

HPSCI Majority Questions for the Record Remote Hearing: “Emerging Trends in Online Foreign Influence Operations” June 18, 2020

From Chairman Schiff

For all witnesses

1. In the course of removing assessed networks engaged in CIB or foreign influence operations, does your company have standing policy or guidance with respect to proactively informing users who engaged with those removed accounts or the content? Why or why not?

YouTube’s goal is to provide context and authoritative information before or during viewer engagement, not after. For example, users will see an information panel pointing them to authoritative sources such as the CDC on the watch page of every video related to COVID-19. We also surface similar information panels for a variety of other topics that are prone to misinformation online.

Similarly, we have invested heavily to make sure that we surface authoritative content in Google Search results, which significantly reduces the spread of misinformation. This is in addition to the work that we do to develop and enforce policies against mis- and dis-information on our services, which we continue to invest in and iterate on.

Furthermore, the state of research is still open on the best ways to provide such notifications without inadvertently furthering harms – and we’d want to be very cautious about unintended consequences.

2. Can you please describe your company’s relationships or engagements with the national political parties, state parties, and individual campaigns, generally, and in the event you discovered a covert foreign influence operation targeting a specific candidate or political party?

Google regularly engages with national parties, state parties and campaigns through product training conducted by our Civic Outreach team. All of these trainings are currently virtual. Additionally, we have an email alias that is distributed to all civic entities where they can contact Google. In the event of an external party reaching out about a foreign influence group, we would route this immediately to our Trust &

Safety teams for review.

We also work with campaigns, elections officials, journalists, and others to ensure the security of the online platforms that they depend on. This includes the Advanced Protection Program (APP), which is available to all users and provides the strongest account protection that Google offers.

Additionally, we recently created a form for at-risk elections users to indicate that they'd like our account security team to more closely monitor their accounts for unusual and suspicious activity during the 2020 election cycle. This process is in addition to protections that APP users receive, and we are encouraging at-risk users to take advantage of both of these initiatives.

In regards to influence operations: our Threat Analysis Group (TAG), working with Google's Trust & Safety teams, identifies bad actors, disables their accounts, warns our users about them, and shares information with other companies and law enforcement officials.

a) Are these interactions regular, or would they depend on identification of a specific threat?

Please see response to Question 2.

b) If an individual candidate suspects they are being subjected to malign online activity, do they know who and how to contact at your company?

Google invests heavily in keeping our own platforms secure and have dedicated Trust & Safety teams working around the clock to prevent abuse.

People can use several tools to report instances of inappropriate content and our teams will review those reports and take appropriate action. On YouTube, any user may flag inappropriate content via our in-product flagging tool. In addition, we have an external email alias that is distributed to all civic entities and may be used to contact Google in such a case.

3. We've seen China in particular engage in overt use of its official diplomatic accounts and state-controlled media to shape the information space online and promote misleading or false narratives that advance its state strategic interests in an identifiably coordinated manner. Beyond mere labeling of

state-controlled media or identification of official foreign or diplomatic account as such:

a) Can you please describe your company’s approach to fact-checking or adding context to misleading or outright disinformation posted by these overt, foreign- linked accounts in a coordinated manner, which might allow users to readily understand the broader context or be directed to authoritative, credible sources about the claims?

We highlight third party fact-checks in Google and YouTube Search results in response to relevant user searches. Additionally, we display information panels on Google and YouTube Search result pages as well as on the watch pages of YouTube videos, to provide context from authoritative sources to users. For instance, we have long shown information panels indicating that a YouTube news channel receives government funding underneath the videos of that channel.

More recently, in response to the COVID-19 pandemic, we launched health information panels in YouTube Search results. These panels feature information on COVID-19 symptoms, prevention, and treatment, and links to local health authorities on the watch pages of COVID-19 related videos. Users will also see an information panel pointing them to authoritative sources such as the CDC on the watch page of every video related to COVID-19 (including those uploaded by state-funded news channels).

b) If a Facebook post, Tweet, or YouTube video created by a state-controlled media outlet promotes misleading or provably false narratives in apparent coordinated manner reasonably assessed to be in the service of that state’s interests, what steps might your respective platforms consider in terms of labeling, fact-checking, or providing context to users about such material?

Our approach to tackling disinformation in our products and services is based around a framework of three strategies: make quality count in our ranking systems; counteract malicious actors; and give users more context. In 2017, we began raising up authoritative voices on our platform, including news sources like CNN and Fox News, for news and information in search results and “watch next” panels. Millions of search queries and recommendations are getting this authoritative ranking treatment today, and we’re continuing to improve and expand our systems.

With regard to countering malicious actors, when we find attempts to conduct coordinated influence operations on our platforms, we swiftly remove such content from our platforms and terminate these actors' accounts. We take steps to prevent possible future attempts by the same actors, and routinely exchange information and share our findings with others in the industry. Over the past few years we have shared updates about our work to that end on the Google blog, and in May 2020, we introduced a more streamlined way of sharing such updates via a new, quarterly bulletin that describes actions we take against accounts that we attribute to coordinated influence operations (both foreign and domestic). Our actions against coordinated influence operations from January, February and March 2020 can be found in our [Q1 Bulletin](#).

