PRIVILEGED AND CONFIDENTIAL
ATTORNEY WORK PRODUCT

Questions For the Record From:
Blumenthal
Coons
Durbin
Feinstein
Grassley
Klobuchar
Whitehouse

Deadlines	Action Items
11/15/17 noon	**All information provided by teams**
11/1/17 EOD	**Send to Kent and Rick**
11/17/17	**Finalize Drafts**
11/21/17	**Electronic versions of responses due**

Questions for the Record
for Mr. Richard Salgado, Director of Law Enforcement and Information Security, Google
Submitted by Senator Richard Blumenthal
November 7, 2017

A. Data Production to the Judiciary Committee

1. What is the status of your efforts to identify Russian advertisements intended to influence the election?

We have conducted an extensive review of this issue including developing a list of actors we know or suspect were involved in this effort from our research of publicly available information, the work of our security team, as well as leads we received from others in the industry, and applying those leads to nearly twenty of Google's products, including all Ads products. We identified limited activity on our platforms, but did identify two Ads accounts with approximately $4,700 of spend. In order to validate our findings, we broadly reviewed all political ads from June 2015 until the election last November that had even the loosest connection to Russia, which substantiated that we had identified the ads connected to this effort.

Our investigation is ongoing, we continue to request and receive leads from peers in our industry, and will be happy to continue cooperating with Congressional investigations on this topic.

2. When can you provide these advertisements to this Committee?

We previously provided both electronic and hardcopy versions of the ads associated with accounts we identified as connected to this effort to the Committee.

B. Role of Social Media Consultants/Social Media Management Companies

3. Have you done any analysis to determine the degree to which Russia relied on social media consultants/management companies to purchase ads designed to influence the election?

We found limited activity on our platforms. Our investigation has revealed there were only two Ads accounts associated with this effort. We saw no evidence that those accounts used media consultants or management companies. To the extent any type of consultancy or agency is involved in purchasing Election Ads on our systems, our new transparency policies will require that the consultant or agency identify the end advertiser placing the ad.

4. To what degree will your new transparency policies help the public identify ads purchased by foreign governments if these ads are purchased through social media consultants/management companies?

Our customers generally do not use social media consultants or social media management companies with our platforms, but do use agencies. To the extent any type of consultancy or agency is involved in purchasing an Election Ad on our system, our new transparency policies will require that the consultant or agency identify the end advertiser placing the ad.

5. Are you working with social media consultants/management companies to ensure that they cannot be used to shield political ads from transparency efforts?

To the extent any type of consultancy or agency is involved in purchasing an Election Ad on our system, our new transparency policies will require that the consultant or agency identify the end advertiser placing the ad.

C. Responding to Search Engine Optimization/Search Engine Manipulation

6. It is my understanding that you have systems for detecting attempts to manipulate search results. Are you using these detection systems to identify manipulation originating in Russia? If so, what have you been able to identify?

Google's mission is to organize the world's information and to make it universally accessible and useful. Google Search's ability to surface authoritative content from the web responsive to a search query is integral to that mission. Google has long had numerous systems in place, both automated and manual, to detect and address manipulative and deceptive behavior in Search (what we call "webspam") and we apply those systems broadly. We enforce our policies whether the efforts to forge the appearance of signals that Google and other search engines use to rank content are at the hands of a small business in the United States seeking a competitive advantage or a state-sponsored actor with larger ambitions. For Google Search, we also updated our Search Quality Rater Guidelines and our evaluation test sets to help identify misleading information and have used this data to improve our search algorithms, resulting in more responsive, higher quality, and more authoritative Search results.

We have confronted attempts at webspam from Russia over the years, just as we have from countries all over the globe, and have dedicated Russian-language analysts to help combat them. The sites we were able to identify as potentially associated with those whom we believe were part of this effort were carefully scrutinized by our webspam team as part of our investigation. We found no attempts at manipulation. We did find that some of the sites had low-quality content according to algorithms we have developed to detect content of minimal use to our users. Our algorithms automatically reduced their ranking without human intervention. For more information on our webspam process as well as a review of other common webspam techniques, please see Google's Webmaster Guidelines at: https://support.google.com/webmasters/answer/35769.

7. How are you addressing the challenge of search engine optimization or search engine manipulation in this context? Are you prioritizing this issue?

Addressing the challenge of search engine optimization and related efforts to artificially inflate search result rankings ("webspam") is a priority for Google. We continue to invest in and enhance our numerous systems, both automated and manual, for detecting and correcting webspam. For example, we have a vast array of algorithmic and manual methods to detect the use of "link farms" to make content appear to be more popular than it really is. We publish guidelines for webmasters that go into more detail, which you can find at https://support.google.com/webmasters/answer/35769. For YouTube, we employ a sophisticated spam and security-breach detection system to detect anomalous behavior and catch people trying to inflate view counts of videos or numbers of subscribers.

8. Is there a "paper trail" for this sort of manipulation? What other challenges do you face in identifying search engine manipulation?

Attempts to deceive or manipulate signals that we use to rank potential Search results present a massive challenge in light of the 130 trillion pages we crawl on the web, particularly when combined with the efforts of spammers who continue to find new techniques to deceive signals that play a role in rankings. We are constantly investing in resources and updating our guidelines in order to stay ahead of would-be bad actors. For example, we recently updated our Search Quality Rater Guidelines and evaluations to help identify misleading information, and thereby surface more responsive and authoritative content from the web.

D. Transparency for Issue-Based Advertisements

9. How do you intend to bring greater transparency to issue-based advertisements that potentially originate in Russia?

Google is very concerned about attempts to undermine democratic elections and deeply committed to getting this right. We have updated our advertising guidelines to prohibit ads on sites that misrepresent themselves. We are committed to working with Congress, law enforcement, others in our industry, and the NGO community to strengthen protections around elections, ensure the security of users, and help combat disinformation.

