January 8, 2018

VIA EMAIL

Chairman Richard Burr
U.S. Senate Select Committee on Intelligence
211 Hart Senate Office Building
Washington, D.C. 20510

Dear Chairman Burr, Vice Chairman Warner, and Members of the Committee:

Thank you for your questions for the record from the November 1, 2017 Committee Hearing on Social Media Influence in the 2016 U.S. Elections. Per your request, attached are the answers for the record to your questions, as well as our hearing testimony.

Please let us know if you have any further questions.

Regards,

Colin Stretch
General Counsel, Facebook

facebook

Address: 1601 Willow Road
Menlo Park, CA 94025

Chairman Burr

1. What procedures must the Russian government follow to compel the production of customer-created content or personally identifiable information from your company?

As part of official investigations, government officials sometimes request data about people who use Facebook. We have strict processes in place to handle these government requests, and we disclose account records solely in accordance with our terms of service and applicable law. We require officials to provide a detailed description of the legal and factual basis for their request, and we push back if the request appears to be legally deficient or is overly broad, vague, or otherwise inconsistent with our policies. Further, with respect to government requests for disclosure from outside the United States, a Mutual Legal Assistance Treaty request or letter rogatory may be required to compel the disclosure of the contents of an account. Facebook has never produced customer-created content or personally identifiable information in response to a request from the Russian government.

2. Has the Russian government compelled the production of customer created content or personally identifiable information from your company?

As part of our ongoing effort to share information about the requests we have received from governments around the world, Facebook regularly produces a Transparency Report about government requests to Facebook. Our Transparency Report contains historical information about Russian requests for data going back to 2013. In summary, we received 28 requests from the Russian government between 2013 and 2017. We did not provide any data in response to these requests.
See https://transparency.facebook.com/country/Russia/2017-H1/.

3. If so, has your company complied with such efforts by the Russian government to compel the production of customer-created content or personally identifiable information?

See response to question 2.

4. Has your company ever refused to comply with efforts by the Russian government to compel the production of customer-created content or personally identifiable information? If so, have any of these efforts been successful?

See response to question 2.

5. Has your company provided any content created by a U.S. person or personally identifiable information about a U.S. person to the Russian government?

See response to question 2.

6. More specifically, has your company provided to the Russian government the content of any direct messages sent to or from a U.S. person?

See response to question 2.

7. Has your company provided to the Russian government any information that could be used to determine the location of a U.S. person?

See response to question 2.

Vice Chairman Warner

1. Facebook took some action to curb activity on tens of thousands of nefarious accounts in the lead-up to this year's French election.

a. What were the lessons learned by Facebook and the French election authorities in the aftermath of that election, where Russian propaganda did not appear to be particularly effective?

While we cannot speak to lessons learned by French election authorities, we are encouraged that improvements to our fake account detection systems allowed us to disable more than 30,000 additional accounts in connection with the French election—though these accounts were most commonly used for financially-motivated spam, not organized propaganda campaigns from a particular country. In addition, after reports of foreign interference in the run-up to the French—and U.S.—elections, we worked closely with German officials on a number of initiatives to fight disinformation and make Facebook a safer and more secure environment for genuine civic engagement. As with all security threats, we have been continuously incorporating new insights into our models for detecting fake accounts, including information specific to fake accounts focused on social, political, or election issues. We believe that we were more effective at taking down fake accounts connected to the French and German elections as a result.

2. Does Facebook intend to notify unwitting American users who shared, liked, or commented on content published by Russian-backed users? For example, informing users that they followed or interacted with a Facebook community that was actually a Russian front. Why or why not?

We recently launched a portal to enable people on Facebook to learn which of the IRA Facebook Pages or Instagram accounts they may have liked or followed between January 2015 and August 2017. This tool is available in the Facebook Help Center at https://www.facebook.com/actionplan.

In the coming weeks, we will take significant steps to make users aware of this new tool, including a notification in their News Feed. We will be promoting this tool alongside our efforts to educate people about changes that we are making to strengthen our platform's resistance to inauthentic and malicious behavior, including increasing ads transparency, implementing a more robust ads review process, imposing tighter content restrictions, and exploring how to add additional authenticity safeguards. These efforts complement ongoing work to help promote news literacy by giving people information so they can make smart choices about the news they read. This work includes partnerships with organizations like the News Literacy Project, support for the News Integrity Initiative, and collaboration with academics and researchers around the world.

We have also worked to notify people about this issue, broadly, through our white paper in April 2017, *Information Operations on Facebook*, and our disclosures this past fall.

3. What are your policies regarding the hosting and distribution of hacked and stolen emails from your platform?

We prohibit any content that is claimed or confirmed to have come from a hacked source. In rare situations and on a case-by-case basis, we may choose to allow content that is newsworthy, significant, or important to the public interest even if it otherwise violates our policies.

4. Reports suggest that in 2014 and 2015, Facebook was alerted by Ukranian activists and politicians—and reportedly even Ukraine's President—about an active campaign by Russian trolls to push disinformation on Facebook.

a. What actions did you take in response to address the threat at the time?

Those reports are incorrect to the extent that they imply that we were warned of information operations similar to those that we have described to this Committee in connection with the 2016 election.

In Russia and Ukraine, like in many countries, we've seen people use slang to refer to people of other countries. For example, some people in Ukraine refer to Russians as "moskal" (literally "Muscovites") and some in Russia call Ukrainians "khokhol" (literally "topknot"). After conflict started in the region in 2014, people in both countries started to report posts containing these words as hate speech. We conducted a review and concluded that these words were indeed being used in some cases as hate speech that violates our Community Standards. We removed reported content, a decision that was initially unpopular on both sides. Anyone can report content to us if they think it violates our policies, and one report is enough for us to remove something. Multiple reports will not lead to the removal of content if it meets our standards.

