. And as part of a collaboration with HUD, last year we launched a new policy to prohibit housing, employment, and credit advertisers from targeting ads based on age, gender, family status, or ZIP Code, in addition to our longstanding policies prohibiting targeting based on sensitive categories like race.

Beyond content moderation, we are also strengthening our commitment to diversity, equity, and inclusion to help Google build more helpful products for our users and the world, and we recently released an update on our progress towards our racial equity commitments (<https://blog.google/inside-google/company-announcements/progress-racial-equity-commitments/>). Last June, we announced our goal to improve leadership representation of underrepresented groups by 30% by 2025, increasing our investment in diverse talent markets such as Atlanta, Washington D.C., Chicago, and London. And this past October, we expanded on our commitments, including setting a goal to spend \$100M with Black-owned businesses, and committing to adding an additional 10,000 Googlers across our sites in Atlanta, Washington D.C., Chicago, and New York by 2025.

We also believe in the importance of transparency and are proud of the tools, systems, and processes that we currently have in place to examine and share our efforts. For example, we were the first company in Silicon Valley to publicly release our workforce diversity data, setting

an example that other tech companies later followed. Since 2014, we have regularly shared our progress on diversity, even when that progress was less than we had hoped to see. In fact, in subsequent years, we've actually erred on the side of sharing more data—disaggregating by intersectional identities and working closely with internal and external experts to dig deeper into our data to inform our planning and resourcing.

But we are not stopping there. We are committed to continuing and deepening our engagement with external experts, and investing in and embedding our internal experts as we address these issues. Last year, in addition to announcing our racial equity commitments, we were the first Silicon Valley company to elevate our internal leadership civil and human rights to head our civil and human rights program. We also launched our Human Rights Executive Council.

2. What is Google doing to ensure that you have appropriate staffing and mechanisms to avoid AI bias?

Google is committed to the responsible development and use of AI, including protecting against bias. In 2018, we announced Google's AI Principles (available at <https://ai.google/principles>), which govern our research and product development and impact our business decisions in this area. Specifically, Google believes AI should: be socially beneficial; avoid creating or reinforcing unfair bias; be built and tested for safety; be accountable to people; incorporate privacy design principles; uphold high standards of scientific excellence; and be made available for uses that accord with these principles. We also identify four areas we will not pursue—technologies: (1) that are likely to cause overall harm; (2) whose principal purpose is to cause or directly facilitate injury to people; (3) that gather or use information for surveillance violating internationally accepted norms; and (4) whose purpose contravenes widely-accepted principles of international law and human rights.

We continue to be one of the leading companies investing heavily in responsible AI research and implementation to ensure we can live up to our AI Principles. This includes hundreds of people working in culture, education, and participation; technical progress; internal review processes; and community outreach and exchange. More details can be found in our AI Principles 2019 and 2020 progress update reports (available at <https://ai.google/static/documents/ai-principles-2019-progress-update.pdf> and <https://ai.google/static/documents/ai-principles-2020-progress-update.pdf>).

Google has also taken multiple approaches to mitigating discriminatory bias in large-scale language models, including by promoting and supporting research in this area. We have developed and released tools and resources to mitigate discriminatory bias in large-scale language models directed at both Google's own products, and more generally, including, for example, the What-If Tool (please see <https://ai.googleblog.com/2018/09/the-what-if-tool-code-free-probing-of.html>), and the Language Interpretability Tool (please

see <https://ai.googleblog.com/2020/11/the-language-interpretability-tool-lit.html>), both of which are available to the general public. We have similarly developed the Responsible AI Toolkit (please see https://www.tensorflow.org/responsible_ai), which is a suite of tools we use internally and have open sourced to, among other things, improve fairness and eliminate or mitigate bias in machine learning systems. And we have mitigated potential bias in our own products, such as by refining the Perspective API developed by Alphabet subsidiary Jigsaw to more accurately identify toxic language in online forums (please see <https://developers.google.com/machine-learning/practical/fairness-indicators>) as well as making its process and findings available to developers so that others could learn from the company's experience and by making improvements to Google Translate to reduce gender bias (please see <https://ai.googleblog.com/2020/04/a-scalable-approach-to-reducing-gender.html>).

We're committed to building a workforce that is more representative of our users and a workplace that creates a sense of belonging for everyone. As part of that commitment, we are focused on supporting Black+ Googlers throughout their careers. One goal is to improve leadership representation of underrepresented groups by 30% by 2025; a related goal is to more than double the number of Black+ Googlers at all other levels by 2025. We have set aspirational goals for other underrepresented groups as well. These goals will be tracked and reviewed by the Alphabet Board.

With respect to representation in our AI work, while there are dozens of Black+ and Hispanic/Latinx+ Googlers on the AI team and hundreds of female Googlers on the AI team, we continue to expand this representation. We are working to attract additional talent and build out tech education and partnerships in New York City and Atlanta, which offer more diverse talent pools. We continue to invest in the external research community to support new voices entering into the field of computer science and more equitable outcomes. In 2020, we supported external research with \$37M in funding, including \$8M in research inclusion and equity and \$2M in responsible AI research.

3. Social media platforms have been used to suppress the ability of many users to exercise their civil rights both online and offline. What systems have been implemented on your platforms to prevent this from happening? Do you have any data as to how successful these systems are?

Google is deeply committed to civil rights. Our products are built for everyone, and we design them with extraordinary care to be a trustworthy source of information without regard to a user's demographic, socioeconomic background, or political viewpoint. Billions of people use our products to find information, and we help our users, of every background and belief, find the high-quality information they need to better understand the topics they care about. While our business is quite different from that of social media platforms, we have long been focused on making improvements to our products to protect civil rights.

We are engaged in extensive discussions with civil rights experts and leadership, and we are proud that we have civil and human rights expertise on staff, internal frameworks like our AI Principles and YouTube Community Guidelines in place, and governance structures through groups like our Responsible Innovation and Trust & Safety teams, working to help build civil and human rights considerations into our work. Our civil and human rights leads will continue to develop a structure to provide the transparency that the civil rights community needs, and we have confidence that we can demonstrate our long-term commitment to getting this right.

We aim to be as clear as possible to our users about how our products and policies work. We were the first platform to have a publicly-available transparency report in 2010. Since then, we have launched a number of different transparency reports to shed light on how the policies and actions of governments and corporations affect privacy, security, and access to information for our users. Our current reports cover topics such as security and privacy, content removal, political advertising on Google, and traffic and disruptions to Google. We also have a report specifically focused on YouTube community guidelines enforcement, including data on removal by the numbers, source of first detection, views, removal reason, and country/region. For example, please see our YouTube Community Guidelines Enforcement FAQs, <https://support.google.com/transparencyreport/answer/9209072>.

We will continue to approach this thoughtfully and to provide transparency about our products and policies, and we are always open to feedback on our efforts.

The Honorable Anna G. Eshoo (D-CA)

1. As I stated at the hearing, journalists and researchers suggest that YouTube has a problem with its recommendation engine. I acknowledge YouTube has done work to begin to resolve some of the issues, but how do you respond to the Anti-Defamation League's (ADL) report from last month that states that when participants watch videos from extremist channels, they are more likely to see and follow recommendations to similar videos?

a. Do you dispute the findings in the ADL report?

b. Will you rectify the issue at a product level (e.g., changing the recommendation algorithms)?

c. Have you taken actions against the channels ADL finds to be problematic? Why or why not?

As an initial matter, the ADL report concluded that there is no clear evidence that people with neutral or mixed views on issues such as race frequently view videos from alternative or extremist channels on YouTube. In fact, a recent study by Cornell University (please see <https://arxiv.org/abs/2011.12843>) similarly found no evidence that “echo chambers” are caused by YouTube recommendations. Rather, they found that the consumption of radical content on YouTube appears to reflect broader patterns of news consumption across the web.