In addition, we have announced a number of measures to enhance transparency within election advertising:

- **Transparency Report.** In 2018, we'll release a transparency report for Election Ads, which will share data about who is buying election-related ads on our platforms and how much money is being spent.
- **Creative Library.** We'll also introduce a publicly accessible database of Election Ads purchased on AdWords and YouTube (with information about who bought each ad). That means people will not only be able to learn more about who's buying election-related ads on our platforms; they'll be able to see the ads themselves, regardless of to whom they were shown.
- **In-ad disclosures.** Going forward, we'll identify the names of advertisers running election-related campaigns on Search, YouTube, and the Google Display Network.
- **Verification program.** U.S. law restricts entities outside the United States from running

election-related ads. We'll reinforce our existing protections by requiring that advertisers proactively identify who they are and where they are based before running any election-related ads. As they do, we'll verify that they are permitted to run U.S. election campaigns through our own checks.

10. Do you believe your recent transparency policies go far enough, or do you intend to build on them?

We think that the additional measures we've described will considerably enhance the transparency of election advertising on our platforms. Our commitment to transparency is ongoing and we will continue to explore ways to increase Election Ads transparency, including by working with others in the industry, researchers, NGOs, and Congress on ongoing transparency efforts.

E. Speed of Transparency Efforts

11. I want to impress upon you the importance of dealing with this issue swiftly—otherwise we may be facing the same situation in November 2018 as we did in November 2016. What assurances can you provide this Committee that your company is working as fast as possible to implement new transparency features?

We consider transparency to be a high priority and have dedicated resources across the company to ensure we meet our recently announced goals to increase transparency around Election Ads. We are committed to implementing new transparency features before the November 2018 election season.

F. Impact of Disinformation Campaigns

12. Do you have information regarding how many voters were impacted by posts that were part of foreign disinformation campaigns? Do you have data on how these posts may have impacted election results?

It is difficult to measure what, if any, impact the limited activity on our platforms had on the election. But we are committed to doing our part, and recognize that we must work together across government, civil society, and the private sector to address these complex issues at their root. We look forward to continuing to work with this Committee as it takes on this important issue.

13. When can you provide this information to this Committee?

We do not anticipate being able to make a determination as to what, if any, impact the limited activity we found on our platforms had on the election. But, again, we are committed to doing our part to combat foreign disinformation campaigns. We look forward to continuing to work with this Committee as it takes on this important issue.

Back to Top

Richard Salgado – Extremist Content and Russian Disinformation Online
Questions for the Record
Submitted November 7, 2017
QUESTIONS FROM SENATOR COONS

1. On January 6, 2017, the U.S. Intelligence Community released a public report that concluded, "Vladimir Putin ordered an influence campaign in 2016 aimed at the US presidential election." The documents provided to the Committee confirm the intelligence community's conclusion and highlight the need for online platforms to work with the government to prevent threats going forward.

   a. Who is your contact person at the Department of Justice and/or the FBI as you work to counter these ongoing threats?

Who we work with depends on the particular investigation, and we understand the need to work with law enforcement, including the FBI, the Department of Justice (including U.S. Attorney's Offices), and other agencies, as is consistent with our policies and the law.

   b. Have the DOJ or FBI made any recommendations to you for preventing interference going forward?

Google is deeply concerned about attempts to undermine democratic elections. We are committed to working with Congress, law enforcement, others in industry, and the NGO community to strengthen protections around elections. To date, we have not received specific recommendations on this issue, but we welcome input from law enforcement and Congress on preventing and detecting any abuses of our platforms.

2. Do you believe that computer algorithms and machine learning are sufficient to catch foreign political ads, fake accounts, and false information?

Threats to our systems are continuously evolving. We serve billions of users every day, so our solutions need to work at scale. We rely on highly-trained individuals from our Trust and Safety and Security teams who work closely with machine learning tools and our algorithms to ensure our platforms are protected and there is adherence to our policies.

   a. What new technologies or capabilities will you introduce to prevent these abuses?

We face motivated and resourceful attackers, and we are continually evolving our tools to stay ahead of ever-changing threats. We will continue to build industry-leading security systems and deploy those tools in our products. Those tools will be aimed at protecting our physical and network security, but also detecting and preventing the artificial boosting of content, spam, and other attempts to manipulate our systems. As threats evolve, we will continue to adapt to understand and prevent new attempts to misuse our platforms and will continue to expand our use of cutting-edge technology to protect our users. We will also be making political advertising more transparent, easier for users to understand, and even more secure.

b. At the hearing, you testified that Google has thousands of people across various teams to manage problematic content. What do these employees do to improve safety and security?

We have a global team of thousands of policy experts, reviewers, product managers, and data scientists focused on creating, maintaining, and enforcing our policies. Through a combination of sophisticated algorithms and other technologies and human review, we both proactively look for violations and respond to complaints. We take this work very seriously; in 2016 alone we removed 1.7B ads for violating policies. Our policies evolve to deal with emerging issues.

c. How many employees will be dedicated to preventing foreign political ads from being published on the platform?

We will rely on the thousands of policy experts, reviewers, product managers, and data scientists focused on creating, maintaining, and enforcing all of our policies and ensuring the security of our systems. This is in addition to the significant resources we dedicate to automated/machine learning methods we use to protect our systems from manipulation and misuse.

3. Undetected bot accounts can quickly disseminate false news. What improvements are you making to counter increasingly sophisticated bots?

We have developed robust protections over the years to address attempts to manipulate our systems by bots or other schemes, like link farms. (You can learn more in our webmaster guidelines: https://support.google.com/webmasters/answer/35769.) We use both algorithmic and manual methods and we deploy these across our products including Search and YouTube. We have not, however, seen the same type of social media bots that have been reported on other platforms.

4. Russian operatives were able to increase their influence by hacking or purchasing online accounts that were originally authentic but no longer maintained by their owners. In fact, buying unmaintained accounts has become a cottage industry. What steps are you taking to prevent unmaintained accounts from falling into the hands of inauthentic users?

Our systems rely on a host of inputs about historical use and pattern recognition across various services in an effort to detect if an account creation or login is likely to be abusive. The system operates to block “bad” account creation or to close groups of such accounts. We prevent users from creating a large number of Google Accounts in a short time period if our systems detect that the user might be abusive. We also require verification, aimed at detecting if a bot is attempting to access or create an account, if we detect suspicious conduct.

5. As we saw with the Comet Pizza incident, where a man brought a gun into a D.C. pizza restaurant based on false reports that criminal activity was occurring there, fake news can stoke hatred and risk violence. What is Google doing to prevent the proliferation of fake news across the site?