Senator Collins

1. **Several independent researchers have said that Facebook has the ability to search for content or metadata that could substantiate or disprove allegations of possible collusion between the Russian disinformation operation and the Trump campaign's own social media efforts, such as timing of certain posts and sharing of content.**

    a. **Is this true, and if so, has Facebook found any information relevant to these allegations?**

Facebook does not believe it is in a position to substantiate or disprove allegations of possible collusion. Facebook is, however, providing investigators, including this Committee, with information it has regarding the scope and nature of Russian information operations on our platform so that those investigators have information that may be relevant to their inquiries. We are happy to schedule a meeting with your staff to discuss our findings in more detail.

2. **Facebook has an automated engine for recommending content to users.**

    a. **Did this recommendation engine suggest to any Facebook user that they view, follow, or join any of the Russian-linked pages?**

This happened in some cases. Because we were not aware that these Pages were not legitimate, they were sometimes recommended when people followed similar Pages, for example. However, these recommendations were not the primary way that the Pages attracted their audience.

3. **What provision in your Terms of Service ensures that political advertisements targeted towards the United States are purchased by an American citizen?**

Facebook's Statement of Rights and Responsibilities (the terms that govern all use of our services) prohibit using Facebook to do anything that is unlawful, misleading, or malicious. In addition, advertisers must comply with Facebook's Advertising Policies, including acknowledging that they are responsible for understanding and complying with all applicable laws and regulations. Therefore, violating the Federal Election Campaign Act also violates our terms.

We also support efforts to promote greater transparency in political advertising online and recently announced steps to make advertising on Facebook more transparent, increase requirements for authenticity, and strengthen our enforcement against ads that violate our policies. *See* https://newsroom.fb.com/news/2017/10/update-on-our-advertising-transparency-and-authenticity-efforts/.

4. **Do your Terms of Service prohibit users from influencing elections in other countries?**

Facebook requires advertisers and users to comply with all applicable laws and regulations, including laws related to electoral activity. We also require advertisers and users to comply with our policies, which require that people use their authentic name and identity. Our advertising policy also prohibits misleading or false content, stating: "Ads, landing pages, and business practices must not contain deceptive, false, or misleading content, including deceptive claims, offers, or methods."

5. If a foreign national working on behalf of a foreign intelligence service was an authentic user in real name on your platform, could he post divisive, but non-violent content related to a U.S. election without violating your Terms of Service? Would he be able to purchase political advertising?

If the person was engaged in coordinated inauthentic behavior on behalf of a foreign intelligence service, that would violate our terms even if he or she was using his or her real name. In addition, our terms prohibit ads that are illegal, so if the advertising was prohibited by the Federal Election Campaign Act or any other law, that would also violate our terms.

Senator Feinstein

1. Facebook has conceded that the number of people exposed to content from foreign groups online is far more pronounced through organic traffic and fake accounts than it is through paid advertising. Troublingly, it does not appear there is a proven method for combating the spread of fake accounts created to sow division in society. Although Facebook has indicated authenticity activity measures are in development, as recently as August 2017, divisive foreign unpaid content designed to polarize and anger the American people could be found on Facebook.

a. What specific actions is Facebook taking to combat this type of divisive, unpaid activity on an on-going basis?

We are constantly improving our technical systems to identify and remove inauthentic accounts and reduce the distribution of material that can be spread by accounts that violate our policies, and we do have a track record of success and improvement in this area. Each day, we block millions of fake accounts at registration, as our systems examine thousands of account attributes and focus on detecting behaviors that are difficult for bad actors to mask or fake, including their connections to others on our platform. For example, we are encouraged that recent improvements to our fake account detection systems focused on social and political content allowed us to disable more than 30,000 additional accounts in connection with the French election. However, these accounts were most commonly used for financially-motivated spam, not organized propaganda campaigns from a particular country.

Facebook also removes hate speech, which includes content that directly attacks people based on their protected characteristics. We don't allow any organizations or individuals that are engaged in terrorist activity, organized violence or criminal activity, or organized hate groups to have a presence on Facebook. We also remove content that expresses support for groups that are involved in violent or criminal behavior. And, we remove credible threats of physical harm to individuals and specific threats of theft, vandalism, or other financial harm.

We have deployed a variety of tools in this fight to find and remove bad content, including artificial intelligence, specialized human review, and industry cooperation, as well as supporting important corrective measures such as counter-speech training. We are more than doubling the number of people who work on safety and security at Facebook.

2. One of the more troubling findings from this investigation is the number of targeted voter disengagement efforts promoted through social media.

a. Can you say with certainty that foreign actors did not use the U.S. voter registration data to target individuals through both paid and unpaid activity?

Facebook is not in a position to know everything that foreign actors did in their online activities, but to date we have uncovered no evidence that the IRA used U.S. voter

registration data for ad targeting on Facebook. The targeting for the IRA ads that we have identified and provided to the Committee was relatively rudimentary, targeting broad locations and interests, and did not use a tool known as Contact List Custom Audiences. That tool (which can be used for ad targeting, but not organic, unpaid content) is the one that a candidate's campaign might use (in connection with their own vendors) to develop customized targeting based on voter profile information. More information about Facebook's targeting options in general is publicly available on our website at https://www.facebook.com/business/products/ads/adtargeting.