That said, and while we encourage our creators to have a voice on YouTube, we have always had rules of the road. We take our responsibility to protect our users seriously. We have never allowed pornography, incitement to violence, or content that would harm children, and we strictly prohibit harassment, terrorist content, and hate speech. We know the vast majority of creators follow those rules. But when they don’t, we use a combination of human reviewers and cutting-edge machine learning to enforce our policies at scale, taking down half of extremist content in YouTube within two hours, and nearly 70% in eight hours. In Q4 2020, we removed 9.3M videos for violating our community guidelines and approximately 74% of those first flagged by our systems received 10 or fewer views.

In addition to removing content that violates our policies, we have systems in place designed to reduce recommendations of borderline content. YouTube defines borderline content as content that comes close to violating our policies but doesn’t cross the line. Since January 2019, when we announced the latest of our improvements to our recommendation systems, we have greatly reduced recommendations of borderline content and content that could misinform users in harmful ways. Specifically, we saw a 70% average drop in watch time of this content coming from non-subscribed recommendations in the U.S that year. And 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, 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.

To measure our effectiveness, and to create transparency around our enforcement, YouTube publishes quarterly data in our Community Guidelines enforcement report (<https://transparencyreport.google.com/youtube-policy/removals?hl=en>), a section of the Google Transparency Report. The YouTube report provides public data about the number of videos we remove from our platform for each of our policy violation categories, as well as additional information about channel removals, views before removals, appeals and reinstatements, and human and machine flagging. The report also has a “featured policy” page that is specific to our work to combat violent extremism.

Because this content is ever-evolving, we will continue to be vigilant in evaluating our responses to removing harmful content from our platforms and will continue to ensure that our recommendations systems raise up authoritative content versus borderline information.

2. An oft-cited figure states that 70% of viewing time on YouTube is attributable to recommendations. Yet this figure is now over three years old. What percentage of YouTube views and of total viewing time is attributable to 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. Our recommendation systems take into account many signals. We 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.

In January 2019 we announced that the latest of our improvements to our recommendation system is to greatly reduce borderline content and content that could misinform users in harmful ways. This is content which comes close to but doesn't quite violate our policies and represents less than 1% of the content that is watched on YouTube. The result was a 70% average drop in watch time of this content coming from non-subscribed recommendations in the U.S that year. While algorithmic changes take time to ramp up and you might see consumption of borderline content go up and down, our goal is to have views of non-subscribed, recommended borderline content below 0.5%. We have now rolled out

changes to our recommendation system to reduce borderline content in every market where we operate.

We raise up authoritative channels for newsworthy events in the “Up next” recommendation 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.

3. If a user is banned from YouTube for posting misinformation, are they banned from posting the same content to other Google services (e.g., podcasts in the Google Play store)? Why or why not?

In the event a channel or account has been terminated, an individual is prohibited from creating or owning other YouTube channels and/or accounts, and there are processes in place to enforce this policy. Note that it is also a violation of our policies and penalties to start a new channel to circumvent suspensions or terminations. To measure our effectiveness, and to create transparency around our enforcement, YouTube publishes quarterly data in our Community Guidelines enforcement report (<https://transparencyreport.google.com/youtube-policy/removals?hl=en>), a section of the Google Transparency Report. The YouTube report provides public data about the number of videos we remove from our platform for each of our policy violation categories, as well as additional information about channel removals, views before removals, appeals and reinstatements, and human and machine flagging.

We design the rules of the road across all of our products and services to protect users from harm while supporting the purpose of the product. For each product and service, we tailor these policies to strike the appropriate balance between providing access to a diversity of voices and limiting harmful content and behaviors. This balance can differ from one product to the next, in part because harm manifests differently in each service and context. While a universally recognized harm may be prohibited across all our products and services, it can appear on each product and service differently. So, we evaluate the potential for harm specific to each product and design our policies accordingly.

4. Why is it that a channel on YouTube TV can air content that would be removed on YouTube?

YouTube TV is not available for free to all users, like the main web-based YouTube platform, and our approach to the content appearing on YouTube TV also differs from the controls we have in place for the main YouTube platform. YouTube TV is a paid membership that brings users live TV from major broadcast networks, popular cable networks, and premium networks, along with popular shows from YouTube creators. YouTube TV will show related and trending content from YouTube on certain pages and search results. Users can watch these YouTube videos directly on YouTube TV. YouTube TV allows users to filter content to TV-Y, TV-Y& & TV-G shows, as well as G and PG films.

We take our responsibility to protect our users seriously, which is why we have policies prohibiting hate speech, terrorist content, and other content that violates our policies, as well as strict standards for who can monetize their content on YouTube.

5. The Center for Countering Digital Hate recently published a report titled “The Disinformation Dozen” that identifies 12 individuals responsible for 65% of all anti-vaccine content on social media. For each of the 12 individuals, please identify the following:

a. Have you taken adverse actions against any content the individual’s account(s) posted? If so, how many times?

b. Have you banned or taken other adverse actions against the individual’s account(s)?

Since the outbreak of COVID-19, teams across Google have launched 200 new products, features, and initiatives and are contributing over \$1B in resources to help our users, clients, and partners through this unprecedented time. That includes our Homepage “Do the Five” promotion, launch of the COVID-19 site, and amplifying authoritative voices through ad grants. As the COVID-19 situation has evolved, we have partnered closely with the World Health Organization and local health authorities to ensure that our policy enforcement is effective in preventing the spread of harmful misinformation relating to COVID-19.

As a result of this work, we have a clear COVID-19 medical misinformation policy for YouTube (<https://support.google.com/youtube/answer/9891785?hl=en>) that prohibits content which contradicts local health authorities’ guidance on the treatment, prevention, diagnosis, and transmission of COVID-19. Our policies prohibit, for example, content that explicitly disputes the efficacy of WHO or local health authority advice regarding social distancing that may lead people to act against that guidance.

We also prohibit claims that there are guaranteed prevention methods for COVID-19 or that an approved vaccine will cause death or contraction of other infectious diseases. YouTube’s policies specifically prohibit claims about COVID vaccines that contradict expert consensus from the World Health Organization or local health authorities.

We have consistently enforced these policies. For example, we’ve removed more than 850,000 YouTube videos with dangerous or misleading COVID-19 information, and in 2020, we blocked or removed over 275M COVID-related ads (globally), including Shopping ads, for policy violations including price-gouging, capitalizing on global medical supply shortages, making misleading claims about cures, and, most recently, fake vaccine doses. And there have also been more than 400B impressions on our information panels for COVID-19 related videos and searches.

To measure our effectiveness, and to create transparency around our enforcement, YouTube publishes quarterly data in our Community Guidelines enforcement report (<https://>

transparencyreport.google.com/youtube-policy/removals?hl=en), a section of the Google Transparency Report. The YouTube report provides public data about the number of videos we remove from our platform for each of our policy violation categories, as well as additional information about channel removals, views before removals, appeals and reinstatements, and human and machine flagging.

Our Ads policies are designed not only to abide by laws but to ensure a safe and positive experience for our users. Our Ads services prohibit sites and pages from deceiving users about their identity, products, or services, and prohibit posting of dangerous or derogatory content. Last year, we blocked and removed 3.1B bad ads—almost 5,900 bad ads per minute.

This includes policies that prohibit ads that potentially profit from or exploit a sensitive event with significant social, cultural, or political impact, such as a public health emergency. Videos that promote harmful health or medical claims, for example, are a violation of our long standing Harmful or Dangerous acts policy, and are not permitted to show ads. On July 17, 2020, we put additional safeguards in place by expanding our harmful health claims policies for both publishers and advertisers to include content about a current, major health crisis that contradicts authoritative, scientific consensus (e.g., claims that the COVID-19 vaccine includes a microchip).

6. Have you observed a rise in hate directed at the Asian American and Pacific Islander (AAPI) on YouTube in the prior year relative to the same 12-month period a year earlier?

a. If so, please share any quantifications of this change.

b. What have you done to combat anti-AAPI hate on YouTube?

