Written Testimony of Richard Salgado
Director, Law Enforcement and Information Security, Google LLC

House Permanent Select Committee on Intelligence

"Emerging Trends in Online Foreign Influence Operations: Social Media, COVID-19, and Election Security"

June 18, 2020

Chairman Schiff, Ranking Member Nunes, and members of the Committee: Thank you for inviting me to testify today to provide an update on Google's efforts to protect election integrity and prevent platform abuse. I appreciate the opportunity to discuss our work in this space.

My name is Richard Salgado and I am the Director of Law Enforcement and Information Security at Google. In this role, I work with thousands of people across the company to protect the security of our networks and user data. Previously, I served as Senior Counsel in the Computer Crime and Intellectual Property Section at the Department of Justice, focusing on computer network cases involving hacking, illegal computer wiretaps, denial of service attacks, malicious code, and other computer crimes.

Google created its search engine in 1998 with a mission to organize the world's information and make it universally accessible and useful. In the years since, we have remained focused on this mission and the importance of providing greater access to information. This access is fundamental to helping people make sense of the world around them, exercise their own critical thinking, and make informed decisions as citizens. This is especially true during this unprecedented moment in our nation's history. As we cope with a global pandemic and are reminded once again of the injustices and structural challenges that continue to exist in our institutions and society, our responsibility to help people access high-quality, relevant, and useful content is more important than ever.

To this end, our business model is dependent on being a useful and trustworthy source of information for everyone. We have a natural incentive to prevent anyone from interfering with the integrity of our products. This is why we have undertaken a wide range of approaches to prevent platform abuse. In my testimony today, I will provide an update on our efforts to combat election-related interference, and describe in more detail three areas in which we are preventing platform abuse and helping our users:
(i) empowering people with authoritative information; (ii) providing account and site security; and (iii) transparency and accountability on advertising.

An Update on Our Efforts to Combat Election-Related Interference

We take our responsibility to protect our users seriously, and have been working to address disinformation, harassment, and misuse of our platforms for years, predating the 2016 election. As we reported during our last testimony to the Committee in 2017, our investigation found relatively little violative foreign government activity on our platform in connection with the 2016 election.

Entering the 2018 midterms, we remained vigilant and continued to improve our ability to detect and prevent election-related threats. We also continued to share information with industry partners, including those represented in the hearing today, as well as governments and other organizations. As we reported, we saw limited misconduct linked to state-sponsored activity in connection to the midterms. Nevertheless, we've continued to improve and be alert for new threats.

We recently launched a new regular bulletin to provide additional information about our findings. Looking ahead to the November general election, we know that the COVID-19 pandemic, the widespread protests relating to systemic racism, and other significant events can provide fodder for nation-state sponsored divisive disinformation campaigns. We remain steadfast in our commitment to combat such abuse and protect our users.

Empowering People with Authoritative Information

We have been building products for over a decade that provide timely, authoritative information to help voters around the world make decisions that affect their

communities, their cities, their states, and their countries. These efforts can have very practical effects on elections.

For example, in 2018, we helped people in the US learn how to register to vote, find their polling places, and understand how to vote absentee. We intend to continue this effort in 2020. We also provided information about all US congressional candidates on the Search page via Knowledge Panels. On election day, we surfaced election results for US congressional races directly in Search in over 30 languages. We have also partnered with organizations like the Voting Information Project since 2008 to help millions of voters get access to details on where to vote, when to vote, and who will be on their ballots. This project has been an ongoing collaboration with the offices of 46 Secretaries of State to ensure that we are surfacing fresh and authoritative information to our users.

In addition to Search results about election information, we have made voting information freely available through the Google Civic Information API, which has allowed developers to create useful applications with a civic purpose. Over 400 sites have embedded tools built on the Civic Information API; these include sites of candidates, campaigns, government agencies, nonprofits, and others who encourage and make it easier for people to get to the polls.

On YouTube, we prominently surfaced information panels and official candidate YouTube channels in related YouTube search queries to help ensure that voters could easily access authoritative content related to the candidates they were interested in ahead of the US midterm elections. Additionally, we surface information panels indicating if a channel is owned by a news publisher that is “funded in whole or in part" by a government and include a link to the publisher's Wikipedia page. Similarly, for channels that are publicly funded, an information panel will surface that indicates the publisher “is a publicly funded broadcaster.”

Combating Disinformation

Over the past several years, we have continued to improve our proactive and reactive efforts to detect and counter coordinated influence operations. Our approach is based around a framework of three strategies: make quality count in our ranking systems, counteract malicious actors, and give users more context. This work has included designing better ranking algorithms, implementing tougher policies against monetization of misrepresentative content, and making authoritative information more easily accessible and useful to people.