3. Between June 2015 and August, an estimated 126 million Americans were exposed to Facebook content generated by Russian troll farms.

a. Are you keeping a database of accounts exposed?

We recently launched a portal to enable people on Facebook to learn which of the IRA Facebook Pages or Instagram accounts they may have liked or followed between January 2015 and August 2017. This tool is available in the Facebook Help Center at https://www.facebook.com/actionplan.

In the coming weeks, we will take significant steps to make users aware of this new tool. We will be promoting this tool alongside our efforts to educate people about changes that we are making to strengthen our platform's resistance to inauthentic and malicious behavior, including increasing ads transparency, implementing a more robust ads review process, imposing tighter content restrictions, and exploring how to add additional authenticity safeguards.

4. Facebook confirmed in the House Intelligence committee hearing that they found no overlap in the groups targeted by the Trump campaign's advertisements, and the advertisements tied to the Russia-linked accounts identified thus far.

a. Does this targeting assessment extend to the content used by the Trump campaign and the Russia-related accounts?

b. Does this assessment extend to both the content used and groups targeted by the companies associated with the campaign—like Cambridge Analytica—and Russian accounts?

We have seen only what appears to be insignificant overlap between the targeting and content used by the IRA and that used by the Trump campaign (including its third-party vendors). We are happy to schedule a meeting with your staff to discuss our findings in more detail.

Senator Cotton

1. Do Facebook's Terms of Service prohibit collaboration with Russian intelligence services intended to influence a U.S. election?

Such conduct would violate our policies requiring users not to use Facebook to do anything unlawful, misleading, or malicious, and prohibiting the use of inauthentic accounts. We have processes designed to identify inauthentic and suspicious activity, and we also maintain a sanctions compliance program to screen advertisers.

That said, the challenge of attributing online activities to specific governments or organizations is widely recognized in the intelligence and law enforcement communities. It is extremely difficult to definitively attribute online activity to particular threat actors, and we often rely on information from others, like information included in the January 6, 2017 DNI report, to identify actors behind abuse that we observe and to better understand these issues.

2. Provided an individual or entity does not violate Facebook's Terms of Service, will they be allowed to use your platform to work with hostile, foreign intelligence services to potentially influence the 2018 and 2020 U.S. elections?

We do not believe an individual or entity could engage in this kind of activity without violating our policies—including our policy prohibiting coordinated inauthentic activity – or the law, and we are working to improve detection and enforcement in this area. We are also making significant investments across all of our safety and security teams, which means that we will have more people dedicated to finding this type of abuse among many others.

We support efforts to promote greater transparency in political advertising online and recently announced steps to make advertising on Facebook more transparent, increase requirements for authenticity, and strengthen our enforcement against ads that violate our policies. *See* https://newsroom.fb.com/news/2017/10/update-on-our-advertising-transparency-and-authenticity-efforts/. We'll require more thorough documentation from advertisers who want to run election-related ads. We are starting with federal elections in the U.S., and will progress from there to additional contests and elections in other countries and jurisdictions. As part of the documentation process, advertisers may be required to identify that they are running election-related advertising and verify both their entity and location. Once verified, these advertisers will have to include a disclosure in their election-related ads, which reads: "Paid for by." When users click on the disclosure, they will be able to see details about the advertiser, and we will maintain a searchable archive of information. Like other ads on Facebook, they will also be able to see an explanation of why they saw that particular ad. For political advertisers that do not proactively disclose themselves, we are building machine learning tools that will help us find them and require them to verify their identity.

3. What is Facebook's justification for allowing entities and individuals such as WikiLeaks, Julian Assange, and Edward Snowden to maintain Facebook pages?

Facebook is a place where people are empowered to communicate, and we make our services available to everyone who complies with our policies. We take seriously our role in keeping abuse off our services. These individuals and organizations can maintain a presence on Facebook as long as they comply with our policies. We take action on activity that violates these policies, including blocking the accounts of repeat offenders. In addition, when governments believe that something on Facebook violates their laws, they may ask us to restrict access to that content. We scrutinize these requests, and if we determine the specified content does indeed violate local laws, we may make it unavailable in the relevant country or territory.

Senator Heinrich

1. What percent of Facebook content reviews are conducted by an actual human being rather than via automated review?

We don't have an either-or approach to reviewing content. All content goes through some degree of automated review, and we use human reviewers to check some content that has been flagged by that automated review or reported by people that use Facebook. We also use human reviewers to perform reviews of content that was not flagged or reported to check the accuracy and efficiency of our automated review systems. The percentage of content that is reviewed by a human varies widely depending on the type and context of the content, and we don't target a specific percentage across all content on Facebook.

2. Are Facebook's content review processes the same now as they were during the 2016 election? If not, how have they changed?

Our content review processes are fundamentally the same, but we are more than doubling the number of people who work on safety and security at Facebook and have already hired thousands more content reviewers. They will be engaged in processes that we are continuously refining, but this significant investment of resources will help us to perform those processes more accurately, quickly, and thoroughly. One improvement that we believe will help to address more subtle kinds of abuse is that our ad review team will do more to assess not just the content, but also the overall context of an ad, including the buyer and intended audience. We will also significantly expand the number of people who work specifically on election integrity before the 2018 U.S. federal elections this fall, including people who investigate this specific kind of abuse by foreign actors. Additionally, we have begun testing a program where people will be able to click "View Ads" on a Page and view advertisements a Page is running on Facebook, Instagram and Messenger—whether or not the person viewing it is in the intended target audience for the ad. All Pages will be part of this effort, and we will require that all ads be associated with a Page as part of the ad creation process. Finally, we continue to make improvements to our efforts to more effectively detect and deactivate fake accounts to help reduce the spread of spam, false news, and misinformation. We continually update our technical systems to identify, checkpoint, and remove inauthentic accounts, and we block millions of attempts to register fake accounts every day. These systems examine thousands of detailed account attributes and prioritize signals that are more difficult for bad actors to disguise, such as their connections to others on our platform. As with all security threats, we have been incorporating new insights into our models for detecting fake accounts, including information specific to election issues.