We strongly condemn racism, sexism, and violence in all its forms, and we stand with our Asian and Pacific Islander communities in the fight against hatred. Our Community Guidelines set the rules of the road for what content is and is not allowed on YouTube, and we review these guidelines to adapt to the evolving nature of potentially violative content. This includes policies that prohibit content that praises, incites, or promotes violence. As noted in our response to Rep. Schakowsky's Question No. 5, there are also other policies that are helpful in combating evolving threats and providing as much coverage as possible. These include:

- **Harassment & Cyberbullying:** YouTube prohibits content that threatens individuals. We also don't allow content that targets an individual with prolonged or malicious insults based on intrinsic attributes. These attributes include their protected group status, which include 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. For example, we don't allow content in which someone simulates violence toward an identifiable individual, such as by using a weapon on a dummy representing that person.

- **Harmful or Dangerous Content:** YouTube prohibits content that encourages dangerous or illegal activities that risk serious physical harm or death. This policy specifically prohibits, for example, content that provides instructions to build a bomb meant to injure or kill others. This policy also prohibits promoting or glorifying violent tragedies, such as school shootings.
- **Hate Speech:** We remove 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.

To measure our effectiveness, and to create transparency around our enforcement, YouTube publishes quarterly data in our Community Guidelines enforcement report (<https://transparencyreport.google.com/youtube-policy/removals?hl=en>), a section of the Google Transparency Report. The YouTube report provides public data about the number of videos we remove from our platform for each of our policy violation categories, as well as additional information about channel removals, views before removals, appeals and reinstatements, and human and machine flagging. The report also has a “featured policy” page that is specific to our work to combat violent extremism.

7. What disinformation campaigns did you track in the lead up to the November 2020 elections that targeted people of color, women, LGBTQ+, veteran, or older American users?

a. How did you respond to these campaigns?

b. Did you notify the users they had been targeted and share corrective information?

Because the answers to these questions are related, we have grouped together our response to Questions Nos. 7 and 9.

Elections are a critical part of the democratic process, and we are committed to helping all voters find relevant, helpful, and accurate information. Google takes election integrity very seriously, and we have a number of policies to protect our users from disinformation related to elections. For example, we have zero tolerance for ads that employ voter suppression tactics or undermine participation in elections. When we find those ads, we take them down. Moreover, given that an unprecedented amount of votes were counted after this past election

day, we implemented a sensitive event policy for political ads after polls closed on November 3rd. For more information, please see https://support.google.com/adspolicy/answer/6015406?hl=en&ref_topic=1626336.

Our misrepresentation policy for advertising (https://support.google.com/adspolicy/answer/6020955?hl=en&ref_topic=1626336) also prohibits claims related to politics, social issues, or matters of public importance that are “demonstrably false and could significantly undermine participation or trust in an electoral or democratic process.” Examples include false information about public voting processes and procedures, political candidate eligibility based on age or birthplace, election results, or incorrect claims that a public figure has died, or been involved in an accident.

We also provide transparency regarding election ad spending. Our publicly accessible, searchable, and downloadable Transparency Report contains information about election ad content and spending on our platforms (<https://transparencyreport.google.com/political-ads/region/US>). The report provides information about when election ads ran, how they were targeted, how many impressions they served, and the advertiser who paid for the ads.

We undertook a wide range of approaches to protect election integrity and prevent platform abuse during the 2020 election. They were focused on four different areas:

1. Elevating authoritative election-related content: When people sought information on elections, polling places, or candidates, Google protected against abuse and manipulation by safeguarding products like WebAnswers and Knowledge Panels from being used to mislead voters. Collectively, our election info panels were shown more than 8B times. In addition to the election results information panel, we displayed a number of other election-related information panels leading up to and throughout the course of the 2020 U.S. election, including on where and how to vote. With more than a dozen U.S. publishers participating, the information we presented was from independent third parties from across the political spectrum — such as The Dispatch, FactCheck.org, PolitiFact, and The Washington Post Fact Checker.
2. Combating coordinated influence operations: Our Threat Analysis Group and Trust & Safety teams monitored and disrupted account hijackings, inauthentic activity, disinformation campaigns, coordinated attacks, and other forms of abuse on our platforms 24/7. For more information, please see our Threat Analysis Group’s blog, <https://blog.google/threat-analysis-group/>.
3. Protecting users and campaigns: We trained nearly 4,000 civic entities and more than 4,000 election-related stakeholders on our security tools to improve election integrity. We created Protect Your Election, a suite of free tools to help protect high-risk users from potential digital attacks. We launched the Advanced Protection Program (APP)

(<https://landing.google.com/advancedprotection/>), providing top level account protection for campaigns, and we worked with the non-profit organization Defending Digital Campaigns (<https://www.defendcampaigns.org/>) to provide free Titan Security keys. For more information about our efforts to promote election and civic integrity, please see our recent blog posts, including <https://www.blog.google/technology/ads/update-our-political-ads-policy/>, and <https://blog.google/outreach-initiatives/civics/following-2020-us-election-google/>.

4. Continuing our work with law enforcement and industry partners on identifying and combating coordinated influence operations.

With respect to corrective information, on YouTube, our goal is to provide context and authoritative information before or during viewer engagement, not after. Our work in this area is guided by four strategic pillars: removing violative content, raising up authoritative content, reducing the spread of borderline content, and rewarding trusted creators—what are referred to as the four “Rs” of responsibility. In addition to these strategic initiatives, we utilize a number of other product features, such as information panels and fact checks to provide context and links to authoritative sources. We also have invested heavily to make sure that we surface authoritative content in our search results, which significantly reduces the spread of disinformation.

We also actively work to provide users with more information about the content they are seeing to allow them to make educated choices. On YouTube, Our Community Guidelines (<https://www.youtube.com/howyoutubeworks/policies/community-guidelines/>) prohibit spam, scams, deceptive practices, and content that seeks to incite violence. We also work to make sure that the line between what is removed and what is allowed is drawn in the right place. For example, our policies prohibit misleading viewers about where and how to vote. For more information on YouTube’s efforts to protect the integrity of the 2020 elections, please see our YouTube 2020 Election Retrospective at <https://www.youtube.com/howyoutubeworks/our-commitments/supporting-political-integrity/#election-news-and-information>.

We are proud of these processes that help protect against abuse and manipulation across our products and that help ensure the integrity and transparency of our nation’s elections.

8. What categories of information related to the January 6th attacks have you shared with law enforcement agencies?

a. Have you shared any of this information with researchers?

b. Have you permanently deleted any of this information?

c. Will you commit to retaining any such information until all legitimate law enforcement and research requests are fulfilled?

The attack on the Capitol on January 6 was unprecedented, and any attack on our democracy or incitement to violence is something we have always worked hard to prevent. Protecting our users is always our first priority. While we generally apply our policies independently, we of course rely on information from our peers and law enforcement to ensure the integrity of our platform. From Counter Terrorism, to our work to fight against CSAM, to protecting our elections, we work closely with law enforcement. Just recently, we were thanked by law enforcement for the work we did to protect the 2020 elections.

We have a robust law enforcement response process with analysts and lawyers dedicated to ensuring that we appropriately respond to legal process from law enforcement and make referrals to law enforcement when we identify problematic or illegal activity on our platform. A law enforcement agency may ask Google to preserve specific information while the agency applies for legal process to compel the disclosure of that information. We have a dedicated team that responds to law enforcement around the clock, every day of the year, and we respond to such requests in accordance with applicable law, including requests that may be related to the January 6 events. Legal requests related to ongoing investigations are often subject to non-disclosure requirements. As such, we are generally unable to comment further on specific matters.

We are proud of the work we do with law enforcement to protect our democracy and our users, and we describe our work with law enforcement in our publicly available Transparency Report. For more information, please see <https://transparencyreport.google.com/>.

9. What did you do to combat voter suppression targeted at Black Americans in the recent elections?

Please see the response to Question 7 above.

10. If a user's post is designed to meet the letter of your policies but clearly infringes on the spirit of your misinformation policies, how do you treat the content and the content creator?