We take misinformation on our platforms very seriously, and we have put a lot of effort into curbing misinformation in our products–from better ranking algorithms that prioritize authoritative sources, to tougher policies against monetization of misrepresentative content. 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 are looking at. Publishers who review third-party claims or rumors can showcase their work on Google News through fact-check labels and in Google Search through fact-check cards.

To help ensure Google does not monetize content designed to mislead users, we have implemented a new policy for our AdSense publishers that explicitly bans ads on any site that misrepresents, misstates, or conceals information about the publisher, the publisher's content, or the primary purpose of the site. For Google Search, we updated our Search Quality Rater Guidelines and our evaluation test sets to help identify misleading information and unexpected offensive results, and have used this data to improve our search algorithms. This results in higher quality and more authoritative Search results. We will continue to work on preventing the spread of misinformation by partnering with the journalism industry and civil society to help people understand what they see online and to support the creation of quality content.

6. Does Google support the Honest Ads Act? If you do not support this bill or are unable to commit to a position, please explain why.

We support the goal of the Honest Ads Act, namely to enhance the integrity of American elections via transparency and accountability of digital political advertisements, and we look forward to working with Congress as they address these challenges. We also fully support more transparency, which is why we will (i) identify the names of advertisers running election-related campaigns on Search, YouTube, and the Google Display Network within Election Ads; (ii) publish our first ever transparency report for Election Ads that will detail who is buying election-related ads on our platforms and how much is being spent; and (iii) introduce a database, accessible for public research, of Election Ads purchased on AdWords and YouTube (with information about who bought each ad).

7. Can you assure us that electioneering ads will include permanently displayed disclosure notifications like in print or television ads? If you cannot, please explain why.

Yes. As we recently announced, we plan on identifying the names of advertisers running election-related campaigns on Search, YouTube, and the Google Display Network via permanent disclosures.

8. What reforms will Google enact to address issue ads that do not mention political candidates by name?

We are committed to working with Congress, law enforcement, others in our industry, and the NGO community to strengthen protections around elections, ensure the security of users, and help combat disinformation. We are dealing with difficult questions that require the balancing of free expression, access to information, and the need to provide high quality content to our users. There are no easy answers here, but we are deeply committed to getting this right.

9. Foreign entities will continue to try to use social media to interfere with U.S. elections. Has Google identified attempts by foreign entities to interfere with post-2016 elections? Please describe such attempts.

While we identified very limited activity on our platforms prior to the 2016 election, we agree this threat is evolving. Our investigation and monitoring efforts to prevent and detect this misuse are ongoing. We have and will continue to work with the Committee in this effort.

Back to Top

Written Questions from Senator Richard J. Durbin
Senate Judiciary Subcommittee on Crime and Terrorism
Hearing on Extremist Content and Russian Disinformation Online: Working with Tech to Find Solutions
November 7, 2017

For questions with subparts, please answer each subpart separately.

<u>Questions for Richard Salgado, Google</u>

1. On January 6, the U.S. Intelligence Community issued a report on Russian election interference and described what happened last year as the "new normal in Russian influence efforts." The IC Report said "we assess Moscow will apply lessons learned from its campaign aimed at the U.S. presidential election to future influence efforts in the United States and worldwide."

We are less than a year away from Election Day in 2018. The campaign season will be upon us before we know it. We do not have much time to safeguard our nation's social media platforms against Russian disinformation efforts and election propaganda.

a. **Will your company be ready before Election Day 2018 to reassure Americans that your platform is not tainted by foreign disinformation or influence efforts?**

The activity on our platforms was limited during the 2016 election cycle, and we believe that was in large part due to the controls we had in place prior to the 2016 election. We understand the importance of maintaining and enhancing those controls as we go into the 2018 election season. We will continue to expand our use of cutting-edge technology to protect our users and will continue working with others to help ensure that our platforms are not misused going into the 2018 election season.

For example, we recently introduced the Advanced Protection Program, a new level of account protection designed for those with an elevated risk of attack. We are offering this to all political campaigns and elected officials in the United States, to minimize the risk of future election-related hacking attacks. We have also introduced a suite of digital tools designed to help election websites and political campaigns protect themselves from phishing, unauthorized account access, and other digital attacks.

We have partnered with the National Cyber Security Alliance to fund and advise on security training programs, such as Lock Down Your Login, that focus specifically on elected officials, campaigns, and staff members. We are also increasing our long-standing support for the bipartisan Defending Digital Democracy Project at the Belfer Center for Science and International Affairs at Harvard Kennedy School.

Over the past 18 months, we have undertaken a broad effort to highlight authoritative sources and minimize the spread of misinformation on our platforms. 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 are looking at. Publishers who review third-party claims or rumors can showcase their work on Google News through fact-check labels and in Google Search through fact-check cards. To help ensure Google does not monetize content designed to mislead users, we have implemented a new policy for our AdSense publishers that explicitly bans ads on any site that misrepresents, misstates, or conceals information about the publisher, the publisher's content, or the primary purpose of the site. For Google Search, we updated our Search Quality Rater Guidelines and our evaluation test sets to help identify misleading information and unexpected offensive results, and have used this data to improve our search algorithms. This results in higher quality and more authoritative Search results. We will continue to work on preventing the spread of misinformation by partnering with the journalism industry and civil society to help people understand what they see online and to support the creation of quality content.

And we have recently announced a number of enhancements around Ads transparency:

- **Transparency Report.** In 2018, we'll release a transparency report for Election Ads, which will share data about who is buying election-related ads on our platforms and how much money is being spent.
- **Creative Library.** We'll also introduce a publicly accessible database of Election Ads purchased on AdWords and YouTube (with information about who bought each ad). That means people will not only be able to learn more about who's buying election-related ads on our platforms; they'll be able to see the ads themselves, regardless of to whom they were shown.
- **In-ad disclosures.** Going forward, we'll identify the names of advertisers running election-related campaigns on Search, YouTube, and the Google Display Network.
- **Verification program.** U.S. law restricts entities outside the United States from running election-related ads. We'll reinforce our existing protections by requiring that advertisers proactively identify who they are and where they are based before running any election-related ads. As they do, we'll verify that they are permitted to run U.S. election campaigns through our own checks.