We are determined to do everything that we can to protect our platform. The investments that we are making to address these issues and other security issues will be so significant that we have informed investors that we expect that the amount that we will spend will impact our profitability.

3. How does accountability work when an algorithm makes bad recommendations? For example, as recently as mid-September, anybody

could target ads towards anti-Semites using terms provided by Facebook such as "Jew Hater" and "History of 'why Jews ruin the world.'" Was anybody held accountable for this?

In September 2017, we temporarily disabled some of our ad tools following news reports that slurs or other offensive language could be used as targeting criteria for advertising. In order to allow businesses—especially small ones—to find customers who might be interested in their specific products or services, we offered them the ability to target profile field categories like education and employer. So, if someone on Facebook self-identified as a "Jew-hater" or said that they studied "how to burn Jews" in their education or employer fields on their profile, those terms showed up as potential targeting options for advertisers. These deeply offensive terms were used so infrequently in these write-in fields that we did not discover this until a journalist brought it to our attention.

We have long prohibited hate on Facebook, and although we are not aware of instances in which these terms were ever used to actually target ads, we take our failure to enforce that policy with adequate caution and care in this instance extremely seriously. We never intended or anticipated that this functionality would be used this way, and we did not find it ourselves. We are accountable for these failures. We have tried to learn everything that we can from this painful incident so that we can do better in the future. We have tightened our ad policies and have taken steps to improve our enforcement, including by adding more oversight of our automated review processes, and have been exploring how best to implement tools for people to tell us when our systems may inadvertently enable abuse. We have also used human reviewers to manually check existing targeting options and reinstate the roughly 5,000 most commonly used targeting terms—terms like "nurse" or "dentistry"—to ensure that they meet our Community Standards. We will do more manual review of new targeting options going forward to help prevent offensive terms from appearing.

Targeted advertising on Facebook has helped millions of businesses grow, find customers, and hire people. Our systems match organizations with potential customers who may be interested in their products or services. The systems have been particularly powerful for small businesses, who can use tools that previously were only available to advertisers with large budgets or sophisticated marketing teams. Our ads help make meaningful connections between business and people, and the improvements we are making to our ad policies will help us do this more effectively.

4. In hiring more content reviewers, are your companies simply throwing bodies at a specific problem, or are you fundamentally rethinking how to prioritize which user interactions require additional human oversight and review. If so, how? What other changes have you made in this regard?

We work continuously to make our platform safer and more secure, and our effort to do so is a holistic one that involves not only hiring additional employees when issues arise, but also a continual evaluation of our processes and policies. We rely on both automated and manual ad review, and we're now investing in and taking steps to strengthen both. Reviewing ads means assessing not just what's in an ad but also the context in which it

was bought and the intended audience—so we're changing our ads review system to pay more attention to these signals. We're also adding more than 1,000 people to our teams dedicated to reviewing ads around the globe, and their work will be used to train our automated review systems to be more efficient and effective at finding improper ads and enforcing our policies. Enforcement is never perfect, but we will get better at finding and removing improper ads.

Senator Manchin

1. Does Facebook or any Facebook affiliate use information security products or services of Kaspersky Lab or any Kaspersky Lab affiliate?

We work with anti-malware companies to make available free malware cleanup software to Facebook users when Facebook detects that the user is accessing Facebook from a device that may be infected with malware. In October 2017, we removed Kaspersky's anti-virus software from the list of products that we make available to these users.

We are also in the process of phasing out internal use of a different Kaspersky anti-virus product. That product does not transmit data back to Kaspersky, and is not subject to the security and privacy concerns that have been identified in recent months. We continue to use a Kaspersky service that provides us with information about threat activity as a one-way feed.

2. How many Facebook subscribers took advantage of your offer to provide Kaspersky Lab malware detection products to clean up their computer systems?

We no longer make available Kaspersky's anti-virus software to people with infected devices. Unfortunately, we are unable to easily reconstruct how many Facebook users downloaded Kaspersky software.

3. Does Facebook or any Facebook affiliate sell network space to RT or Sputnik news agencies?

RT and Sputnik can maintain Pages on Facebook and use our advertising tools as long as they comply with Facebook's policies, including complying with applicable law.

4. If you recently terminated any agreements with RT or Sputnik, on what date did the termination become effective?

Not applicable.

5. Do either RT or Sputnik need to purchase advertising space on your platform, or can they freely maintain a Page or distribute web content via their own or affiliated Facebook accounts?

RT and Sputnik can maintain Pages on Facebook for free, as long as they comply with Facebook's policies, including complying with applicable laws.

6. Does Facebook prohibit, or have any concern about, foreign state-sponsored news organizations positing content via the Facebook platform?

Many media organizations, including state-sponsored media organizations, have a presence on Facebook and use our advertising services to promote their content. Like everyone on Facebook, these organizations are required to comply with Facebook's

policies, including complying with applicable laws. They must also comply with our policy requiring them to use their authentic names and identities. And, no entities that advertise on Facebook, including media organizations, can use deceptive, false, or misleading content in their ads. If we become aware that our policies are being violated, we will take action.