In Search, ranking algorithms are an important tool in our fight against disinformation. Ranking elevates relevant information that our algorithms determine is the most authoritative and trustworthy above information that may be less reliable. Our ranking systems are specifically designed to identify sites with high indicia of expertise, authority, and trustworthiness.

We continue to develop and improve our ranking system to ensure we are delivering on our commitment to surfacing high quality search results. This work involves ongoing rigorous testing and evaluation processes to help us benchmark the quality of our results and make sure these meet a high bar all around the world. For example, in 2019 alone, we ran more than 464,000 experiments with trained external Search Raters and live tests, resulting in more than 3,600 improvements to Search.

Similarly on YouTube, our work focuses on what we call the "The Four Rs of Responsibility" where we work to:

- **Remove** content that violates our policy as quickly as possible. We're always looking to make our policies clearer and more effective, and we consult a wide variety of experts to inform our thinking. We report on our content removals in our quarterly Community Guidelines enforcement report. In the fourth quarter of 2019, we removed 5.8 million videos for violating our guidelines; 90% were first flagged by machines rather than humans and of those, 64.7% had no views.
- **Raise** up authoritative voices when people are looking for breaking news information. Our Breaking and Top News shelves, which feature content from authoritative news channels on our home page and in search results, are available in 40 countries, including the United States. We've also launched information panel features to provide users with additional context about the videos they're watching, including for topics prone to misinformation and for channels that receive government funding.
- **Reduce** the spread of content that brushes right up against our policy line. After making changes in the beginning of 2019 to our recommendations systems, we have seen a 70% drop in non-subscribed watch time from recommendations of this type of content in the US.
- **Reward** trusted, eligible content creators who clear a high bar for what channels can make money on our site.

We also invest in automated tools to tackle a broad set of malicious behaviors, and of course we have people who act as reviewers and help train the tooling. Our policies across Google Search, Google News, YouTube, and our advertising products clearly

outline behaviors that are prohibited, such as attempts by spammers to deceive our ranking systems on Google Search, misrepresentation of one's ownership or primary purpose on Google News and our advertising products, or impersonation of other channels or individuals on YouTube. We make these rules of the road clear to users and content creators, while being mindful not to disclose so much information about our systems and policies as to make it easier for malicious actors to circumvent our defenses.

As we work hard to combat disinformation, we also strive to provide users with easy access to context and a diverse set of perspectives, which are key to providing users with the information they need to form their own views. Our products and services expose users to numerous different sources in response to their searches, allowing them to see diverse perspectives or viewpoints. There is no silver bullet to deal with disinformation, but we will continue to work to get it right, and we rely on a diverse set of tools, strategies, and transparency efforts to achieve our goals.

COVID-19

We also applied these strategies to the worldwide COVID-19 response by developing new resources to connect users to helpful information. These resources include a dedicated experience for COVID-19 on Google Search that provides easy access to authoritative information from government health authorities along with data, news, and locally-relevant information from trustworthy sources, YouTube Information Panels, and homepage promotions around the world, including a global "Do the Five" campaign to remind people to follow five simple practices to help stop the spread of COVID-19. Moreover, we developed a website with resources dedicated to COVID-19 education and prevention.

Similarly, we have worked to remove misinformation that risks individual harm and undermines efforts to reduce infection rates. On YouTube, for example, we have clear policies prohibiting content that promotes medically unsubstantiated cures or treatments. We similarly remove content that disputes the existence or transmission of coronavirus or that promotes conspiracy theories such as claims that symptoms are caused by 5G and not COVID-19. We also worked to reduce recommendations of borderline content or videos that could misinform users in harmful ways.

On Google Ads, we treat the COVID-19 pandemic as a sensitive event, which means that as part of our enforcement of this policy, we do not allow ads to run on our platforms where it appears an advertiser may be exploiting the pandemic. This policy

includes engaging in price gouging or where the ads contain or target certain keywords, regardless of ad or site content. We are currently allowing COVID-19 related ads for selected advertiser segments such as government organizations and others who want to get relevant information out to the public.

These efforts have had tangible results: On YouTube, we've removed over 200,000 videos with dangerous or misleading coronavirus information and since launch, there have been over 200 billion impressions on our information panels for coronavirus related videos and searches. We've also removed over 100 million coronavirus related ads globally.

Providing Account and Site Security

Although the sophistication and determination of malicious actors has expanded the electoral threat landscape, Google offers a broad array of services and tools to help campaigns, candidates, and election officials reduce the likelihood of security breaches. We have devoted significant resources to help them improve their cybersecurity posture in light of existing and emerging threats. We strive to apply security protections automatically without requiring user intervention, but when we need users to take affirmative steps themselves, we offer clear recommendations and actions.