This is a high priority for Google. We are committed to implementing these features before the 2018 election season, and we will continue to explore additional measures to enhance election integrity.

b. **Will you be ready before then to ensure that consumers can quickly identify who is truly responsible for election ads or election-related content that they see on your platform?**

We view this as a high priority and can assure Congress that we are committed to implementing these features before the 2018 election season.

c. **If you cannot provide reassurance that you will be ready before Election Day 2018, what else needs to happen in the next year to provide that reassurance?**

N/A

2. We've heard a lot about the Russian "troll farm" model best exemplified by the Internet Research Agency in St. Petersburg. It is astonishing that we are seeing these types of businesses sprout up for the purpose of spreading disinformation and sowing division online. Your company has taken steps to remove some accounts and ads created by these troll farms, but I fear that a reactive strategy is not going to be good enough.

a. **What additional legislative or administrative actions do you think Congress or federal agencies should pursue against these troll farms to prevent them from spreading lies and discord across the internet?**

We agree this issue raises serious concerns and welcome the opportunity to work with Congress and federal agencies to address the misuse of our platforms and prevent similar issues going forward.

b. **Should there be a special designation or "watch list" set up by the government for troll farms which would carry certain penalties or obligations for companies that fit this designation?**

We welcome information sharing from Congress and law enforcement on these issues, including input from law enforcement and Congress on preventing and detecting this and similar abuses of our platforms.

3.

a. **Is it your view that the federal Departments of Justice and Homeland Security are taking the problem of Russian disinformation over social media seriously?**

We are unable to opine on the efforts of the Departments of Justice and Homeland Security, but do welcome assistance around these issues from those agencies and from Congress, including input on preventing and detecting this and similar abuses of our platforms.

b. **Is your company getting support, guidance and collaboration from those two agencies?**

To date, we have not received recommendations on preventing interference going forward.

We understand the importance of combating foreign disinformation campaigns on our platforms and therefore welcome input from law enforcement and Congress on preventing and detecting any abuses of our platforms.

c. **Who are the point people in those agencies dedicated to working with your company on this challenge?**

Who we work with depends on the investigation, but we take all of these issues seriously and understand the need to work with law enforcement, including the FBI, the Department of Justice, (including U.S. Attorney's Offices), and other agencies, consistent with our policies and the law.

4. Much of the discussion about combatting extremist content on social media has centered around the global terrorism threat. However, we are also facing a rising threat posed by white supremacist and other domestic extremist groups, who are all too often motivated by bigotry and hate.

An unclassified May 2017 FBI-DHS joint intelligence bulletin found that "white supremacist extremism poses [a] persistent threat of lethal violence," and that white supremacists "were responsible for 49 homicides in 26 attacks from 2000 to 2016 … more than any other domestic extremist movement." And *Politico* reported recently that "suspects accused of extreme right-wing violence have accounted for far more attacks in the U.S. than those linked to foreign Islamic groups like al Qaeda and ISIS, according to multiple independent studies."

What steps is your company taking to address extremist content from white supremacists and other domestic terrorist threats?

Our policies prohibit content that promotes or condones violence against individuals or groups based upon race or ethnic origin, religion, disability, gender, age, nationality, veteran status, or sexual orientation/gender identity, or whose primary purpose is inciting hatred on the basis of these core characteristics.

YouTube uses a combination of user flagging and technology to enforce its policies. For example, our global community flags more than 250,000 videos per day, and our enforcement teams around the world review flagged content 24 hours a day, 7 days a week. We also have a number of technological tools driven by machine learning that help identify extremist content and prevent re-uploads of that content. We are committed to further research and development to identify new ways in which technology, machine learning in particular, can help us in the fight against extremism of all kinds online.

Finally, we recently announced the creation of the Global Internet Forum to Counter Terrorism, along with Facebook, Microsoft, and Twitter, to help us continue to make our hosted consumer services more efficient at enforcing our policies and helping to curb the pressing global issue of terrorist content online. Through the forum, we will collaborate with other companies to develop and share the best of technology, policies, and enforcement tactics to tackle these issues.

Back to Top

Extremist Content and Russian Disinformation Online:
Working with Tech to Find Solutions
Richard Salgado, Google
Questions for the Record
Submitted November 7, 2017

QUESTIONS FROM SENATOR FEINSTEIN

Information on Russian-Connected Accounts

1. Accounts and ads created by the Internet Research Agency (IRA) or other Russia-linked entities have been identified to varying degrees by your company.

   a. How do you know whether all accounts tied to the IRA or other suspected Russian-connected entities that are using your platforms have been identified?

   b. Have you found other troll farms or other organizations like the IRA? (If so, please describe those organizations and their use of your social media platform.)

   c. What criteria do you use to identify inauthentic accounts?

   d. What do you do with inauthentic accounts once you've identified them?

We have conducted an extensive review of this issue including developing a list of actors we know or suspect were involved in this effort from our research of publicly available information, the work of our security team, and leads we received from others in the industry, and applying that list to nearly twenty of Google's products, including all of its Ads products. That effort identified limited activity on our platforms, but did identify two Ads accounts with approximately $4700 of spend. As validation of our findings, we broadly reviewed all political ads from June 2015 until the election last November that had even the loosest connection to Russia, which confirmed that we had identified all of the relevant ads we believe to be connected to this effort.

We have removed the actors we suspected to have been involved in this effort from our platforms, suspending their accounts in light of this investigation.

Our investigation is ongoing, we continue to request and receive leads from peers in our industry, and will continue working with Congressional investigations on this topic.

AD Purchasing

2. Did your company have any restrictions before the 2016 election on who could buy ads?

We had several controls in place prior to the 2016 election, and understand the importance of maintaining and enhancing those controls as we go into the 2018 election season. Our existing safeguards included policies designed to prevent foreign nationals from buying U.S. Election Ads. We also already had tight restrictions in place before 2016 limiting which advertisers can serve ads to audiences based on their political leanings. Moving forward, among other enhancements, we will go further by verifying the identity of anyone who wants to run an election ad or use our political-interest-based tools in order to prevent foreign nationals from buying U.S. Election Ads.