Senator Harris

1. Facebook has produced information about Russian propaganda advertisements. Your company has also produced information about Russian propaganda that appeared as ordinary user content. You have not, however, provided information about the legitimate advertisements that accompanied Russian content.

a. How long do you retain placement and billing records for advertisements on your services?

Facebook generally retains core billing and advertisement records for as long as it is commercially reasonable and necessary for our business.

b. Have you instructed relevant business units to retain these records of advertisements that accompanies Russian propaganda? If you have not, will you immediately issue that instruction?

Facebook has taken appropriate steps to retain relevant information related to IRA activity on Facebook.

c. How much revenue do you estimate that you earned from the advertising that accompanied Russian propaganda?

Ads generally did not run on IRA Pages, and we expect that any revenue from such ads would be immaterial. Ads that appear in News Feed are not connected to or endorsed by other pieces of content in an individual's News Feed.

d. Have you notified the advertisers whose advertisements accompanied Russian propaganda?

Facebook has notified all its users and vendors via publicly accessible postings about general activity associated with the IRA. Ads generally did not run on IRA Pages, and ads that appear in News Feed are not connected to or endorsed by other pieces of content in an individual's News Feed.

e. What do you plan to do with the revenue that you earned from the advertisements that accompanied Russian propaganda?

Ads generally did not run on IRA Pages, and we expect that any revenue from such ads would be immaterial. Ads that appear in News Feed are not connected to or endorsed by other pieces of content in an individual's News Feed. We have contributed hundreds of thousands of dollars to the Defending Digital Democracy Project, which is a bipartisan project launched by the Harvard Kennedy School that works to secure democracies around the world from external influence. In addition, the investments that we are making to address election integrity and other security issues will be so significant that we have informed investors that we expect that the amount that we will spend will impact our profitability.

2. The problems of inauthentic, false, and hyper-partisan content are much broader than Russian propaganda.

a. How many of the accounts on your service do you estimate are inauthentic?

We regularly evaluate metrics to estimate the number of "false" accounts among our monthly active users. We divide "false" accounts into two categories. The first category includes user-misclassified accounts, where users have created personal profiles for a business, organization, or non-human entity such as a pet. Such entities are permitted on Facebook using a Page rather than a personal profile under our terms of service. The second category includes undesirable accounts, which represent user profiles that we determine are intended to be used for purposes that violate our terms of service, such as spamming. We estimate that in the third quarter of 2017, user-misclassified and undesirable accounts may have represented approximately two to three percent of our worldwide monthly active users. Our estimation of false accounts can vary as a result of episodic spikes in the creation of such accounts. Additional information relating to our estimate of false accounts is included in our quarterly filings with the Securities and Exchange Commission.

b. How much of the activity on your service do you estimate is inauthentic or false?

See response to question 2(a).

c. How much of your annual revenue do you estimate is attributable to inauthentic or false content?

We believe that annual revenue that is attributable to inauthentic or false accounts is immaterial.

d. Do you have a policy of notifying advertisers when their advertisements accompany inauthentic or false content?

Ads do not run on Pages, so there is no advertising that "accompanies" Pages on Facebook. Ads that appear in News Feed are not connected to or endorsed by other pieces of content in an individual's News Feed.

e. What do you do with the revenue that you earn from advertisements that accompany inauthentic or false content?

See response to question 2(d).

f. If you are aware of independent estimates of inauthentic or false content on your platform, please provide those estimates. If you disagree with the estimates, please explain why.

We estimate that in the third quarter of 2017, user-misclassified and undesirable accounts may have represented approximately two to three percent of our worldwide monthly active users. The estimates of these accounts are based on an internal review of a sample of accounts, and we apply significant judgment in making this determination. Our estimates may change as our methodologies evolve, including through the application of new data signals or technologies, which may allow us to identify previously undetected duplicate or false accounts and may improve our ability to evaluate a broader population of our users.

g. If the independent estimates were accurate, how much of your annual revenue would be attributable to inauthentic or false content?

Not applicable.

h. How much of the news content that is shared on your services do you estimate is false?

We are taking steps to better understand the prevalence of content people find to be wrong or misleading based on the sources they trust. We must better understand if people trust that the information they see on Facebook is credible, and one of the areas we want to better understand is the broader category of information that people find wrong or misleading. Defining what is false news content is difficult and controversial. While it is difficult to estimate a formal percentage, we believe that only a very small amount of content on Facebook is false news.

i. How much of the news content that is shared on your services do you estimate is hyper-partisan?

Defining what is hyper-partisan is difficult and controversial, and we do not have an estimate.

j. Have you conducted any studies of how false content performs on your services? If yes, please describe those studies and provide copies.

As part of our efforts to reduce the spread of false news and misinformation, we routinely track and analyze how such content is disseminated on our platform based on various signals. We do not have formal written studies or reports.

k. Have you conducted any studies of how hyper-partisan content performs on your services? If yes, please describe those studies and provide copies.

We are aware of the concern that our platform may contribute to polarization, and we have been working to understand the role that we play in discourse and information diversity. Defining what is hyper-partisan is difficult and controversial, and we have not conducted any formal studies on how such content spreads on the platform.

3. In the area of state-sponsored hacking, each of your companies has a responsible senior executive and dedicated technical experts.

a. Who is the senior executive responsible for countering state-sponsored information operations? When did that executive assume that responsibility, and what is the scope of the responsibility?

Our Chief Security Officer manages a threat intelligence team that is acutely focused on state-sponsored information operations. He reports to Colin Stretch, Facebook's General Counsel, who is ultimately responsible for Facebook's overall response to this threat.

b. As of November 2016, how many of your technical employees had the primary day-to-day task of countering state-sponsored information operations?