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. 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've been investing in the policies, resources, and products needed to protect our users from harmful content. We have clear and public policies that we apply consistently. We've dedicated significant resources to content moderation efforts: we spent nearly \$1.2B on content moderation efforts in 2020. We also have nearly twenty-two 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, for instance by reducing recommendations of borderline content on YouTube. This is content which comes close to but doesn't quite violate our policies, and represents less than 1% of the content that is watched on YouTube. In January 2019, we announced that we would begin reducing recommendations on YouTube of borderline content or videos that could misinform users in harmful ways. 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 (<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. On YouTube, for example, there have been billions of impressions of information panels around the world since June 2018. For more information, 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.

11. Spanish, Cantonese, Mandarin, Tagalog, Vietnamese, French and French Creole, Korean, German, Arabic, and Russian are the most spoken languages in the U.S. What are you doing to combat misinformation in these languages?

As noted in the response to Rep. Schakowsky's Question No. 3, we have clear, public-facing policies that outline what is allowed on our services, and what is not. Our policies are not language-specific, and we do not differentiate when it comes to content on our platforms. We apply our policies equally and protect our users no matter where they are—in the United States

or in a foreign country. We use both automated and manual detection and enforcement efforts to protect our users.

Our business model depends on us being a useful and trustworthy source of information for everyone. We will continue to be vigilant to prevent deceptive content, no matter what language it is in.

12. How often are you personally involved in content moderation decisions?

To enforce our content moderation policies at the scale of the web, we rely on a mix of automated and human efforts to spot problematic content. We also enable users and trusted organizations to flag content that may be problematic. Flagged content is reviewed by our Trust & Safety teams who, along with the relevant product team (e.g., YouTube, Google Play, etc.) make decisions as to whether the content warrants action. We've dedicated significant resources to these content moderation efforts: we spent nearly \$1.2B on content moderation efforts in 2020. We also have nearly twenty-two thousand employees dedicated to ensuring the protection of our platforms.

13. How many employees and contractors at your company are dedicated to content moderation?

a. How does this number compare to one year and four years prior?

b. To what degree are your content moderation team members

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. To enforce our policies at scale, we use a combination of automated and human efforts to spot problematic content.

We've dedicated significant resources to content moderation efforts: we spent nearly \$1.2B on content moderation efforts in 2020. We also have nearly twenty-two thousand employees dedicated to ensuring the protection of our platforms. Content moderation at Google is primarily managed by Trust & 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. These teams also work with our legal and public policy teams, and oversee the vendors we hire to help us scale our content moderation efforts, as well as provide the native language expertise and the 24-hour coverage required of a global platform. Moderating content at scale is an immense challenge, but we see this as one of our core responsibilities and we are focused on continuously working toward removing content that violates our policies before it is widely viewed.

14. How many employees and contractors does your company employ for selling advertising?

As of December 31, 2020, Alphabet had 135,301 employees. The majority of new hires during last year were engineers and product managers. As we recently announced, we plan to invest over \$7B in offices and data centers across the U.S. and create at least 10,000 new full-time Google jobs in the U.S. this year. This includes investing in communities that are new to Google and expanding in others across 19 states.

15. To what degree are your product teams and executives incented – through bonuses, salary changes, or non-monetary incentives – to increase user engagement? Please describe any such incentives.

Our people are critical for our continued success. We work hard to provide an environment where Googlers can have fulfilling careers and be happy, healthy and productive. We offer industry-leading benefits and programs to take care of the diverse needs of our employees and their families, including access to excellent healthcare choices, opportunities for career growth and development, and resources to support their financial health. Our competitive compensation programs help us to attract top candidates and retain employees, and reward talented technical and non-technical Googlers well for the short- and long-term success of our business.

16. Have you studied the effectiveness of targeted or behavioral advertising relative to contextual advertising for each of Google’s advertising products?

a. What is the difference between targeted or behavioral advertising and contextual advertising with respect to user engagement and time spent on your platform?

b. Have you partnered with independent researchers to study this? If so, what did they find?

Users have control over whether and how we use their data to make the ads they see more relevant. They can opt out of interest-based ads at any time, see why each ad was selected, turn off ads personalization, mute ads on third-party sites, and delete data stored in their account. We never sell our users’ personal information, and our policies prohibit ad targeting to individuals based on sensitive categories like religion and personal health conditions.

We are also charting a course for a privacy-first web. Advertising is the economic foundation of the internet. But as our industry has strived to deliver relevant ads to consumers across the web, it has created a proliferation of individual user cross-site data across thousands of companies, typically gathered through third-party cookies.

At Google, we realize that digital advertising must evolve to address the growing concerns people have about targeted ads. That’s why in 2020, we announced our intent to remove

support for third-party cookies. We have approached this issue very thoughtfully and have sought to avoid unintended consequences to users and the web ecosystem. Specifically, we wanted to ensure our continued ability to support publishers—especially during COVID—and to prevent invasive workarounds, such as fingerprinting, which actually reduce user privacy and control. That’s why we’ve been working with the broader industry on the Privacy Sandbox to build innovations that protect anonymity while still delivering results for advertisers and publishers. This year, we made clear that once third-party cookies are phased out, we will not build alternate identifiers to track individuals as they browse across the web, nor will we use them in our products.

Advances in aggregation, anonymization, on-device processing and other privacy-preserving technologies offer a clear path to replacing individual identifiers. In fact, our latest tests of Federated Learning of Cohorts (“FLoC”) show one way to effectively take third-party cookies out of the advertising equation and instead hide individuals within large crowds of people with common interests. For more information on FLoC, please see <https://blog.google/products/ads-commerce/2021-01-privacy-sandbox/>. Chrome made FLoC-based cohorts available for public testing through origin trials with its release in March, and we expect to begin testing FLoC-based cohorts with advertisers in Google Ads this year. As more proposals reach the origin trial stage, they will receive more feedback from end users and the industry.

We remain committed to preserving a vibrant and open ecosystem where people can access a broad range of ad-supported content with confidence that their privacy and choices are respected. We look forward to working with others in the industry on the path forward.

17. As you have in the past, will you commit to continue promoting authoritative information on Google Search users about Open Enrollment?

Our business model depends on us being a useful and trustworthy source of information for everyone. As detailed in our answers to Questions Nos. 7 and 10, we have a number of tools that curb harmful disinformation and promote authoritative content in our products—from ranking algorithms to policies against selling ads with deceptive content. As always, we are committed to promoting authoritative content about Open Enrollment to our users through Google Search.

The Honorable Yvette D. Clarke (D-NY)

1. Mr. Pichai, what is your company doing to address the amplification and discrimination in ad-targeting?

We are committed to ensuring that our policies protect users from content that incites hatred against, promotes discrimination of, or disparages individuals or groups based on

characteristics associated with systemic discrimination or marginalization. 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.

Our Ads policies are designed not only to abide by applicable laws but also to ensure a safe and positive experience for our users. Our Ads services prohibit sites and pages from deceiving users about their identity, products, or services, and prohibit posting of dangerous or derogatory content. Last year, we blocked and removed 3.1B bad ads — almost 5,900 bad ads per minute. We also give users control over the kinds of ads they see, including the ability to opt out of seeing any personalized ads (<https://adssettings.google.com/authenticated>). Our Ads policies apply to all of the ads we serve. If we find ads that violate our policies, we take action.

For over a decade, we've also had personalized advertising policies (<https://support.google.com/adspolicy/answer/143465?hl=en>) that prohibit advertisers from targeting users on the basis of sensitive categories related to their identity, beliefs, sexuality, or personal hardships. This means we don't allow advertisers to target ads based on categories such as race, religion, ethnicity, or sexual orientation, national origin, or disability to name a few. We regularly evaluate and evolve our policies to ensure they are protecting users from behaviors like unlawful discrimination.