3. Is there any way for your company to tell if an ad buyer is a mere intermediary or proxy for someone else? For example, can your company detect when an ad buyer is serving as a proxy for the Russian government or a Russian troll farm that actually paid for the ad campaign?

Our existing safeguards include policies that prohibit foreign nationals from buying U.S. Election Ads. We already tightly restrict which advertisers can serve ads to audiences based on their political leanings. Moving forward, among other enhancements, we will go further by verifying the identity of advertisers who want to run an election ad or use our political-interest-based tools in order to prevent foreign nationals from buying U.S. Election Ads.

4. What are you doing to make sure that you know when foreign state actors buy ads? What are you doing to disclose that fact to other users?

Our policies only permit U.S. entities to buy U.S. Election Ads or use our political affiliation targeting (which we currently limit to only two categories–left-leaning and right-leaning). We have also updated our publisher policies to prohibit ads on sites that misrepresent, misstate, or conceal information about the publisher, the publisher's content, or the primary purpose of the site. And we are committed to transparency. We will publish our first ever transparency report for Election Ads, detailing who is buying election-related ads on our platforms and how much money is being spent. We will also introduce a publicly accessible repository of Election Ads from across our Ads products (with information about who bought each ad). And we'll include information about the identity of advertisers, accessible from the ads themselves.

5. What is your company doing to identify businesses and organizations that run election ads?

We will verify the country of origin of any entity that wants to run an election ad or use our political-interest-based tools by looking at multiple signals such as billing address and currency in order to prevent foreign nationals from buying U.S. Election Ads.

6. Do you believe that other platform users should be notified regarding the identity of individuals or entities purchasing election ads on your platform?

We are committed to transparency and have announced a suite of enhancements to ensure transparency around Election Ads on our systems. As explained above, we will identify advertisers running Election Ads; publish our first ever transparency report for Election Ads, detailing who is buying election-related ads on our platforms and how much money is being spent; and introduce a publicly accessible repository of Election Ads from across our Ads products (with information about who bought each ad). Our commitment to transparency is ongoing and we will continue to explore ways to increase transparency on our platforms.

7. What specific documentation are you using to verify that ad purchasers are who they say they are?

We will verify the country of origin of any entity that wants to run an election ad or use our political-interest-based tools by looking at signals such as billing address and currency.

<u>Shutting Down Suspect Accounts</u>

8. What criteria does your company use to determine if an account should be shut down?

Our policies cover a vast array of problematic behavior and through a combination of sophisticated algorithms and other technologies and human review, we both proactively look for violations of our policies and respond to complaints. We take this work very seriously; in 2016 alone we removed 1.7B ads for violating policies.

9. What steps are being taken to prevent your platforms from being used to incite violence or lawlessness?

Our policies prohibit content that promotes or condones violence against individuals or groups based upon race or ethnic origin, religion, disability, gender, age, nationality, veteran status, or sexual orientation/gender identity, or whose primary purpose is inciting hatred on the basis of these core characteristics. We both proactively look for violations of our policies and respond to complaints.

We use a combination of user flagging and technology to enforce these policies. For example, our global community flags more than 250,000 videos per day, and our enforcement teams around the world review flagged content 24 hours a day, 7 days a week. We also have a number of technological tools driven by machine learning that help identify extremist content and prevent re-uploads of that content.

We are committed to further research and development to identify new ways in which technology, and machine learning in particular, can help us in the fight against extremism of all kinds online.

Finally, we recently announced the creation of the Global Internet Forum to Counter Terrorism, along with Facebook, Microsoft, and Twitter, to help us continue to make our hosted consumer services more efficient at enforcing our policies and helping curb the pressing global issue of extremist content online Through the forum, we will collaborate with other companies to develop and share the best of technology, policies, and enforcement tactics to tackle these issues.

10. Are you considering changes to your terms of service to address this content?

We continually evolve our policies to address new threats, just as new threats continue to evolve. Our policies already prohibit the misuse of our platforms and we've already suspended accounts we believe were associated with this effort pursuant to our existing Terms of Service.

Russian State-Sponsored Media

11. What steps did your company take to evaluate how its platform is being exploited by Russian organizations before and after the Intelligence Community Assessment was released in January 2017?

Protecting our platforms from state-sponsored interference is a challenge we began tackling as a company long before the 2016 presidential election. We've dedicated significant resources to help protect our platforms from such attacks by maintaining cutting-edge defensive systems and by building advanced security tools directly into our consumer products.

With respect to the 2016 election, we have been looking across our products to understand whether individuals who appear to be connected to government-backed entities were disseminating information in the United States for the purpose of interfering with the election. That effort identified limited activity on our platforms, but did identify two Ads accounts with approximately $4,700 of spend. As a validation of our findings, we broadly reviewed all political ads from June 2015 until the election last November that had even the loosest connection to Russia, which confirmed that we had identified all of the relevant ads we believe to be connected to this effort.

While the activity on our platforms was relatively limited, we believe that was in large part due to the controls we had in place prior to the 2016 election, and understand the importance of maintaining and enhancing those controls as we go into the 2018 election season. We will continue to expand our use of cutting-edge technology to protect our users and will continue working with governments to ensure that our platforms are not used misused going into the 2018 election season.

12. How does your identify state-sponsored propaganda? What steps does your company take once such propaganda is identified?

Google enforces policies that prohibit a range of misconduct by those who place content on its platforms, including misrepresenting the owner's origin or purpose, engaging in harassment, or posting hateful, extremist, or violent content.   We are in the process of enhancing the transparency of Election Ads by permitting users to find the name of any advertiser running an election ad on Search, YouTube, and the Google Display Network. We also will be releasing a transparency report for Election Ads, sharing data about who is buying Election Ads on our platforms and how much money is being spent. We will pair our transparency report with a publicly available repository of election ad creatives from across our Ads products. And we will make the database available for public research.

Over the past 18 months, we have also undertaken a broad effort to highlight authoritative sources and minimize the spread of misinformation on our platforms. 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 are looking at. Publishers who review third-party claims or rumors can showcase their work on Google News through fact-check labels and in Google Search through fact-check cards. To help ensure Google does not monetize content designed to mislead users, we have implemented a new policy for our AdSense publishers that explicitly bans ads on sites that misrepresent, misstate, or conceal information about the publisher, the publisher's content, or the primary purpose of the site. For Google Search, we updated our Search Quality Rater Guidelines and our evaluation test sets to help identify misleading information and unexpected offensive results, and have used this data to improve our search algorithms. This results in higher quality and more authoritative Search results. We will continue to work on preventing the spread of misinformation by partnering with the journalism industry and civil society to help people understand what they see online and to support the creation of quality content.