In 2014, we stood up a threat intelligence team dedicated to reviewing and monitoring for, among other things, attacks from threat actors tied to nation states. That work was mostly directed at traditional cybersecurity, such as account compromise, surveillance, and the dissemination of stolen information, but we did have some people addressing state-sponsored information operations. We are expanding our threat intelligence team, and more broadly, we are working now to ensure that we will more than double the number of people working on safety and security at Facebook, from 10,000 to 20,000, by the end of 2018. Many of the people we are adding to these efforts will join our ad review team, and we also expect to add at least 3,000 people to Community Operations, which reviews content that our users and automated tools flag as inappropriate, dangerous, abusive, or otherwise violating our policies. These investments will help us to enforce our policies, including our authenticity policy, and help us to counter threats from malicious actors, including those who are state-sponsored. We will also significantly expand the number of people who work specifically on election integrity before the 2018 U.S. federal elections this fall, including people who investigate information operations by foreign actors.

c. As of today, how many of your technical employees have the primary day-to-day task of countering state-sponsored information operations?

Many people at Facebook have a role in countering state-sponsored information operations and other threats. We are working now to ensure that we will more than double the number of people working on safety and security at Facebook, from 10,000 to 20,000, by the end of 2018. Many of the people we are adding to these efforts will join our ad review team, to assess not just the content but the entire context of an election ad, including the buyer and the intended audience. We expect to add at least 3,000 people to Community Operations, which reviews content that our users and automated tools flag as inappropriate, dangerous, abusive, or otherwise violating our policies. These investments will help us to enforce our policies, to identify new kinds of abuse on our platform, and to respond quickly to reports from our community and from law enforcement. We will also significantly expand the number of people who work specifically on election integrity

before the 2018 U.S. federal elections this fall, including people who investigate information operations by foreign actors.

4. Much of what we know about Russian propaganda is because of academic researchers and investigative journalists. These groups do not currently have access to the data that they need to inform the public and to build tools for detecting state-sponsored information operations. For example, these groups generally cannot assess the full set of public user activity associated with state-sponsored information operations. Providing access to this data need not come at the expense of user privacy, since these groups could be bound by non-disclosure agreements and use privacy-preserving algorithms to conduct their studies.

a. Will you commit to, by the end of the year, providing five or more independent non-profit entities with access to the data they need to understand and counter state-sponsored information operations? If you will, please provide specifics and a timeline for how you plan to honor the commitment. If you will not, please explain why.

Information operations can affect the entire information ecosystem, from individual consumers of information and political parties to governments, civil society organizations, and media companies. We recognize that an effective response requires a whole-of-society approach and collaboration on matters of security, education, governance, and media literacy. We are committed not only to addressing information and other threats that directly involve our platform, but also supporting the efforts of others to understand and counter these threats. That's why we are supporting the Defending Digital Democracy Project, which is a bipartisan project launched by the Harvard Kennedy School that works to secure democracies around the world from external influence. We are the founding sponsor of Defending Digital Democracy's first big project: building an information sharing and analysis organization with members in various critical areas of the democratic process. We would need to evaluate any request for data on a case-by-case basis that would include understanding the information security practices and capabilities of a third-party organization.

5. Similarly, much of what we know we now know about inauthentic, false, or hyper-partisan content is because of independent groups.

a. Will you commit to, by the end of the year, providing five or more independent non-profit entities with access to the data they need to understand the prevalence and performance of inauthentic, false, or hyper-partisan content on your services? If you will, please provide specifics and a timeline for how you plan to honor the commitment. If you will not, please explain why.

At Facebook, we have been working on the issue of inauthentic and false content for a long time, and we believe that tech companies, media companies, newsrooms, and

educators all need to work together to address these societal problems. We are engaged with partners across these industries to help create a more informed community.

Though we are committed to doing everything we can to reduce the spread of false news, we also need to make sure we take steps to address the problem when people do encounter hoaxes. To that end, we are exploring ways to give people more context about stories so they can make more informed decisions about what to read, trust, and share; we are also exploring ways to give people access to more perspectives about the topics about which they're reading. We are committed to collaborating with news organizations to develop products together, providing tools and services for journalists, and helping people get better information so they can make smart choices about what they read. We have also joined a group of over 25 funders and participants—including tech industry leaders, academic institutions, non-profits and third-party organizations—to launch the News Integrity Initiative, a global consortium focused on helping people make informed judgments about the news they read and share online.

We are also committed to working with others to protect the political process more generally. That's why we are supporting the Defending Digital Democracy Project, which is a bipartisan project launched by the Harvard Kennedy School that works to secure democracies around the world from external influence. We are the founding sponsor of Defending Digital Democracy's first big project: building an information sharing and analysis organization with members in the various critical areas of the democratic process.

We would need to evaluate any request for data on a case-by-case basis that would include understanding the information security practices and capabilities of a third-party organization.

6. Addressing state-sponsored information operations will continue to require cooperation among private sector entities and with the government.

a. Have you established a formal mechanism for promptly sharing actionable information about the state-sponsored information operations with other online services, similar to the mechanisms that already exist for sharing information about state-sponsored cybersecurity threats? If not, will you commit to developing such a mechanism?

We agree that information-sharing among companies is critical to combating constantly evolving cyber threats, including state-sponsored information operations. We have been working with many others in the technology industry, including Google and Twitter, on this issue, building on our long history of working together on issues like child safety and counterterrorism. We are already sharing more but agree that there must be a formal, enduring effort in this area, and are working to establish an independent organization dedicated to these efforts.

b. The FBI is the federal agency responsible for countering foreign propaganda. Do you have a written policy of promptly sharing what you learn about state-sponsored information operations with the FBI? If not, will you commit to developing such a policy?