To further improve access to housing, employment, and credit opportunities, as part of a collaboration with the U.S. Department of Housing and Urban Development, in June 2020 we introduced a new personalized advertising policy for certain types of ads. This policy prohibits impacted employment, housing, and credit advertisers from targeting or excluding ads based on gender, age, parental status, marital status, or ZIP code, in addition to our longstanding policies prohibiting personalization based on sensitive categories like those described above.

Google is committed to working with the broader advertising ecosystem to help set high standards for online advertising, and we will continue to strive to set policies that improve inclusion and access for users. For more information, please see <https://www.blog.google/technology/ads/upcoming-update-housing-employment-and-credit-advertising-policies/> and <https://support.google.com/adspolicy/answer/143465?hl=en>.

The Honorable Robin Kelly (D-IL)

1. At the hearing you said that you do not profit off harmful misinformation, conspiracy theories, and violent content.

a. How do you manage to avoid collecting revenue from ads served on misinformation content?

b. Are you claiming that your company has not received any payments from ads that displayed before or adjacent to harmful content and was later removed for violating your policies?

Our business model depends on us being a useful and trustworthy source of information for everyone. We have a number of tools that curb harmful disinformation and promote authoritative content in our products—from ranking algorithms to policies against selling ads with deceptive content. We have strong policies across our products to protect our users, and we use both automated and manual detection and enforcement efforts. 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.

Our policies are designed not only to abide by laws but to ensure a safe and positive experience for our users. For example, our Publisher Policies (<https://support.google.com/adsense/answer/9335564>) prevent ads from running alongside certain types of problematic content, including alongside content that is illegal, dangerous or derogatory, or misrepresentative. And we're serious about enforcing these policies. In 2020 alone, we took action against 1.3B pages and 1.6M domains. This included taking down 168M URLs due to dangerous or derogatory content. For more information, please see our 2020 Ads Safety Report, available at https://services.google.com/fh/files/misc/ads_safety_report_2020.pdf.

In addition to our policies and enforcement, Google provides advertisers with robust controls over what kind of content their ads show on. We allow them to exclude specific websites (https://support.google.com/google-ads/answer/2454012?hl=en&ref_topic=3123061) as well as entire topics that they wish to avoid (<https://support.google.com/google-ads/answer/2456096>). For additional information, please see <https://support.google.com/google-ads/answer/3306596>. We recently improved these tools to allow advertisers to upload dynamic exclusion lists as part of their advertiser controls. While advertisers could already upload exclusion lists to their Google Ads account, this new tool now allows these lists to be scheduled for updates as new websites or web pages are added by the advertisers themselves or by a third party they trust (<https://support.google.com/google-ads/answer/10602822>).

The Honorable Kathleen M. Rice (D-NY)

1. There have been some successful partnerships between big tech companies like yours and law enforcement. However, the sales of illicit live animals and their parts and products continue to grow in online marketplaces and communities of violence promoting illegal animal abuse content continue to be prevalent problems. Some online crime experts have raised concerns that the widespread company policies of removing illicit wildlife for sale may be hindering law enforcement efforts by prematurely alerting

criminals who then open new accounts under different usernames. It also poses the problem of deleting public evidence that could help authorities catch traffickers. Company policies also differ between platforms, creating an inconsistent landscape for both companies and law enforcement to work in.

How could Congressional mandates assist your company in building stronger and more comprehensive partnerships with law enforcement?

We have robust policies addressing animal cruelty, such as ads and merchant policies that prohibit content that promotes cruelty or gratuitous violence towards animals, or content that may be interpreted as trading in, or selling products derived from, threatened or extinct species (please see <https://support.google.com/adspolicy/answer/6015406> and <https://support.google.com/merchants/answer/6149990>).

But as noted in our response to Rep. Eshoo's Question No. 8, we also have a robust law enforcement response process with analysts and lawyers dedicated to ensure that we both appropriately respond to legal process from law enforcement to protect our users, and make referrals to law enforcement when we identify problematic or illegal activity on our platform. A law enforcement agency may ask Google to preserve specific information while the agency applies for legal process to compel the disclosure of that information. We respond to such requests in accordance with applicable laws.

As our Transparency Report shows, we work cooperatively with law enforcement, while respecting the privacy of our users. For more information, please see <https://transparencyreport.google.com/>.

The Honorable Gus M. Bilirakis (R-FL)

The tech industry can often be a metrics-driven and data-driven world. Frequently, you're trying to improve metrics that measure engagement: video watch time for YouTube, the number of daily active users on Facebook, or the number of daily tweets on Twitter.

However, sometimes you can increase engagement by pushing both sides into their echo chambers. This echo chamber effect leads to less civil discourse and more political polarization, and it also allows misinformation to more easily to spread on both sides.

1. With respect to the problem of echo chambers, what metrics has your company developed which can measure this problem?

We are committed to making Google a safe and inclusive space for people to share their viewpoints. We understand your concerns and are deeply troubled by any attempts to use our

platforms to sow division. As noted in our response to Rep. Eshoo’s Question No. 1, a recent report by the Anti-Defamation League concluded that there is no clear evidence that people with neutral or mixed views on issues such as race frequently view videos from alternative or extremist channels on YouTube. In fact, a recent study by Cornell University (<https://arxiv.org/abs/2011.12843>) similarly found no evidence that “echo chambers” are caused by YouTube recommendations. Rather, they found that the consumption of radical content on YouTube appears to reflect broader patterns of news consumption across the web.

Nonetheless, we have put significant effort into combating harmful activity across our platforms. This includes, for instance, ranking algorithms in Search that prioritize authoritative sources. Our Search algorithm ranks pages to provide the most useful and relevant information by matching search terms against available web pages and looking at factors like the number of times the words appear and the freshness of the page. A user’s viewpoint is not a relevant factor. We also seek to ensure that our Search results are providing the most authoritative and responsive results by using external quality raters from across the vast majority of U.S. states. In addition, we have robust systems in place to ensure that our policies are enforced in a politically impartial way across all of our products and services, including Search.

Additionally, on Google News, we mark up links with labels that help users understand what they are about to read—whether it is local content, an op-ed, or an in-depth piece, and encourage them to be thoughtful about the content they view. Publishers who review third-party claims or rumors can showcase their work on Google News and in Google Search through fact-check labels. People come across these fact checks billions of times per year. For more information, please see <https://blog.google/products/search/fact-check-now-available-google-search-and-news-around-world/>.

Google News also has a Full Coverage feature that provides a complete picture of how a story is reported from a variety of sources. Through this feature, you can see top headlines from different sources, videos, local news reports, FAQs, social commentary, and a timeline for stories that have played out over time. Having a productive conversation or debate requires everyone to have access to similar information. That’s why content in Full Coverage is the same for everyone—it’s an unpersonalized view of events from a range of trusted news sources. Meanwhile, to find out what the world is reading, Headlines provides an unfiltered view of news from around the world. Additional sections let you dig into more on technology, business, sports, entertainment, and other topics.

We also have increased transparency around news sources on YouTube, including disclosure of government funding. When a news channel on YouTube receives government funding, we make that fact clear by including an information panel under each of that channel’s videos. There have been billions of impressions of information panels on YouTube around the world since June 2018. For more information, please see <https://support.google.com/>

[youtube/answer/7630512](https://www.youtube.com/watch?v=7630512), and <https://blog.youtube/news-and-events/greater-transparency-for-users-around>. Our goal is to equip users with additional information to help them better understand the sources of news content that they choose to watch on YouTube. For well-established historical and scientific topics that are often subject to misinformation, such as “Apollo 11,” general reference articles linking to third-party sources appear alongside related search results and videos to provide more context. We do this by surfacing contextual information from third-party sources including Encyclopedia Britannica and Wikipedia alongside videos and search results. Likewise, for searches related to COVID-19 or COVID-19 vaccine info, you may see information panels on YouTube featuring links to guide you to the World Health Organization for more information. For more information, please see <https://www.youtube.com/howyoutubeworks/product-features/news-information/#topics-prone-to-misinformation>.