Finally, we believe that providing additional context on content is an essential part of our work in combatting disinformation, and we provide that context through a number of product features on YouTube. For example, YouTube was an early industry leader in providing information panels below videos that indicate funding sources from publishers that receive public or government funding. In addition, during breaking news events, we may provide short previews of text-based news articles in search results along with a reminder that developing news can rapidly change.

In May, we launched fact check information panels for users in the U.S. to provide additional context by surfacing relevant, third-party fact-checked articles above search results for relevant queries. When a user enters a query seeking information that relates to a specific claim for which we have a relevant fact-checking article, we may display an info panel at the top of the search results that includes:

- The fact check article title
- A link to the article
- The publisher's name

If more than one relevant fact checking article exists, we will show a carousel that allows users to scroll through the available articles.

Our fact check information panel relies on an open network of third-party publishers and leverages the ClaimReview tagging system. All U.S. publishers are welcome to participate as long as they follow the publicly-available ClaimReview standards and are either a verified signatory of the International Fact-Checking Network's (IFCN) Code of Principles or are an authoritative publisher. Over a dozen U.S. publishers are participating today, including The Dispatch, FactCheck.org, PolitiFact, and The Washington Post Fact Checker, and we encourage more publishers and fact checkers to explore using ClaimReview.

4. Graphika’s June 16, 2010 report about the so-called “Secondary Infektion” group assessed it as having links to Russia and attempted to use false stories and outright forged materials to advance narratives favorable to Moscow.

a) Does your company have a policy governing the removal of “genuine,” provably hacked or stolen materials found on your platform, similar to the episode involving the hacked-and-dumped emails of Clinton Campaign Chair John Podesta in 2016? If so, please provide it in writing.

Google takes the threat of election interference via the release of hacked materials seriously. We recently updated our Community Guidelines for YouTube and will remove content that distributes hacked materials that are likely to interfere with democratic processes.

Under this policy, “hacking” includes traditional hacking methods such as exploiting security vulnerabilities, as well as other methods like phishing, spearphishing and other forms of social engineering. Examples of hacked materials that may pose a risk of interference in democratic processes include information disclosed to disparage a candidate or information that may delegitimize election results by questioning the integrity of the democratic process.

In addition, we may remove hacked material content from YouTube that includes certain private user information or poses a serious risk of physical harm to an individual. That said, we do not remove hacked material content if it is uploaded along high educational or documentary context, for example when included in a news reporting context by authoritative news sources.

b) Does this policy include or account for the posting of suspected or proven forgeries that were presented as genuine and was linked to a foreign influence operation? Or would your company otherwise prevent the sharing or re-posting of such forged content?

In addition to our policies related to hacked materials as outlined in response to Question 4(a), YouTube has clear policies that prohibit deceptive practices that take advantage of the YouTube community. Such deceptive practices include, but are not limited to, content that has been technically manipulated in a way that misleads users and poses risk of egregious harm, content that misleads voters about how to participate in an election or census, and content with misleading metadata or

thumbnails.

c) Do these or other policies cover content that might otherwise be illicitly obtained, e.g. a phone conversation that was recorded by a third party without the knowledge or consent of the calling or the called party, and then posted to Facebook, Twitter, or YouTube?

As with all content decisions, any enforcement of the YouTube Community Guidelines would depend on the full content and context of the video. As described in response to 4(a), we have policies that prohibit content that distributes hacked materials that are likely to interfere with democratic processes, with appropriate carve outs for educational or documentary content.

5. What changes has your company made to algorithms deployed on its internet platforms since 2017, especially with respect to limiting the reach or potential virality of extremist content and conspiracy theories?

Engineering teams at Google and YouTube continuously improve our ranking systems to elevate content from authoritative sources and to reduce the spread of borderline content or content that could misinform users in harmful ways. To that end, we invest in consistent development and experimentation. For example, in 2019, we ran over 464,065 experiments with trained external Search Raters and live tests, which resulted in more than 3,620 improvements to Google Search.

On YouTube, of the 6.1 million videos removed for violating our guidelines in Q1 2020, 93% were first flagged by machines rather than humans, and 77% of videos removed had 10 or fewer views. In addition, since making changes to our recommendations systems at the beginning of 2019, we've seen a 70% reduction in watch time from non-subscribed recommendations to borderline content and harmful misinformation content in the US, and we continue to hone our systems to further reduce this number.

a) How do you measure your success? b) Would you make public metrics so that we in Congress can judge these issues in a non-anecdotal fashion?

We publicly document our process for testing and assessing what changes are approved to launch when it comes to our Search algorithms, including the [detailed guidelines](#) that are used to train the external raters that are an integral part of these assessments.

Similarly, our [political ads Transparency Report](#) provides insights into our platform that any user, reporter, or researcher can easily use. It includes details about who is buying ads, how much they are spending, and a library of ad creatives with the dates they ran, the ad spend, targeting information and impressions.