We have a long history of working successfully with the DOJ, the FBI, and other law enforcement to address a wide variety of threats to our platform, including threats emanating from Russia. When appropriate, we share our understanding of abusive behavior on our platform with these authorities, and when we learn of a situation involving an imminent threat of harm to a person, we immediately report the situation to first responders. Further, as part of official investigations, government officials sometimes request data about people who use Facebook. We have strict processes in place to handle these government requests, and we disclose account records solely in accordance with our terms of service and applicable law.

7. You currently have automated systems in place to detect spam and abuse.

a. Do you have an automated system in place to detect state-sponsored information operations? If yes, will you provide this to the Committee with a private briefing on the system's design and performance? If no, why not?

We are constantly refining our technical systems to detect and remove abuse on our platform. We have systems that identify and remove inauthentic accounts, and we prioritize signals that are difficult for sophisticated bad actors to disguise. We have also incorporated signals specifically related to accounts connected to social or political issues. And we use automated systems to help identify hate speech, terrorist and extremist content, and other forms of abusive activity on our platform.

These systems do not differentiate on the basis of whether the bad actor is or may be state-sponsored. The challenge of attributing online activities to specific governments or organizations is widely recognized in the intelligence and law enforcement communities. It is extremely challenging to definitively attribute online activity to particular threat actors, and we often rely on information from others, like information included in the January 6, 2017 DNI report, to identify actors behind abuse that we observe and to better understand these issues.

We do not share detailed descriptions of how our tools work in order to avoid providing a road map to bad actors who are trying to avoid detection.

8. You have promised to adopt additional transparency and verification requirements for political advertising.

a. Please detail the new requirements and your timeline for implementing those requirements

We support efforts to promote greater transparency in political advertising online and recently announced steps to make advertising on Facebook more transparent,

increase requirements for authenticity, and strengthen our enforcement against ads that violate our policies. *See* https://newsroom.fb.com/news/2017/10/update-on-our-advertising-transparency-and-authenticity-efforts/. We'll require more thorough documentation from advertisers who want to run election-related ads. We expect these reforms to be in effect by the 2018 U.S. federal elections this fall, and will progress from there to additional contests and elections in other countries and jurisdictions. As part of the documentation process, advertisers may be required to identify that they are running election-related advertising and verify both their entity and location. Once verified, these advertisers will have to include a disclosure in their election-related ads, which reads: "Paid for by." When users click on the disclosure, they will be able to see details about the advertiser, and we will maintain a searchable archive of information. Like other ads on Facebook, they will also be able to see an explanation of why they saw that particular ad. For political advertisers that do not proactively disclose themselves, we are building machine learning tools that will help us find them and require them to verify their identity.

b. How do you define the political advertisements that are covered by the new requirements? Why did you adopt the definition that you did?

We support efforts to promote greater transparency in political advertising online and recently announced steps to make advertising on Facebook more transparent, increase requirements for authenticity, and strengthen our enforcement against ads that violate our policies. But our commitment to ad transparency is not limited to political ads. While our most recent announcements have focused on election-related ads—although not necessarily only ads that mention candidates by name—we are bringing greater transparency to all ads by making sure that people can see all of the ads run by any Page, regardless of whether those ads are targeted to them.

c. Will you commit to including within your definition, at a minimum, advertisements that advocate for or against a specific candidate, political party, piece of legislation, regulatory action, or ballot referendum? If not, why not?

Our commitment to ad transparency is not limited to political ads. While our most recent announcements have focused on election-related ads—although not necessarily only ads that mention candidates by name—we are bringing greater transparency to all ads by making sure that people can see all of the ads run by any Page, regardless of whether those ads are targeted to them.

9. Your platform offers a range of advertisement targeting criteria.

a. Which types of targeting criteria, such as demographic, behavioral, lookalike, or email matching did Russia use for its information operations?

We have provided the targeting criteria to this Committee along with the ads themselves.

10. Have you seen any evidence of state-sponsored information operations associated with American elections in 2017, including the gubernatorial elections in Virginia and New Jersey?

We have learned from the 2016 election cycle and from elections worldwide this last year. We have incorporated that learning into our automated systems and human review and have greatly improved in preparation for the upcoming elections. We hope to continue learning and improving through increased industry cooperation and dialogue with law enforcement moving forward.

11. User reports are an important signal of when an account is not authentic.

a. How frequently do you receive user reports about inauthentic accounts?

User reports are an important signal, and we rely on our community to help identify inauthentic accounts. Facebook's Community Operations team receives millions of reports each week about content that may violate our community standards, including about potential inauthentic accounts.

b. What is your process for responding to those reports? How often does that process usually take?

Our community of users helps us by reporting accounts or content that may violate our policies. Our Community Operations team—which is growing significantly over this year—works around the world, 24 hours a day, and in dozens of languages to review these reports.

c. What proportion of those reports result in an account restriction, suspension, or removal?

When we review accounts that have been reported, we look for any violation of our policies, not just for the violation that may have been identified by the reporter. When our review identifies a policy violation, remedies may include feature disabling, content removal, checkpointing, or account disabling. The severity of the remedy is dictated by the policy violation in question, as well as its frequency and whether the account is a repeat offender.

d. Among the reports that you decline to take action on, what proportion involve reported accounts that you subsequently identify as inauthentic?