13. How does Google treat content from state-sponsored propaganda accounts or suspect accounts in its News Feed and search results?

As noted above, we have introduced the Fact-Check Label to provide useful context for people as they explore information online. This feature is now available globally in Search and on Google News. We have also introduced new policies against misrepresentative content for AdSense and Ad Exchange publishers and have since taken action against hundreds of publishers. We'll continue to build on these efforts. For example, we are considering ways to provide greater transparency around news sources, including disclosure of government funding.

14. Identify how much money Google has made, directly or via third-party intermediaries, through its relationships with RT, Sputnik, and any other Russian state-run media entities, whether by: (i) selling these entities' ads; (ii) placing ads on these entities' websites or webpages; or (iii) in any other way. Please provide this information broken down by year, Russian entity, and Google product.

The confidentiality and non-disclosure provisions of our contracts prohibit us from publicly disclosing revenue attributable to a particular customer. We are happy to discuss this with you further.

<u>Gmail Accounts Used To Exploit Other Social Media Platforms</u>

15. I understand that Russians set up Gmail accounts so that they could then establish accounts,

using that Google email credential, on other social media platforms like Facebook or Twitter.

a. What efforts has Google taken to identify Gmail accounts that were used for malicious purposes on other social media platforms?

We have conducted an extensive review of this issue including developing a list of actors we know or suspect were involved in this effort from our research of publicly available information, the work of our security team, and leads we received from others in the industry, and applying those leads to nearly twenty of Google's products, including Gmail. While we had relatively limited activity on our platforms, we can confirm that Gmail accounts appear to have been used to open accounts on social media platforms.

b. What does Google do when it learns that a Gmail account that was used for malicious purposes on other company's platforms?

The Gmail accounts associated with the actors who we believe were involved with this effort have been suspended in light of this investigation, per our policies. We have also shared our leads with those other companies to further their investigative and enforcement efforts.

c. Do your companies share intelligence about malicious accounts with each other?

Yes. We do share intelligence about malicious accounts with other companies. Combating disinformation campaigns requires efforts from across the industry. We will continue our long-established policy of routinely sharing threat information with our peers, and work with them to better protect the collective digital ecosystem. We also welcome input from law enforcement and Congress.

Targeting Voters

16. It has been reported that social media companies offered to embed their personnel with the presidential campaigns so that they could make more effective use of your ad buying tools. ("How Facebook, Google and Twitter 'embeds' helped Trump in 2016" Politico, 10/26/17.) For example, your personnel could assist the campaigns in refining their voter targeting to maximize the effectiveness of their ads.

a. What tools did your company offer the campaigns to target voters?

Our employees did not provide voter profiles to allow campaigns to "microtarget" prospective voters. As noted above, we offered limited political affiliation audience targeting (limited only to targeting left-leaning and right-leaning voters). Beyond political affiliation targeting, we did not provide any targeting tools to campaigns that were not generally available to all AdWords advertisers.

b. Did your company's employees provide voter profiles to the campaigns to allow for "microtargeting" of prospective voters?

No. Our company does not collect or provide any information on voting history of users.

c. Did your company's employees provide input on the content of ads to make them more effective?

Our employees do not help build content, but do assist campaigns with a high-level understanding of our Ads tools, just as we do for all other larger customers, and provide examples of ads used in major corporate branding campaigns.

17. We know that Russian operatives used Facebook, Twitter, and Google platforms to build deceptive online presences. We also know that Russia-linked ads targeted U.S. users in various ways, including interests and location.

a. Did the Russia-linked advertisers target people in similar ways – by similar interests, locations, etc. – as the Trump campaign?

The activity on our platforms was limited, and we believe that was at least in part due to the restrictions we had on political ad targeting going into the 2016 election. The two Russian state-sponsored accounts that we identified during the course of our investigation (which spent $4,700 combined) neither utilized our political leanings targeting nor targeted their ads toward "swing states" or any discernible political audience.

White Nationalism and Online Cyberhate

18. During the last Presidential election (from August 2015-July 2016), the Anti-Defamation League found 2.6 million tweets that had anti-Semitic language, with nearly 20,000 tweets directed at 50,000 U. S. journalists. One Jewish reporter received threats over twitter, including a photoshopped picture of her face on a corpse in a concentration camp. (USA Today, "Massive Rise in Hate Speech on Twitter during Presidential Election," 10/21/16.) The photo included a message saying, "Don't mess with our boy Trump, or you will be first in line for the camp." This type of cyberhate has targeted other minority communities as well, including Muslim and immigrant communities.

a. What is your company doing to take down these types of messages and advertisements?

We have specific policies that target this type of offensive content and hateful rhetoric. Pursuant to these policies, we removed over 1.7 billion ads for policy violations in 2016 alone, including violations of our policies around "hate speech", harassment and bullying, and user security and physical safety.

Back to Top

Questions for the Record

Senate Committee on the Judiciary,
Subcommittee on Crime and Terrorism Hearing on
Extremist Content and Russian Disinformation Online:
Working with Tech to Find Solutions

October 31, 2017

QUESTIONS FOR THE RECORD - Chairman Grassley

Richard Salgado, Law Enforcement/Information Security Director, Google

1. To follow up on a request made during the hearing, please provide a detailed written update on what internal investigations have found regarding all accounts, advertisements, and posts with connections to Russia that relate to the lead-up and aftermath of the 2016 presidential campaign.

We have conducted an extensive review of this issue. We developed a list of actors we know or suspect were involved in this effort from our research of publicly available information, the work of our security team, and leads we received from others in the industry. We applied those leads to nearly twenty of Google's products, including all Ads products. We identified limited activity on our platforms, but did identify two Ads accounts with approximately $4,700 of spend. As validation of our findings, we broadly reviewed all political ads from June 2015 until the election last November that had even the loosest connection to Russia, which confirmed that we had identified all of the ads we believe to be connected to this effort. We have removed the actors we suspected to have been involved in this effort from our platform, suspending their accounts in light of this investigation. Although the investigation has continued, we have identified no additional activity.