Finally, on YouTube, we release a quarterly report on our [Community Guidelines Enforcement](#).

For Google

1. Mr. Salgado referenced the YouTube transparency report, as well as the quarterly bulletins on influence operations.

a) Are Google or YouTube contemplating taking similar steps as Twitter to create public archives of activity tied to foreign-linked influence campaigns? Why or why not?

There are significant differences between the service Twitter offers and the many services offered by Google, which warrant for different approaches towards transparency. We have already developed multiple ways to inform the public and third party experts of the nature and scope of policy violative activities taking place on our services, such as the YouTube Community Guidelines Enforcement Report, our annual Bad Ads reports, and our recent Threat Analysis Group quarterly bulletin, which describes actions we take against coordinated influence operations (both foreign and domestic).

We also publish and make available to download a [political ads Transparency Report](#) covering many countries including the U.S., where anyone can see the [election ads](#) themselves and can look up information such as who purchased a particular ad, how much they spent, or how many impressions the ad received.

We are always looking for ways to be more transparent and improve on the mechanisms we have in place, and will continue to make advancements.

b) Are Google or YouTube contemplating similar steps as Facebook, to collaborate with third-party, independent experts or research groups, such as Graphika or the Stanford Internet Observatory, and allow those entities to release coordinated detailed reports describing and contextualizing the

removals? Why or why not?

Google's Threat Analysis Group works with a variety of others, including law enforcement and third-party companies, to identify, receive, and disable threats. We have made major investments internally to address these issues, and we look to draw upon other valuable sources of expertise and intelligence as well. Similarly, we identify additional activity on our own platforms that we may share with our external partners to assist them in their investigations. We are always evaluating our approach and considering the best way to partner with organizations across the spectrum.

c) The YouTube transparency reports don't appear to have categories or any metrics specifically addressing coordinated inauthentic behavior or influence operations. Are Google or YouTube considering creating a separate category to quantify and publicly document removals tied to CIB or influence operations? Or provide additional granularity to users or the public about aspects such as countries assessed to be responsible for CIB or influence operations. Why or why not?

On any given day, Google's Threat Analysis Group (TAG) is tracking more than 270 targeted or government-backed attacker groups from more than 50 countries. Our team of analysts and security experts is focused on identifying and stopping issues like phishing campaigns, zero-day vulnerabilities, and hacking against Google, our products and our users.

The quarterly TAG bulletin we referenced in our response to Question A should help others who are also working to track these groups, such as researchers studying this issue, and we hope these updates can help confirm findings from security firms and others in the industry. We will also continue to share more detailed analysis of vulnerabilities we find, phishing and malware campaigns that we see, and other interesting or noteworthy trends across this space. As such, we are continuing to evaluate the appropriate transparency measures and how much detail we can provide. Input from the Committee is very valuable to us as we determine how best to adapt and expand the bulletin moving forward.

d) Will future quarterly bulletins on foreign influence published by the Threat Analysis Group offer more detailed breakdowns or other data about videos removed due to CIB or influence operations, and include information about the assessed sponsoring foreign entity? Why or why not?

Please see response to Question C.

2. Will Google commit to providing regular briefings or updates to the Committee as new foreign-linked CIB efforts, foreign influence operations, or other foreign malign cyber activity targeting the election are uncovered and removed from its properties, including from YouTube?

We remain vigilant in our efforts to provide users with authoritative information, protect their sites and accounts, and provide increased transparency about our efforts. We certainly can't do this important work alone and recognize that preventing platform abuse, combating disinformation, and protecting elections requires concerted effort and collaboration across the industry and government. We work with many others, including government agencies, academics, researchers and other companies, and we would be happy to work with the Committee as well.

3. Starting from January 1, 2018, how many videos and/or channels has YouTube removed because they were engaged in coordinated influence operations? Which nation-states or state-linked actors were implicated in those removals? How many views and/or subscribers did each video/channel have in those removals?

Google has terminated more than a thousand channels connected to various coordinated influence operations associated with countries such as Russia, China, and Iran. We have historically shared information via blog posts. Going forward, we are also using our quarterly TAG bulletin to share information about action taken against coordinated influence operations in a more systematic way. Input from the Committee is very valuable to us as we determine how best to adapt and expand the bulletin moving forward.

4. A 2019 Google blog described the removal of 210 channels associated with "coordinated influence operations" that "related to the ongoing protests in Hong Kong." It otherwise lacked any detail.

a) Why did Google release relatively little specific information about its removal, as compared with Twitter's and Facebook's related announcements about the same activity?

In our public blog, we described our work to disable 210 channels on YouTube when we discovered that the network behaved in a coordinated manner while uploading content related to the protests in Hong Kong. We also publicly disclosed tactics— e.g. the use of VPNs and other methods to disguise the origin—and worked closely with other companies to share information related to this operation. We believe it is important to increase transparency related to our work, which is why we recently introduced a quarterly bulletin that provides detailed information on our findings across our products.

b) Will Google release a list of the 210 channels removed? Or metrics about how many videos were hosted on each channel, the nature of the content hosted on those channels, how many subscribers each channel had, how many views the channels' videos received, what languages were used, etc.? Why or why not?