Phishing of accounts remains the single biggest threat to account security. Phishing is a fraudulent practice that tricks users into providing their account credentials and is an ongoing threat that Google takes very seriously. Google's Safe Browsing helps protect more than four billion devices from phishing across the web, by hunting, flagging, and disabling malicious extensions in the Chrome Web Store, helping to block harmful ads, and helping to power Google Play Protect (Google's built-in malware protection for Android). Safe Browsing continues to show millions of warnings about websites it considers dangerous or insecure in multiple browsers (Chrome, Firefox, Safari) and across many different platforms, including iOS and Android.

Our improving technology in this area thwarts many account hijacking efforts from ever reaching the inboxes of users. In addition, we have multiple internal teams, including Google's Threat Analysis Group, that identify malicious actors wherever they originate, prevent their attacks, and share threat information with other companies and law enforcement officials. We routinely provide public updates about these

operations and issue warnings to users when we believe they may be the targets of government-backed phishing attacks.

In 2017, we introduced the Advanced Protection Program (APP), which provides the strongest account protection that Google offers. As part of the Program, we have conducted extensive outreach to campaigns, candidates, and election officials to promote the use of security keys that protect users from more sophisticated and targeted phishing campaigns. The APP is available to anyone, but we believe it will have particular utility for candidates, campaigns, election officials, journalists, and democracy and human rights activists. Earlier this year, Google entered into a partnership with Defending Digital Campaigns (DDC) to provide an unlimited number of free Titan Security key and APP enrollment to eligible federal campaigns and political committees. We also recently unveiled a way for at-risk users to let us know 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 program is based on protections all users get, but we increase our sensitivity to potentially risky events for these accounts, much in the same way a bank might with a credit card while a customer is traveling.

Separately, Google provides Protect Your Election, a site that offers a suite of tools to help campaigns, candidates, and election-related websites protect themselves online. In addition to the Advanced Protection Program, the Protect Your Election initiative includes Project Shield, a free service that we released in 2016. Project Shield is designed to mitigate the risk of distributed denial of service attacks, which inundate sites with traffic in an effort to shut them down. These attacks can make campaign and election websites inaccessible to voters, often at critical junctures (*e.g.*, when voters are looking for poll hours and poll location information on election day). Project Shield defends websites against attacks by filtering out and rejecting attack traffic.

Moreover, we work with companies represented in this hearing and other industry partners, as well as law enforcement including the Department of Justice and Department of Homeland Security, to share threat information in order to help detect and thwart bad actors. Just a few weeks ago, our Threat Analysis Group shared that it had identified phishing attempts from a Chinese group targeting the personal email accounts of Biden campaign staff and an Iranian group targeting the personal email accounts of Trump campaign staff. Although we didn't see evidence that these attempts were successful, we sent the targeted users our standard government-backed attack warning, and referred the matter to federal law

enforcement. This highlights the need for at risk-users to take steps to protect their online accounts.

In this regard, Google has supported significant outreach to increase security for candidates and campaigns across the United States and other countries. In the leadup to the 2018 election, we provided trainings on email and campaign website security to over 1,000 campaign professionals and the eight major Republican and Democratic committees. We are continuing this effort for the 2020 elections.

Improving Transparency and Accountability on Advertising

Google has been working hard to make election advertising more transparent. In 2017 we committed to making improvements in this important area and delivered on our commitment by rolling out a verification program for advertisers purchasing U.S. federal election ads, building in-ad disclosures for those buying these ads, and launching a transparency report with a searchable ad library.

Given recent concerns and debates about political advertising, and the importance of shared trust in the democratic process, in November 2019 we announced that we would only allow targeting based on general geographic location, age, gender, and context for election ads in the US. This step aligned our approach to election ads with long-established practices in media such as TV, radio, and print, and resulted in election ads being more widely seen and available for public discussion. In addition, we expanded the coverage of our election ads transparency to include U.S. state-level candidates and officeholders, ballot measures, and ads that mention federal or state political parties. We also clarified our ads policies and added examples to show how our policies prohibit things like “deep fakes” (doctored and manipulated media), misleading claims about the census process, and ads or destinations making demonstrably false claims that could significantly undermine participation or trust in an electoral or democratic process.

Further demonstrating our commitment regarding advertising transparency, this April we announced that we will extend identity verification to all advertisers on our platforms. As part of this initiative, advertisers will be required to complete a verification program in order to buy ads on our network. We have already started with verifying the first wave of advertisers in the US, and you will start seeing in-ads disclosures this summer.

We are thinking hard about elections and how we continue to support democratic processes around the world, including by bringing more transparency to election advertising online, helping connect people to useful and relevant election-related information, and working to protect election information online.

Conclusion

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. Preventing platform abuse, combating disinformation, and protecting elections requires concerted effort and collaboration across the industry. We'll continue to work with tech partners, law enforcement, and others to better protect the collective digital ecosystem and, even as we take our own steps, we are open to working with governments on legislation that promotes electoral transparency.

Thank you for the opportunity to discuss these issues.