We also have taken a tougher stance on removing hateful and supremacist content and have reduced borderline content by reducing recommendations of content that comes close to violating our guidelines. We are proud that we have a wide variety of views expressed on our platforms and are committed to ensuring we continue to enforce our policies in a fair and impartial manner.

Mr. Pichai, in the book *How Google Works*, Eric Schmidt and Jonathan Rosenberg wrote, “You cannot be gender-, race-, or color-blind by fiat; you need to create empirical, objective methods to measure people.”

2. When you testified before Congress in December 2018, you said, “I lead this company without political bias and work to ensure that our products continue to operate that way.” Put another way, you declared by fiat that Google does not have political bias. But let’s apply your own rules here. You cannot be politically blind by fiat, so how can you empirically prove that Google does not have political bias?

Mr. Pichai, when you testified before Congress in December 2018, you said, “We don’t manually intervene on any particular search result.” But your sworn testimony later proved to be false. In April 2019, the Daily Caller reported that Google does have a manual blacklist which is curated by the Trust & Safety team, and this report was also corroborated by a Wall Street Journal investigation later that year. One document describing this blacklist was approved by Ben Gomes, who reported directly to you at the time.

This manual blacklist doesn’t have any effect on the ten blue links, but it will block a site from almost all search features: knowledge boxes, web answers, side bars, etc. However, according to a study released last month, 65% of Google searches in 2020 were zero-click searches. Users did not get their answers by clicking on one of the ten blue links; they got their answer directly from one of these search features.

Our mission and business model is dependent on being a useful and trustworthy source of information for everyone, so we have a natural, long-term incentive to prevent anyone from interfering with the integrity of our products. We have stood for freedom of expression since our founding—accordingly, our services empower users to access a wide variety of viewpoints and perspectives. Our services also enable voices across the political spectrum to reach broader audiences. Some of our biggest critics on the right and left have had unprecedented success on our platforms.

We have over 100,000 employees around the country and the world, representing a diverse swath of backgrounds and viewpoints. For example, we have over 1M message-board groups at Google, representing all political backgrounds: we have a republicans@ discussion board, in addition to liberals@, conservatives@, etc. We've been very clear with our employees that we will not allow any form of political bias whatsoever in our products. We have robust controls to prevent employee interference with our products, but if we ever thought that any Googler would attempt to manipulate products in order to achieve their personal political objectives, we'd take very strong action.

Reputable third parties such as The Economist (<https://www.foxnews.com/tech/google-rewards-reputable-reporting-over-left-wing-politics-economist-study-says>) and New York University (<https://bhr.stern.nyu.edu/bias-report-release-page>) have conducted empirical investigations of our search results, with transparent methodologies, and found no evidence of partisan leaning.

Our commitment to our users has and will continue to be consistent: we'll be a useful and trustworthy source of information for everyone, no matter who they are, and no matter what their political beliefs may be.

3. If a site does end up on this blacklist, are they effectively locked out from these 65% of searches that are zero-click searches?

4. According to the Daily Caller's reporting, this blacklist included "Gateway Pundit, Matt Walsh's blog, Gary North's blog 'teapartyeconomist.com,' Caroline Glick's website, Conservative Tribune, a property of The Western Journal and the website of the American Spectator." Which sites are still on that blacklist?

5. For each of these aforementioned sites, what is the reason why they were originally added to this blacklist?

6. If a site is currently on this blacklist, or if a site is added to this blacklist in the future, can you commit to notifying that site that they have been blacklisted, and also explain how they violated your policies?

Because the answers to these questions are related, we have grouped together our response to Questions Nos. 3-6.

We do not use so-called “blacklists” to promote or demote any particular political ideology in Search—not liberal, not conservative, not any other. Public reports on so-called “blacklists” might refer to a number of measures we take to enforce our content policies and ensure that our algorithms are working as intended to surface high-quality information to users. But we do not manually intervene to set the ranking of any particular Search result or manually determine the ranking of our listings.

When we make changes they apply broadly. We remove content from Search results only in very limited circumstances, such as a court order, valid copyright takedown notice, or violation of our webspam policies. Our publicly available Search Quality Rater Guidelines (<https://static.googleusercontent.com/media/guidelines.raterhub.com/en//searchqualityevaluationguidelines.pdf>) provide transparency for ratings, users, and webmasters about how Search works. They make it clear that ratings should never be based on personal opinions, preferences, religious beliefs, or political views.

The Daily Caller story conflates Google’s organic Search results, commonly known as the “10 blue links,” with Google’s curated convenience features. Because these convenience features are more prominent than the 10 blue links, we have more stringent content policies for them. For example, Featured Snippets receive unique formatting and positioning on Google Search and are often spoken aloud by the Google Assistant. Because of this treatment, we do not allow Featured Snippets to contain sexually explicit vulgar language, whereas 10 blue link results may. If our algorithms fail to recognize that a particular Featured Snippet contains inappropriate vulgar language, we may manually block the Featured Snippet. This is simply enforcement of our publicly disclosed policy, available here <https://support.google.com/websearch/answer/9351707#policies>.

With respect to “zero click” searches (i.e., the claim that the majority of searches on Google end without someone clicking off to a website), this claim relies on flawed methodology and misunderstands how people use Search. There are a lot of reasons why not every query results in a click to a website. For example, people reformulate their queries or look for quick facts that are provided directly on the results page (drawing from licensing agreements or tools we’ve developed). When it comes to local businesses, Google provides many ways for consumers to connect directly with businesses through Google Search, many of which do not require a traditional click. Furthermore, some searches take users directly to apps instead of websites. Despite these possibilities, we send billions of visits to websites every day, and traffic sent into the open web has increased every year since Google Search was first created.

We care deeply about the open web and have continually improved Google Search over the years, helping businesses, publishers and creators thrive. Some would argue that we should revert back to showing only 10 blue website links. While we do show website links for many

queries today when they are the most helpful response, we also want to build new features that organize information in more helpful ways than just a list of links. And we've seen that as we've introduced more of these features over the last two decades, the traffic we're driving to the web has also grown—showing that this is helpful for both consumers and businesses.

The Honorable Kelly Armstrong (R-ND)

1. Mr. Pichai, in this hearing and a 2020 House Judiciary Committee hearing (Antitrust Subcommittee hearing entitled, “Online Platforms and Market Power, Part 6” on July 29, 2020), I asked several questions about the potential competition harms from several recent Google actions. The premise of each question acknowledged that the referenced actions are at least partly justified on privacy grounds; however, the question only inquired as to the competition issue. In both hearings, your answers reiterated the privacy justifications while largely disregarding the competition question. Please directly answer to the competition issue in the following question. Will the Federated Learning of Cohorts (FLoCs) proposal, regardless of potential privacy benefits, further entrench Google’s digital advertising market share and harm competition? You testified that Google will not have access to “FLoC data”. However, Google will still control the Chrome-ad ecosystem while third-party competitors are excluded, correct? Please explain whether the exclusion of third-party competitors from the Chrome ad exchange will harm competition.

FLoC is a new web technology, still in development, to enable advertisers to show relevant ads without tracking individuals across the web. FLoC is one component of the Privacy Sandbox, an open initiative to create web technologies that both protect people’s privacy online and give companies the tools to build thriving businesses that keep the web open and accessible to everyone.

FLoC will not give Google Ads special privileges or access to data. Google Ads will connect to FLoC in the same way as other advertising platforms, publishers, or advertisers. FLoC and other Privacy Sandbox innovations are open source and free to use, so other members of the web community are invited to contribute to these technologies and adopt them in their own platforms and browsers.

Chrome does not operate an ad ecosystem or ad exchange. Rather, as a web browser, Chrome enables people to visit websites, and many of those websites choose to display ads to their visitors as a means to support the content they publish. Google believes that this ad-supported web is at risk if digital advertising practices don't evolve to reflect people's changing expectations regarding the collection and use of data. That is why the Privacy Sandbox is focused on developing new privacy-centric technologies like FLoC as alternatives to legacy, data-intensive mechanisms like third-party cookies. These technologies can

improve people's privacy online while maintaining free access to information for everyone, so that the web can continue to support economic growth, now and for the future.