As discussed in responses to other questions from the Committee, Google is committed to publicly sharing our findings through a new quarterly bulletin that describes actions we have taken across our products related to coordinated influence campaigns. We continue to evaluate the appropriate transparency measures and how much detail we can provide, and appreciate the feedback of the Committee in this area as we continue to increase our transparency related to our work. We have extensively briefed Congressional Committees on all details, provided testimony and cooperated with law enforcement inquiries as well. We are happy to continue this discussion with the Committee.

5. Will YouTube publicly release the 1,100 videos and associated metadata about the IRA- linked activity it found from 2016-2017, properly redacted for privacy protection as necessary, to allow researchers and the public to see how this activity was intended to influence the political conversation during the last presidential election? Why or why not?

When we find attempts to conduct coordinated influence operations on our platforms, we work to quickly remove such content from our platforms and terminate these actors' accounts. Following termination of accounts, we then take steps to prevent possible future attempts by the same actors, and routinely exchange information and share our findings with others in the industry. We have previously disclosed findings publicly in blog posts, but beginning in May of this year we introduced our quarterly bulletin as a streamlined and consistent way to publicly disclose our work to find and take action against coordinated influence operations

across our platforms. We continue to evaluate the appropriate transparency measures and how much detail we can provide.

6. Has YouTube made adjustments to its algorithms since 2016 designed to hinder a sophisticated foreign adversary from gaming the YouTube recommendation engine in any way? If so, please describe.

YouTube works each day to provide timely, authoritative, and credible sources to anyone who comes to YouTube. To achieve this objective, our work is guided by four pillars: removing violative content; raising up authoritative content; reducing the spread of borderline content; and rewarding trusted creators—what we refer to as the 4Rs of responsibility.

Content that comes close to violating our Community Guidelines is a fraction of 1% of what’s watched on YouTube in the U.S. To reduce this even further, in January 2019, we launched [changes](#) to our recommendations systems to limit the spread of harmful misinformation and borderline content. The result is a 70% average drop in watch time of this content coming from non-subscribed recommendations in the U.S. We continue to hone and refine our systems, and in the last two years, we have made hundreds of changes to improve the quality of recommendations for users on YouTube.

YouTube has also invested heavily in research and development to stay ahead of new technologies and tactics that could be used by malicious actors, including technically-manipulated content. In 2018, we formed an Intelligence Desk to detect new trends surrounding inappropriate content and problematic behaviors, and to make sure our teams are prepared to address them before they become a larger issue. In addition, YouTube’s Trust & Safety team works closely with Google’s Threat Analysis Group to identify and take action against coordinated influence operations on the platform.

7. Has Google made adjustments to Search since 2016 designed to hinder a sophisticated foreign adversary from gaming Search’s algorithm or ranking system in any way? If so, please describe.

Our engineering teams at Google and YouTube continuously improve our ranking systems so that they continue to elevate content from authoritative sources and reduce the spread of borderline content or content that could misinform users in harmful ways. To that end, we invest in consistent development and

experimentation. For instance, in 2019, we ran over 464,065 experiments, with trained external Search Raters and live tests resulting in more than 3,620 improvements to Google Search.

Although we found limited improper activity on our platforms in 2016, we continue to take the integrity of our elections very seriously. We have a team dedicated to ensuring the integrity of election-related content and ads across our platforms, including combating potential foreign influence.

8. Since applying the state-backed media labels in 2018:

a) Has YouTube noticed any discernible changes in how frequently users watch or engage with that content? Or identified changes in trends in the subscribers?

Adding state-funded media info panels is one of a host of changes YouTube has made over the past couple of years to raise authoritative voices on the one hand, and reduce the spread of borderline content and harmful misinformation on the other. Since January 2019, we've also launched over 30 different changes to reduce recommendations of borderline content and harmful misinformation. None of these factors work in isolation, and it would be very challenging to assess what change in overall watchtime or engagement should be attributed to any specific one of them—including but not limited to the information panels that help users understand that a channel has received public or government funding. However, the overall result of our work is a 70% average drop in watch time of this borderline content coming from non-subscribed recommendations in the U.S. For subjects such as news, science and historical events, where accuracy and authoritativeness are key, we've re-doubled our efforts to raise authoritative sources to the top and introduced a suite of features to tackle this challenge holistically. While we don't track data related to viewer consumption of any one info panel, we have seen consumption on authoritative news partners' channels grow by 60 percent.

b) Put another way, based on the metrics and observations, did the labels have a measurable impact on consumption?

Please see the response above. It is worth noting that the information panel is designed to provide publisher context to allow a better understanding of the sources of news content that an individual may watch on YouTube. Inclusion of the information panel providing publisher context is based on information about the news publisher

made available by Wikipedia and other independent third-party sources. As we have stated publicly, it is not a comment by YouTube on the publisher's or video's editorial direction, or on a government's editorial influence.

c) Are videos posted by identified state-backed media treated differently by YouTube's recommendation engine or in how YouTube surfaces them in video search results? Please explain.