2. Globally, leaders and law enforcement – including in the United States and in Europe – have been highly critical of Facebook and other tech companies for not doing more to combat extremist content online. In June 2017, Google's General Counsel said, "Terrorism is an attack on open societies, and addressing the threat posed by violence and hate is a critical challenge for us all. Google and YouTube are committed to being part of the solution. We are working with government, law enforcement and civil society groups to tackle the problem of violent extremism online. There should be no place for terrorist content on our services."

    a. How exactly is tech working to help (1) government, (2) law enforcement, and (3) civil society groups tackle the problem of extremist content online?

Please see our response under Item 3 below.

b. Please list all law enforcement agencies, domestic and international, with which tech is partnering or assisting and describe the assistance.

i. If no such partnership or assistance exists, please explain why not?

ii. Regardless of current status, what are future plans in this area?

iii. What specific proposals does tech have to assist, or further assist, law enforcement in this area?

We have long encouraged users and authorities to alert us to content they believe violates the law, and we remove content from search results in response to valid legal requests. We also comply with relevant local laws regarding terrorist or extremist content. Additionally, we have a team that works on responding to government requests 24 hours a day across multiple timezones, and we process thousands of removal requests every year, and share that information in our transparency report: https://transparencyreport.google.com/government-removals/overview. As you will see from that report, we work with various law enforcement agencies each year, including the Department of Justice and Federal Bureau of Investigation. Who we work with depends on the nature of the investigation and jurisdiction.

3. What more should tech be doing to incorporate new and existing technologies to independently, as a matter of corporate responsibility, detect, remove, and report – both to their users and law enforcement – extremist content?

We are committed to getting this right, and continue to develop new tactics to combat extremist content. Our hosted platforms have policies in place that prohibit hate speech, incitement to violence and terrorist recruitment. We use a mix of technology and human review to enforce our guidelines. In June of this year, we announced four steps we are taking to further combat terrorism and hate speech on our platforms.

- First, we've invested more in machine learning technologies for detection and removal of extremist content.
- Second, we are focused on improving and expanding our expertise on these issues through partnerships and collaborations with NGOs.
- Third, we are taking a tougher stance on videos that may be offensive, but do not violate our policies, by placing them behind a warning and disabling comments, recommendations, and other features.
- Finally, we are creating programs to promote counter-speech on our platforms and to redirect people away from violent, extremist propaganda and toward content that counters those narratives.

As described in our transparency report which you can find at https://transparencyreport.google.com/government-removals/overview we respond to over 90,000 requests from law enforcement annually. We collaborate with law enforcement pursuant to valid legal process. And we also have worked with other technology companies to create a shared database of previously identified extremist content to help quickly detect terrorist material and prevent it from spreading across platforms. We are committed to doing our part, and recognize that we must work together across government, civil society, and the private sector to address these complex issues at their root.

4. In March 2017, advertisers left Google's video-sharing platform YouTube after finding that their ads were appearing next to extremist content. Google had allowed some videos to remain on YouTube with a "warning" label — this has not been done with other content such as pornography or copyright violations that defy YouTube's Terms of Service. In June 2017, Google and YouTube introduced new measures to curb extremist video online.

    a. How is tech differentiating between extremist content to be removed and content that needs a warning label?

Our community guidelines prohibit incitement to violence and terrorist recruitment. We also have the ability to place certain videos–that may contain violence or other shocking content–behind a warning interstitial. Additionally, YouTube has the ability to restrict content either by age so that users must be signed-in and over 18 to view it, or by putting it behind a warning interstitial and turning off features like comments and recommendations.

    b. What is tech doing to remove extremist content that violates terms of service?

We use a mix of humans and technology to enforce our policies, including the use of machine learning classifiers to detect new videos that may need to be reviewed and removed under our policies.

    c. What is tech doing to define, catalogue, and categorize content it has removed and share it:

        i. within any family of apps

When we remove a video from YouTube, it is removed from the family of YouTube apps.

        ii. other Internet and social media platforms

In 2016, we, along with Facebook, Twitter, and Microsoft, announced a hash-sharing coalition for terrorist content so that we can exchange digital fingerprints of extremist content identified on our respective services.

        iii. and also with law enforcement?

Google proactively and immediately reports to law enforcement videos that reflect imminent risk of harm, and also reports where there appears to be a historical crime.

5. Last year Google, Facebook, Twitter, and Microsoft announced a "hashing coalition" designed to share signatures of known extremist content and to remove this content.

    a. How large is the signature database?"

We have been building up the database for months through extensive collaboration with the aforementioned companies. Today, the database contains tens of thousands of hashes, and continues to grow on a monthly basis.

    b. How much is the signature database growing on a weekly and on a monthly basis?

The database grew substantially as we built it up in its first year and participating companies contributed their existing hashes; we expect the growth to slow as we contribute on an ongoing basis.

    c. How much content is actually being found and removed?

When a match to an existing hash is found, each company evaluates the content under their respective policies to take appropriate action.

    d. Is the content deployed on all platforms (e.g., for Google, on YouTube, Google Drive, Google Photos)?

We have implemented the hash-sharing coalition on YouTube and are exploring using it for other Google products.

    e. What is being done to share this hashing coalition with law enforcement, both domestically and globally?

While we do not have plans to share the database itself beyond the participants in the coalition, when we see information that poses an immediate threat, we report it to law enforcement. And, of course, when we receive legitimate legal process we work with the requesting entity to provide information that we are able to (based on law and what we have).

6. Tech has promised solutions based on artificial intelligence to identifying the problem of extremist content online.

    a. How large are the current teams working on this technology?

We are striving to incorporate artificial intelligence and machine learning tools into all of our products, services, and processes across the board. Optimizing this technology is truly a company-wide effort.

    b. What is the timeline of deployment?

We have already begun rolling out artificial intelligence technology to help identify extremist and terrorism-related videos on YouTube. We are also applying our advanced machine learning research to train new “content classifiers” to help us remove extremist and terrorism-related content.

c. What is the current accuracy of detection and false alarms?