2. Mr. Pichai, Senator Roger Wicker submitted several questions for the record (QFR) pursuant to a Senate Commerce Committee hearing entitled, "Does Section 230's Sweeping Immunity Enable Big Tech Bad Behavior", which was held on October 28, 2020. A question from Senator Wicker asked, "[I]s Google actively blocking competing video players in favor of Chrome-specific products? Google's response did not directly respond to the question, but broadly discussed the integration of "media technologies developed by third-parties" in Chrome. I would appreciate answers to the following questions. Which non-Google video players are compatible with Chrome? Is ISO's MPEG High Efficiency Video Code compatible with Chrome? If not, why not? If ISO's MPEG High Efficiency Video Code is compatible with Chrome, do users have the ability to access it in Chrome? If not, why not? If "Google Chrome is focused on creating the best possible experience for web browsing", which is Google's stated goal in its aforementioned response to Senator Wicker, what active steps is Google taking to ensure that third-party products and services that directly compete with Google product offerings are easily available and accessible with full functionality to all consumers when using Chrome?

Google Chrome's goal is focused on creating the best possible experience for web browsing. To meet this end, Google has focused on including technologies in Chrome that facilitate the development and delivery of media content on Internet-connected devices and improve the media experience for all users across the browser ecosystem, including OEMs and content providers. These include media technologies developed by third parties outside Google.

The success of Chrome depends on providing users with a fast, secure, and performant browsing experience for websites and services across the web, whether or not they are from Google or other providers. Understanding that is why a number of users choose Chrome from among the many browser options available. We continually evaluate the need to include technologies in Chrome based on feedback from our entire ecosystem of users, OEMs, and content providers, and we add technologies where we identify shared needs across the ecosystem.

The Honorable Debbie Lesko (R-AZ)

Americans are rightfully frustrated with the role of social media and big tech companies in our everyday lives. These platforms play an important role in connecting Americans, but they also play an outsized role in the shaping of political discourse. Some of my colleagues have advocated for more content moderation, while many of my constituents

are concerned that more content moderation of political speech, stifles voices and harms our ability to have transparent debates about the issues that affect all of us.

In addition, big tech companies' business models are predicated on advertising that justifiably raises privacy concerns. It is clear advertising spending is more effective if it can be targeted, but this raises important questions about what information Americans are sharing and are consenting to sharing in the digital economy.

1. Mr. Pichai, in 2018, Google launched Project Strobe, a review of third-party developer access to Google account and Android device data. Google announced updates to its User Data Policy for consumer Gmail application programming interfaces (APIs) to limit the apps that may seek permission to access a Gmail user's data. Specifically, the policy allows only apps that directly enhance email functionality to be authorized to access this data. In addition, these apps will need to agree to new rules on handling Gmail data and will be subject to security assessments by Google-designated third parties.

2. What process does Google use to certify third-party security assessors?

3. You've stated that Google does not "look into users' email contents and we do not share the contents with anyone else without the user asking us to do so." However, these ads are shown to Gmail users based on online activity while signed into Google and is a fully automated process. Does this automated process start as soon as a Gmail user logs into their Gmail account?

4. Does the automated process end once the Gmail user explicitly signs out of their Gmail account? Specifically, at what point does the automated access/monitoring process end and is there any way you or the third-party security assessors continue to collect data after a user has signed out of their Gmail account?

5. Does this automated process require explicit, additional consent on the part of a Gmail user to opt into targeted ads?

Because the answers to these questions are related, we have grouped together our response to Questions Nos. 1-5.

Privacy is for everyone. We have a suite of tools that give users security, transparency, and control over their data. Our goal is to make things as simple as possible while ensuring that users have meaningful control over their data. We make that easy: over 1B users have reviewed their Google privacy settings. More than 20M people every day visit their Google Account page, where they can review their privacy settings and see and delete data stored with their account. And last year, we announced a new auto-deletion default for the Location History, Web & App Activity settings.

Concerning Gmail, in very specific cases where users ask us to and give consent, or where we need to for security purposes (such as investigating a bug or abuse), we may access users' emails. However, we do not use Gmail content to target ads. Before a developer can access a Gmail users' data through the Gmail API, they must obtain consent from the user.

We continuously work to vet developers and their apps that integrate with Gmail API before we allow the developer the ability to request access to user data. Our policies and terms require developers to accurately represent the identity of the application, provide clear and accurate information regarding the types of data being requested, and be honest and transparent with users about the purpose of user data requests. Developers must publish a privacy policy that fully documents how the application interacts with user data. If developers change the way their application uses a Google user's data, they must notify the user and prompt them to consent to an updated privacy policy. Developers are also required to protect against unauthorized or unlawful access, use, destruction, loss, alteration, or disclosure.

We have also implemented an improved user permission flow that provides a finer-grained ability to choose what data users share, limits the types of apps that can request access from Gmail users, and imposes new requirements on how developers must treat Gmail data. Some examples of such apps are those that provide services like helping users manage their travel itineraries or organize their inbox.

Anyone can use our services—including Gmail—for free, without seeing personalized ads from Google at all. We have always made privacy a touchstone of Google's relationship with our users.

6. Regarding the issue of content moderation, does your news and safety team maintain a list of websites that are restricted from appearing in any Search feature or news product?

We design Search and News with extraordinary care to serve our mission without political bias and have robust systems to ensure that employees' personal biases do not impact our products and that our policies are enforced in a politically neutral way. Our Search algorithm ranks pages to provide the most useful and relevant information by matching search terms against available web pages and looking at factors like the number of times the words appear and freshness of the page. For any search for news content, we show news-related results from across the internet, and this can include press releases — from government entities such as NASA, the U.S. Department of Justice, and from interest groups and companies, alike.

We assess sites for inclusion in Google News or Top Stories against a set of publicly disclosed Content Policies (https://support.google.com/news/publisher-center/answer/6204050?visit_id=637546080106902835-3077183141&rd=1). Sites that do not adhere to these policies are not eligible to appear on Google News surfaces or in our news carousels found at the top of search results. Our news content policies target bad behavior, regardless of political

perspective. Google News explicitly prohibits content that incites, promotes, or glorifies violence, harassment, or dangerous activities. Similarly, Google News does not allow sites or accounts that impersonate any person or organization, that misrepresent or conceal their ownership or primary purpose, or that engage in coordinated activity to mislead users. Finally, websites have the ability to opt-out of inclusion in Google News and/or Search results, and some do so.

7. Are websites on this list informed of their inclusion and given opportunities to appeal their inclusion?

8. If they are included, is their inclusion ever revisited?

Because the answers to these questions are related, we have grouped together our response to Questions Nos. 7-8.

As mentioned in our response to Question No. 6, for News, we assess sites for inclusion in Google News or Top News against a set of publicly disclosed Content Policies (https://support.google.com/news/publisher-center/answer/6204050?visit_id=637546080106902835-3077183141&rd=1). Sites that do not adhere to these policies are not eligible to appear on News surfaces or in our news carousels. Should the site rectify the issue causing a policy violation, we would re-assess the site for inclusion in Google News and other news surfaces.

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

9. What do you do to reduce the potential impact of Google employee bias in decisions to include a website on this list?

We build our products for everyone. 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.

As noted above in our response to Question 6, we design Search and other products with extraordinary care to serve our mission without political bias and have robust systems to ensure that employees' personal biases do not impact our products and that our policies are enforced evenly and consistently. We rely on extensive testing and controls, rigorous evaluation, and detailed metrics. Results from these evaluations and experiments go through a thorough review by experienced engineers and search analysts.

These robust controls are designed to prevent employee interference with our products, but if we ever thought that any Googler would attempt to manipulate products in order to achieve their personal political objectives, we'd take very strong action.