Political news and events can be subject to misinformation, so the availability of quality information sources is crucial. That's why we raise up authoritative voices, including news sources, for news and information in search results and "watch next" panels. Millions of search queries and recommendations are getting this authoritative ranking treatment today, and we're continuing to improve and expand our systems.

In addition to our efforts to raise authoritative voices, as noted in our response to Question 6, we also work to reduce recommendations of borderline content or harmful misinformation, which is only a fraction of 1% of what's watched on YouTube.

Part of our raise and reduce efforts include Top News and Breaking News shelves to highlight quality journalism, as well as information panels that indicate funding sources below videos from publishers that receive public or government funding. Both of these features play a key role in our efforts to raise authoritative voices.

From Representative Carson

For all witnesses

1. Can you provide a brief update on the policies that you companies currently use to address the threat deep fakes or other sophisticated manipulated media pose to users? What is your current approach, and how confident are you that you can identify and stop a foreign-connected deepfake as part of an attempted online influence operation?

We take the threat of manipulated media very seriously, whether they are AI-Generated or utilize low-tech edits. YouTube's [Community Guidelines](#) prohibit content that has been technically manipulated or doctored in a way that misleads users—beyond clips taken out of context—and that may pose a serious risk of egregious harm. Similar [policies](#) apply to our advertising services.

We are always looking into the potential for new technologies to cause personal or societal harm, and look to identify and address any gaps in our policies. For example, we updated our Google policy on involuntary pornographic imagery (colloquially referred to as ‘revenge porn’) to cover [fake](#) imagery in addition to [real](#) imagery.

Our research teams are working to develop better models to detect deepfakes, and regularly exchange on this complex technical challenge with other research organizations around the world. In addition, we continue to improve our ranking algorithms on Search, News, and YouTube so that we become better at elevating content from authoritative sources across the board.

2. I know that there was reporting in December about accounts associated with the Epoch Times media outlet as having used faked profile photos on Facebook. Has Facebook, or the other companies, identified any new deployments of deepfakes in a fashion such as this, particularly if linked to a state actor?

While we cannot speak to the Facebook accounts in question, we are cognizant of the growing concerns around AI-generated, hyper-realistic media often referred to as ‘deepfakes’. This is an issue that we take very seriously, and we are extensively investing in research and technical tools to identify such content. In addition to the steps we have taken against harmful uses of manipulated media across our products and services, we are also working with researchers and civil society around the world to help others undertake similar research:

- In [January 2019](#), we made audio synthetic media datasets available to researchers as part of a global, open competition to develop new detection models.
- In [September 2019](#), we expanded upon that work by partnering with Jigsaw to release a large dataset of visual deep fakes through the FaceForensics benchmark for researchers to use in their work as they develop new detectors.

We remain committed to make progress on this issue going into the 2020 elections and beyond, as we recognize its importance.

3. Throughout the recent protests in the wake of George Floyd’s murder, some white nationalist groups have pushed messages of hate and violence, in an attempt to undermine the legitimacy of the protest movement. One such white nationalist group, Identity Evropa, actually created a fake Twitter account, impersonated a left-wing Antifa activist, and explicitly called for violence during some of the most tense moments of the protests. With this example in

mind:

a. How do your companies assess and evaluate any attempts by foreign actors to manipulate the information environment or create chaos during such fast-moving and emotional charged events? Especially when weighed against social media's role as an engine for legitimate civic organizing and the airing of genuine political or social grievances, as we've seen nationwide this month?

YouTube is built on the premise of openness. Based on this open platform, millions of creators around the world have connected with global audiences and many of them have built thriving businesses in the process. But openness comes with its challenges, which is why we also have Community Guidelines that we update on an ongoing basis. For example, we made changes to our hate speech policy and our harassment policy in 2019. When you create a place designed to welcome many different voices, some will cross the line. Bad actors will try to exploit platforms for their own gain.

This is why we invest heavily in our systems and tools and our teams are constantly on the lookout for malicious actors that try to game our platforms. We take action against coordinated influence operations and publish quarterly reports on our Threat Analysis Group's blog. We also continue to communicate with other platforms and with relevant government agencies to enable rapid response to new developments.

b. Can your company provide an update on the procedures that it currently uses to identify content that incites violence? Are those processes automated, or how does that process currently work? What definitions are used, since I imagine it's not always clear-cut?

We have policies that prohibit violent extremism, hate speech, and incitement to violence in content hosted on our services. We rely on a combination of people and technology to flag violative content and enforce our policies. Once potentially problematic content is flagged, human review verifies whether it indeed violates our policies. Although technology has become very good at identifying some kinds of controversial content, people are best at assessing context.

If the content violates our policies, it is removed. The account that posted the content generally receives a strike, and multiple strikes leads to account termination.

Machine learning now helps us take down violent extremist content before it has been widely viewed. On YouTube, between October and December 2019, approximately 90% of the videos uploaded that were removed for violating our Violent Extremism policy were taken down before they had 10 views. The two million videos our teams have manually reviewed provide large volumes of training examples that help improve the machine learning flagging technology.