Our machine learning classifiers have variable levels of precision, but false positives are common, which is why we send content flagged by our machines to human reviewers for evaluation and potential action. Technology has helped us accelerate and scale our removal process. Our most recent data shows that 83% of the violent extremist content we’ve removed was identified proactively by artificial intelligence tools before receiving a human flag.

d. Will these technologies be shared across the tech industry?

We’re already collaborating with our technology and NGO partners to share best practices on addressing violent extremist content online. We are a founding member of the Global Internet Forum to Counter Terrorism along with Facebook, Twitter, Microsoft. This forum works with governments and NGOs to combat spread of terrorist content online.

e. Will these technologies be shared with law enforcement?

We collaborate with law enforcement pursuant to valid legal process.

Back to Top

Senate Judiciary Committee
Subcommittee on Crime and Terrorism
"Extremist Content and Russian Disinformation Online: Working with Tech to Find Solutions"
Questions for the Record
October 31, 2017
Senator Amy Klobuchar

Questions for Richard Salgado, Director of Law Enforcement and Information Security, Google
I want to follow up on our exchange during the hearing regarding the share of advertising revenue that RT earns through its YouTube channel, which – as we discussed – is one of the most popular channels on the platform.

- How much money did RT earn from the advertising revenue generated on its YouTube channel and the RT America channel last year?

The confidentiality and non-disclosure provisions of our contracts prohibit us from publicly disclosing revenue attributable to a particular customer. We are happy to discuss this with you further.

- Is that money paid to RT directly by YouTube or Google? If not, how was that share of the advertising revenue paid to RT?

Revenue earned via monetization of YouTube channels is paid out to creators by Google.

- Do you agree that Google should work to prevent revenue sharing for entities like RT that have facilitated Russia's misinformation campaigns, and is Google taking any measures toward that end?

We are actively working to provide users and advertisers with more information about the content they are seeing to allow them to make educated choices, including whether they advertise on specific sites, such as RT. We have "nutrition labels" on Search describing RT's relationship with the Russian Government and we are currently working on similar disclosures on YouTube. We have contacted RT about its registration under the Foreign Agents Registration Act ("FARA") and specifically FARA's requirement that foreign agents label any informational materials they disseminate. We have requested that RT label informational materials. We allow our advertisers to identify websites where they do not want to place their advertisements. Finally, we want to continue to cooperate with Congress on legislation to combat foreign actors' misinformation campaigns, including potential updates of FARA.

Back to Top

Hearing on Extremist Content and Russian Disinformation Online: Working with Tech to Find Solutions" Judiciary Subcommittee on Crime and Terrorism October 31, 2017

QUESTIONS FOR THE RECORD FROM SENATOR WHITEHOUSE

Mr. Richard Salgado

1) Does Google have an interest in preventing fake news from unverified sources from appearing above news from verified and reliable sources in its search results? If so, what safeguards does it have in place to further this objective?

Google seeks to provide users with links to high-quality, relevant answers to their queries. This is critical to Google's business, because users who obtain low-quality content will begin to take their search queries elsewhere. Fake or misleading content is not helpful to our users and we work on various fronts to rank authoritative content higher in our search results, where feasible, and to show users contextual information that may help them assess the veracity of claims made on pages around the Web.

Our tools do not just protect our physical and network security, but also detect and prevent the artificial boosting of content, spam, and other attempts to manipulate our systems. Among other things:

- On Google News we label links so users can see if the content is locally sourced, an OpEd, or an in-depth piece.
- For Google Search, we have updated our quality guidelines and evaluations to help identify misleading information and thereby helping surface more authoritative content from the web. We have also updated our advertising guidelines to prohibit ads on sites that misrepresent themselves.
- In Search, we seek to provide users contextual information that helps them assess the veracity of a site or understand controversies around its reporting.
- On YouTube we employ a sophisticated spam and security-breach detection system to detect anomalous behavior and catch people trying to inflate view counts of videos or numbers of subscribers.
- We've rolled out features that highlight for users when one site is fact-checking another's reporting, and just this month, launched an improvement to our Knowledge Graph, allowing users to obtain more comprehensive information about news sources, including journalism awards they've won and instances where their content has been fact-checked.

2) *Time* reported in May, 2017 that, while Google saw no evidence of Russian manipulation in its search results, it updated its algorithm "just in case."[2] How did Google determine that Russia-related entities or individuals did not manipulate its search results? What did the update to Google's algorithm entail?

The algorithmic changes referred to by the *Time* story were general in nature, not specifically focused on manipulation.

Google has identified a set of websites associated with the accounts identified during our investigation. Our most experienced webspam fighters reviewed these sites, and found no evidence of manipulation.

3) According to Federal Election Commission records, in 2010 Google received a waiver to the law requiring disclaimers on political advertisements, provided that the ads included a link to the sponsor's site. Given that, since that time, Google's options for advertising have evolved significantly, is it the company's position that it is still entitled to a waiver? If so, under what specific exception?

We are committed to working with the FEC in order to enhance the transparency of digital political advertising. Google did not receive a waiver from the FEC. In a 2010 Advisory Opinion, however, the FEC did confirm that advertisers are not required to include a disclaimer on AdWords because of the size of the ad or impractical nature of including additional language. We are in favor of making election advertising more transparent by implementing the following measures:

- **Transparency Report.** In 2018, we'll release a transparency report for Election Ads, which will share data about who is buying election-related ads on our platforms and how much money is being spent.
- **Creative Library.** We'll also introduce a publicly accessible database of Election Ads purchased on AdWords and YouTube (with information about who bought each ad). That means people will not only be able to learn more about who's buying election-related ads on our platforms; they'll be able to see the ads themselves, regardless of to whom they were shown.
- **In-ad disclosures.** Going forward, we'll identify the names of advertisers running election-related campaigns on Search, YouTube, and the Google Display Network.
- **Verification program.** U.S. law restricts entities outside the United States from running election-related ads. We'll reinforce our existing protections by requiring that advertisers proactively identify who they are and where they are based before running any election-related ads. As they do, we'll verify that they are permitted to run U.S. election campaigns through our own checks.

Even as we take our own steps, we certainly can't do this alone and want to continue working with the FEC and Congress to better protect the collective digital ecosystem.

Back to Top