10. Why does Google suppress content from news websites which have been independently verified as reliable by outside organizations such as NewsGuard? Do you believe potentially biased content moderation by Google's Trust and Safety team is superior to independent verification by an outside panel such as NewsGuard?

As mentioned in response to Question No. 6, we assess sites for inclusion in Google News or other news surfaces against a set of publicly disclosed Content Policies, <https://support.google.com/news/producer/answer/6204050>. Sites that do not adhere to these policies are not eligible to appear on News surfaces.

Google News directly prohibits content that incites, promotes, or glorifies violence, harassment, or dangerous activities. Similarly, Google News does not allow sites or accounts that impersonate any person or organization, that misrepresent or conceal their ownership or primary purpose, or that engage in coordinated activity to mislead users.

The Economist conducted its own analysis and concluded that there was no evidence that ideology influences Google News results (<https://www.foxnews.com/tech/google-rewards-reputable-reporting-over-left-wing-politics-economist-study-says>).

The Honorable Gary Palmer (R-AL)

There have been a lot of discussions about Russian and Chinese attempts to influence our elections and political discussions. I think we need to realize that these disinformation campaigns are not only aimed at us, but also our allies. For example, Radio Free Asia highlighted that China mobilized 300,000 cyber operatives to target our ally Taiwan's 2018 and 2020 elections. A 2016 Harvard study found that a group allied with the Chinese Communist Party (CCP) produced 488 million "fake" social media posts a year to distract other internet users from news and online discussions painting the Communist Party in a negative light.

1. What is your company doing outside of the U.S. to protect our allies like Taiwan from the nefarious activities of the CCP?

2. What actions are you taking to keep individuals associated with the CCP from using your platform to attack the U.S. and our allies?

3. Do you pledge to work with our allies to push back against Russian and Chinese disinformation campaigns?

Because the answers to these questions are related, we have grouped together our response to Questions Nos. 1, 2, and 3.

Recent events have demonstrated how now more than ever it is important that we monitor our platforms to ensure the integrity and preservation of our democracy. Our efforts to combat disinformation apply to both foreign and domestic campaigns. Our misrepresentation policy on ads, for example, prohibits claims related to politics, social issues, or matters of public importance, that are demonstrably false and could significantly undermine participation or trust in an electoral or democratic process (https://support.google.com/adspolicy/answer/6020955?hl=en&ref_topic=1626336). Our Threat Analysis Group and Trust & Safety team monitor and disrupt account hijackings, inauthentic activity, disinformation campaigns, coordinated attacks, and other forms of abuse on our platforms 24/7, involving geographies around the world. We regularly take action against coordinated influence operations and, in May of last year, we began publishing a quarterly bulletin to provide more frequent updates about our efforts. Please see the initial bulletin at <https://www.blog.google/threat-analysis-group/updates-about-government-backed-hacking-and-disinformation/>, and our subsequent quarterly bulletins at <https://blog.google/threat-analysis-group/>.

Our efforts have identified a number of attempts by bad actors based abroad:

- In June 2020, we announced that our teams had identified phishing attempts from the Chinese group APT31 targeting the personal email accounts of Biden campaign staff and the Iranian group APT35 targeting the personal email accounts of Trump campaign staff. While we did not see evidence that these attempts were successful, we warned the targeted users and referred this information to federal law enforcement.
- In September 2020, we disabled 14 accounts linked to a party who was identified as a Russian agent and sanctioned by OFAC for attempting to influence the U.S. elections.
- In Q4 of 2020, we terminated 14,203 YouTube channels as part of our ongoing investigation into coordinated influence operations linked to China.

Our transparency report (<https://transparencyreport.google.com/government-removals/by-country?hl=en>) also provides data on government requests to remove content, including requests from the Chinese and Russian governments.

Finally, Google is also committed to ensuring election integrity and promoting freedoms abroad. For example, when Hong Kong's new national security law that took effect on June 30, 2020, Google, along with Facebook, Microsoft, and Twitter, ceased fulfilling user data requests from Hong Kong authorities. In addition, in December 2020, we announced an update to our Political Content policy to include coverage for Taiwan election ads (<https://support.google.com/adspolicy/answer/10120336?hl=en>). All advertisers running election ads in Taiwan must be verified by Google and observe our policies and applicable laws. We also provide a

publicly available transparency report regarding political advertising in Taiwan (<https://transparencyreport.google.com/political-ads/region/TW?hl=en>). For more information on our verification requirements, please see <https://support.google.com/adspolicy/troubleshooter/9973345>.

As we collectively face evolving threats from foreign actors, we will continue to work together with the U.S. military, government, U.S. allies, and other stakeholders.

The Honorable Michael C. Burgess, M.D. (R-TX)

According to Subsection (f)(3) of Section 230 of the Communications Decency Act, an “information content provider means any person or entity that is responsible, in whole or in part, for the creation or development of information provided through the Internet or any other interactive computer service.”

Based on this almost anyone can be a content provider. One of my biggest concerns over the years has been the accessibility of Terms of Service and whether they have been applied equally to all users – users who must accept them in order to access your services.

1. On your platforms, if you limit distribution of user created content or include additional information to accompany a user’s content or post, are you then a content creator since you were “responsible...in part, for the...development of information”?

No. Google is a technology company that facilitates the speech of a wide range of people and organizations from across the political spectrum, giving them a voice and new ways to reach their audiences. We provide a platform for creators, advertisers, academics, politicians, scientists, religious groups, and myriad others. Section 230 was passed recognizing the unique nature of platforms that host user-generated content and allows us to protect our users in the face of ever-evolving content and threats. Section 230 safeguards open access to information and free expression online and supports platforms’ ability to responsibly manage content.

In some cases, we may also act as an information content provider—for instance, when we publish a blog post on blog.google. In such instances, we would be treated as the information content provider with respect to that specific content, but that is clearly separate from our work providing platforms for other creators’ content. It does not and should not affect how our services overall are treated under the law.

We also recognize the legitimate questions raised by this Committee on Section 230 and would be pleased to continue our ongoing dialogue with Congress.

2. Should Congress be considering an update to the definition of an “information content provider”?

Section 230 has been hailed as the 26 words that created the Internet and vital to the success of small and medium sized online businesses. Your platform has benefited greatly from Section 230 liability protections.

We would be pleased to continue our ongoing dialogue with Congress. We are, however, concerned that changes to Section 230 and the definition of an “information content provider” would have far-reaching impacts and impose substantial burdens on lawful speech and freedom of expression online.

3. What would your service look like without Section 230?

Section 230 is what permits us to responsibly manage content on our services to protect users—and changes could jeopardize removals of terrorist content, spam/malware, scams, misinformation, manipulated media, and hate speech. In fact, Section 230 allows platforms and websites, big and small, across the entire internet, to responsibly manage content to keep users safe and promote access to information and free expression. Without Section 230, platforms would either over-filter content or not be able to filter content at all—in both cases making Internet services less useful.

Thanks to Section 230, consumers and businesses of all kinds benefit from unprecedented access to information and a vibrant digital economy. Today, more people have the opportunity to create content, start a business online, and have a voice than ever before. We are concerned that many recent proposals to change Section 230—including calls to repeal it altogether—would not serve that objective well. In fact, they would have unintended consequences—harming both free expression and the ability of platforms to take responsible action to protect users in the face of constantly evolving challenges.

4. What about when you were first starting out?

It is no accident that the greatest internet companies in the world were created in the United States. Section 230 has worked remarkably well, and we believe a cautious and thoughtful approach to potential changes is appropriate. Our platforms empower a wide range of people and organizations from across the political spectrum, giving them a voice and new ways to reach their audiences. Section 230 has enabled that, and millions of small and large platforms and websites across the internet rely on Section 230 to both keep users safe and promote free expression. Changes to 230 would disproportionately impact up-and-coming platforms without the resources to try and police every comment or defend every litigation, which could deter the next Google, Twitter, or Facebook, as the liability for third-party content would be too great. We would be pleased to continue our ongoing dialogue with Congress on these important questions.