From Representative Swalwell:

For all witnesses

1. Do your platforms have a policy to combat anti-vaccine misinformation in posts by users? Does that policy extend beyond demonetization, if relevant? If so, how?

We understand that trustworthiness is essential to users, so we've been investing in new product features to prominently surface authoritative sources.

On YouTube, we've improved how we suggest videos in our search results and through our "Up Next" recommendation feature so that when viewers watch content from authoritative sources, they'll receive recommendations to watch other videos from additional authoritative sources. This change reduces recommendations of content that can misinform users in harmful ways, including certain types of anti-vaccination videos.

In Google Search, we also strive to give our users the most relevant, reliable and highest quality information as quickly as possible. We work to prioritize results from authoritative sources for queries that pertain to topics like health that may be susceptible to misinformation. When available, we also show fact-check labels in Google Search and News to signpost where claims have been verified by independent fact checkers.

On both YouTube and Search, we highlight information on medical topics through stand-alone features that reflect broad scientific consensus on issues where such consensus exists. On Search, through features like Knowledge Panels, we are able to surface authoritative information about hundreds of medical conditions (including measles and other vaccine-preventable diseases). These Knowledge Panels include information from authoritative sources like the Mayo Clinic and have been evaluated by medical professionals. YouTube also works to empower users by giving them additional contextual information on both search results and watch pages so

consumers can inform themselves about the content with which they engage on the platform. We've expanded this feature to more topics, including the MMR vaccine and a new category of information panel for vaccines in general.

Although we invest heavily in products to address misinformation, we also welcome the efforts of others seeking to create solutions. We will continue to evaluate our efforts, and the efforts of other organizations, to help improve our products.

When it comes to monetization, we have strict policies for both our advertiser and publisher partners. Our policies are designed to protect users on our platforms and we've specifically taken an aggressive approach to preventing harmful and dangerous health content from monetizing. Under our policies, we currently prohibit [ad](#) and [publisher](#) content that makes harmful medical claims about disease diagnosis, prevention and cures. As part of these policies, we prohibit content that encourages users to forgo treatment as well as anti-vaccine promotions, among other examples.

2. Do your platforms have a policy to combat public health misinformation in posts by users? Does that policy extend beyond demonetization, if relevant? If so, how?

On YouTube, our Community Guidelines prohibit content that encourages dangerous or illegal activities that risk serious physical harm or death, including certain types of medical misinformation. As the COVID-19 situation has evolved, we have partnered closely with global and local health authorities to ensure our policy definition and enforcement is effective in removing content where there is a serious risk of egregious harm. Our policies prohibit, for example, content that denies the existence of the coronavirus or encourages the use of home remedies in place of medical treatment. We also prohibit content that explicitly disputes the efficacy of global or local health authority advice regarding social distancing that may lead people to act against that guidance. These policies apply to videos, video metadata, and comments. In addition to the Community Guidelines, our [Advertiser-friendly guidelines](#) set out additional standards for content that is monetized via advertising.

On [Google Ads](#) and [Google for Publishers](#), our Dangerous or Derogatory Content policies prohibit the monetization of content "promoting or advocating for harmful health or medical claims or practices". This policy combats harmful misinformation around health and medical treatments. For example, under this policy, we demonetise publisher content that includes claims about the propagation of COVID-19 that contradicts authoritative sources such as the CDC (e.g. theories involving 5G towers as a transmission vector). On Google Ads, as well as Google Shopping, our [policies](#) do

not allow ads that potentially capitalize on or lack reasonable sensitivity towards a sensitive event, such as a public health emergency – and we began treating the COVID-19 crisis as a Sensitive Event all around the world by end of January 2020 (see [ads policy update section](#) for more details).

On Google Search, our [medical topics policy](#) applies to information we've highlighted via our Knowledge Panels and related Search features that highlight information from our Knowledge Graph. We strive to show information that reflects scientific consensus and evidence-based best practices, since we consider this content high quality. To this end, if this highlighted information runs contrary to general scientific consensus, we reserve the right to correct or remove the information from the feature.

3. Has One American News Network (OANN) had videos or posts removed from your platform? If so, how many and for what reasons?

YouTube removes content that violates our policies whether first flagged through user reports or automated flagging systems. Once a video is flagged for review, our trained human reviewers assess whether the content in question violates our policies and apply, as necessary, appropriate enforcement. Our automated flagging systems are used to identify and remove content such as spam automatically, as well as re-uploads of content we've already reviewed and determined violates our policies. While we are unable to provide a complete history of enforcement decisions for each channel, we do publicly report data specifically related to video and channel removals in our quarterly [Transparency Report](#).

4. Has Fox News had videos or posts taken removed from your platform? If so, how many and for what reasons?

Please see response to Question 3.

5. Has The Epoch Times had videos or posts removed from your platform? If so, how many and for what reasons?

Please see response to Question 3.

6. On June 18, 2020, Facebook, Instagram, and Twitter removed a Trump campaign ad featuring a symbol (a red inverted triangle) used by Nazis to designate political prisoners in concentration camps. Facebook, which owns

Instagram, stated, “We removed these posts and ads for violating our policy against organized hate. Our policy prohibits using a banned hate group’s symbol to identify political prisoners without the context that condemns or discusses the symbol.”

a. How many symbols of hate would a campaign or candidate have to run before the campaign's account or page would be taken down from your platform?

We don’t allow ads or destinations that display shocking content or promote hatred, intolerance, discrimination, or violence. We have strict ad policies to protect users against [Dangerous or Derogatory](#) content. Under this policy we don’t allow, for example, content that incites hatred against, promotes discrimination of, or disparages an individual or group on the basis of their race or ethnic origin, religion, disability, age, nationality, veteran status, sexual orientation, gender, gender identity, or other characteristic that is associated with systemic discrimination or marginalization.

To give you a better understanding about the scale of our enforcement, [in 2019 alone](#), we blocked and removed 2.7 billion bad ads—including 670,000 ads for violating our Dangerous or Derogatory content policy. We also suspended nearly 1 million advertiser accounts for policy violations. On the publisher side, we terminated over 1.2 million accounts and removed ads from over 21 million web pages that are part of our publisher network for violating our policies. Terminating accounts—not just removing an individual ad or page—is an especially effective enforcement tool that we use if advertisers or publishers engage in egregious policy violations or have a history of violating policy.

YouTube has always had [rules of the road](#), including a longstanding policy against hate speech. In 2017, we [introduced](#) a tougher stance towards videos with supremacist content, including limiting recommendations and features like comments and the ability to share the video. This step dramatically reduced views to these videos (on average 80%). In 2019, we took another step in our hate speech [policy](#) by specifically prohibiting 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 provide detailed information about our hate speech policy and enforcement in our quarterly [Community Guidelines Enforcement Report](#).

The first time a channel's content is found to violate any of YouTube's Community Guidelines, they receive a warning. A second violation results in a strike, which stays on a user's account for 90 days, and revocation of certain privileges for one week. Another violation within that 90-day period results in a second strike, and the inability to post content for 2 weeks. Receiving 3 strikes in a 90-day period results in permanent termination of a YouTube channel. Users may appeal strikes, and we make data about appeals and reinstatements available on our Community Guidelines Enforcement Report. We make information about our strike system available to users at <https://support.google.com/youtube/answer/2802032>.

b. How many false or partly false posts, videos, or ads would a campaign or candidate have to run before the campaign or candidate's account or page would be taken down from your platform? Or would consistent posting of false or partly false posts or ads go unenforced?

Please see response to Question A.

c. Have campaign or candidate accounts, pages, or channels associated with U.S. persons been taken down because of repeated posting or advertising of false or partly false information? If so, how many? And if not, have you taken other actions against said accounts, pages, or channels?

We enforce our advertising policies neutrally, without regard to an advertiser's political affiliation, including our [policies for election ads](#). For election ads in the US - ads that feature 1) a current officeholder or candidate for an elected federal office, 2) a current officeholder or candidate for a state-level elected office, 3) a federal or state level political party or 4) a state level ballot measure - we provide both in-ad disclosures and a [transparency report](#) that shows the actual content of the ads themselves, who paid for them, how much they spent, how many people saw them, and how they were targeted. The transparency report also indicates instances where we've removed election ads for policy violations.

d. Have campaign or candidate accounts, pages, or channels associated with U.S. persons been taken down because of repeated use - whether through advertising or not - of symbols of hate and/or violating anti-hate policies? If so, how many? And if not, have you taken other actions against said accounts, pages, or channels?

Please see response to Question C.

e. Are your platforms considering implementing new policies or revising existing ones to address the issues raised in questions 7a through 7d?

Google has a crucial stake in a healthy and sustainable digital advertising ecosystem. Every day, we invest significant team hours and technological resources to minimize content that violates our policies and stop malicious actors. We have a dedicated workforce working around the clock across our teams to make sure we're protecting our users and enabling a safe ecosystem for advertisers and publishers. As mentioned above, we have a wide set of ads policies that we enforce vigorously. We are also constantly evaluating and updating our policies—in 2019 alone we introduced 31 new ads policies.

YouTube systematically reviews and re-reviews all 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.

For Google

1. Do YouTube's comments fall under the same policy as Google Display ads, your comments policy on YouTube?

YouTube's Community Guidelines set the rules of the road for what content is prohibited on YouTube, including user comments. We rely on a combination of smart detection technology and human reviewers to flag content for our review teams. We are transparent about content removal and provide quarterly updates on how we enforce our policies, including as they pertain to comments, in our quarterly [YouTube Community Guidelines Enforcement Report](#). In Q1 2020, we removed 693,579,605 comments, 99.6% of which were detected by automated systems. In addition, we have an appeals process so that users who believe content was removed in error may ask us to reconsider. We provide data on the volume of appeals and reinstatements in our Community Guidelines Enforcement Report.

2. What were the total revenues earned by YouTube content creators associated with RT, OANN, Fox News, and Epoch Times in 2019? And so far this

year?

All channels on YouTube need to comply with our Community Guidelines and, in order to monetize, channels must comply with the YouTube Partner Program policies, which include our Advertiser-Friendly Guidelines. While we are unable to share specific data related to individual content creators, we regularly review and remove channels that are not in compliance with our policies.

3. What are the policies designed to limit public health misinformation advertised on YouTube and Google’s other platforms? And do you have any data or metrics that speak to the effectiveness of YouTube’s and Google’s enforcement of these policies to suppress public health misinformation?

As discussed above, YouTube’s Community Guidelines prohibit certain types of harmful health misinformation, including misinformation relating to COVID-19 such as content that denies the existence of the coronavirus or encourages the use of home remedies in place of medical treatment. We also prohibit content that explicitly disputes the efficacy of global or local health authority advice regarding social distancing that may lead people to act against that guidance. As the COVID-19 situation has evolved, we have partnered closely with global and local health authorities to ensure our policy definition and enforcement is effective in removing content where there is a serious risk of egregious harm. YouTube’s advertiser-friendly guidelines set out additional requirements for content monetized by ads on our platform, and these include guidelines specific to COVID-19 content. For example, we don’t allow ads to run against any content that claims the pandemic is a hoax or that corporations created the virus.

We have also taken additional steps to highlight content from authoritative sources when people search for COVID-19, and have also provided information panels to provide additional context from high-quality sources on content related to COVID-19. Since launch, there have been over 300 billion impressions on our information panels, which we display on our homepage and on videos and searches about COVID-19. We’ve recently expanded our fact check information panels to the United States to help address the challenge of misinformation that comes up quickly as part of a fast-moving news cycle. As a result our work, global watchtime on authoritative news content grew by more than 75% in the first three months of 2020.

On [Google Ads](#) and [Google Publisher](#) our Dangerous or Derogatory Content policies prohibit the monetization of content "promoting or advocating for harmful health or

medical claims or practices". This policy combats harmful misinformation around health and medical treatments. For example, under this policy, we demonetise publisher content that includes claims about the propagation of COVID-19 that contradict e.g., the CDC (such as theories involving 5G towers as a transmission vector).

From Rep. Maloney

For all witnesses

1. Recognizing that strides have been made since 2016 through 2018:

a) Is it your company's stance that the current volume and types of indicators, data, and/or metadata about potential foreign influence activity shared both within the industry and between the industry and the U.S. government are sufficient for protecting our national conversation and elections from foreign influence or interference moving forward?

Google has been sharing information with others in industry and has been receiving leads as well. The communication has been quick and secure, and we've only improved since 2016. Based on threat briefings and discussions about trends and techniques with those companies and government entities, we see no indication that the sharing is deficient. We remain concerned about there being potentially useful information with the government that is unavailable as a result of classification issues or intra-agency sharing complications.

b) What limits imposed by U.S. law or regulations might prevent your company from maximally sharing data or metadata associated with high-confidence foreign influence operations/CIB with U.S. law enforcement? c) How might relevant changes to the Secure Communications Act (SCA), the Electronic Communications Privacy Act (ECPA), Cybersecurity Information Sharing Act (CISA), or the Section 230 Communications Decency Act (CDA) help or harm your companies' efforts to prevent foreign influence from infiltrating your platforms? d) Would considerations such as creating a "safe harbor" provision, or clearly delineating that assessed foreign influence actors don't have claim to the same data privacy protections as genuine users, affect those stances?

Google has been able to share information with agencies under current law in an efficient and effective manner, consistent with current laws and regulations.

Permissions in the SCA allow for this responsible sharing while ensuring appropriate privacy protections for the users of the services.

Section 230 also plays an important role. It has created an internet ecosystem where commerce, innovation, and free expression thrive, while enabling providers to develop innovative content detection mechanisms and take aggressive steps to fight online abuse. The law not only clarifies where services can be held liable for third-party content, but also creates the legal certainty necessary for services to take swift action against harmful content of all types, and thereby have helpful information to share with others to also join the battle against abuse.

The European Court of Justice's ruling in the case of *Data Protection Commissioner v Facebook Ireland and Maximillian Schrems* addresses data sharing and its impacts will need to be further clarified. We urge the US to work with the EU and member states to address the issues raised in the ruling.

2. Would your company find valuable an Information Sharing and Analysis Center (ISAC) or equivalent formalized mechanism devoted specifically to data-sharing about potential foreign-linked influence operations? Would your company support a leading role in an ISAC or equivalent? Why or why not?

Our focus has been on finding the best way to share the information needed to deal with a shifting threat with the entities that need it, and do so with agility and speed. The information and threat sharing arrangements that have developed organically have proven to be effective and efficient for this purpose. There is excellent sharing currently on a day-to-day basis between the industry players with direct links between security and other teams. We have found that these arrangements allow for nimble adjustment of focus, which has proven to be critical for this threat, which is itself quick to morph. We do not see that imposing an artificial ISAC organizational structure, with the regime of bylaws and other inherent trappings that come with formality, would provide any additional ability to allow us to work together to counter these threats. Indeed, we would be concerned that doing so would distract from the real work, with form diluting substance and protocol supplanting action.