Our automated processes for identifying fake accounts typically take effect before we receive reports. Whether an account has been reported multiple times does not affect whether it is disabled. Our reviewers are trained to look for violations and enforce our policies consistently and as objectively as possible.

e. How many of the accounts that you have identified as associated with Russian information operations were the subject of a user report? Please provide all the user reports associated with these accounts and the actions you took in response including the specific time for the report and each action.

We are not aware that any of these accounts were reported as inauthentic.

12. Much of the public discussion about state-sponsored information operations on your platforms has centered on the Internet Research Agency. That is not the only group surreptitiously spreading state-sponsored propaganda.

a. What other groups are you tracking that are affiliated with the Russian government?

We have used "the Internet Research Agency" or "the IRA" to describe a set of actors that were active on our platform during the 2016 U.S. election cycle. We removed accounts linked to those actors in August 2017, and we are continually on the lookout for recidivist actors. These actors may have called themselves by other names, including Glavset.

We have also tracked activity from a cluster of accounts we have assessed to belong to a group, APT28, that the U.S. government has publicly linked to Russian military intelligence services and the "DCLeaks" organization.

b. What other countries do you believe are conducting state-sponsored information operations on your platforms? Please describe the groups that you are tracking for each country, including both government agencies and affiliates.

It is extremely difficult to definitively attribute online activity to particular threat actors, including state actors. We are constantly monitoring for threat actors who, like the IRA, are engaged in state-sponsored information operations.

13. According to news reports, Russia's propaganda is not just online. Russia appears to have organized pro-Trump rallies across Florida during the 2016 campaign, for example, convened an anti-immigrant, anti-Muslim rally in Idaho, and orchestrated a protest and a counter-protest at an Islamic center in Texas.

a. What data does Facebook possess about the real-world gatherings that were associated with Russian propaganda?

A total of 129 events were created across 13 IRA Pages. Approximately 338,300 unique accounts viewed these events. About 25,800 accounts marked that they were interested in an event, and about 62,500 marked that they were going to an event. We do not have data about the realization of these events.

b. Will Facebook make that data public to increase awareness of the threat?

We believe we have a responsibility to be transparent about what we know regarding foreign influence in our election, which is why, as we have discovered information, we have come forward with it to this Committee and to share information publicly. In an effort to provide additional transparency, we recently launched a portal to enable people on Facebook to learn which of the IRA Facebook Pages or Instagram accounts they may have liked or followed between January 2015 and August 2017. This tool is available to users in the Facebook Help Center.

c. How many individuals checked-in at, liked, or shared event pages associated with Russian propaganda?

See response to question 13(a).

14. Recent news reports indicated that Facebook collects, analyzes, and uses a range of information from consumers for purposes of the "Find Friends" feature.

a. What information does Facebook collect, analyze, and use for this feature? How is the data collected, analyzed, and used?

b. Is the information visible to a consumer? If not, will you commit to making this information visible?

c. Can a consumer edit or delete the information? If not, will you commit to enabling consumers to edit or delete the information?

d. Can an advertiser target ads based on this information?

e. Is there any indication Russia may have used or attempted to use this information to advance its information operations?

People You May Know can help Facebook users find friends on Facebook. People You May Know suggestions come from things such as having friends in common, or mutual friends; being in the same Facebook group or being tagged in the same photo; users' networks (for example, school or work); and contacts users have uploaded. We give people context when we suggest someone with mutual friends. Users may delete contacts that they have uploaded to Facebook, in which case that information is no longer used for People You May Know. Facebook does not allow advertisers to target ads based on People You May Know.

15. Inauthentic accounts can be disabled subsequent to automated or manual review.

a. What role do automated and human review play in your decision to disable a suspected inauthentic account?

We rely on both automated and manual review, and we are now taking steps to strengthen both. We continually update our technical systems to identify, checkpoint, and remove inauthentic accounts. We block millions of attempts to register fake accounts every day. These systems examine thousands of detailed account attributes and prioritize signals that are more difficult for bad actors to disguise, such as their connections to others on our platform. We do not share detailed descriptions of how our tools work in order to avoid providing a road map to bad actors who are trying to avoid detection.

b. Do you require that a human employee review a suspected inauthentic account before it is disabled?

When we suspect that an account is inauthentic, we typically enroll the account in a checkpoint that requires the account holder to provide additional information or verification. We view disabling an account as a severe sanction, and we want to ensure that we are highly confident that the account violates our policies before we take permanent action. When we have confirmed that an account violates our policies, we remove the account.

c. If so, given the rate at which inauthentic accounts can be regenerated, how do you anticipate remaining ahead of the problem?

We continually update our technical systems to identify, checkpoint, and remove inauthentic accounts, and we block millions of attempts to register fake accounts every day. These systems examine thousands of detailed account attributes and prioritize signals that are more difficult for bad actors to disguise, such as their connections to others on our platform. We also look for recidivists who repeatedly create inauthentic accounts. We do not share detailed descriptions of how our tools work in order to avoid providing a road map to bad actors who are trying to avoid detection.

d. What are you doing to improve automation in the process of detecting and disabling inauthentic accounts?

As explained above, we continue to learn common traits of bad actors and to update our systems to better identify and remove inauthentic accounts.

e. What are you doing to make it more difficult to establish inauthentic accounts?

The combination of our constantly updated technical systems and human review aids us in preventing the creation of fake accounts and in discovering and removing them if they are created. It is an evolving process and as new methods are used by bad actors creating inauthentic accounts, we will develop new solutions for identification and prevention.

Senator McCain

1. **Current campaign finance law establishes disclosure standards for television, radio, and print media. The Pew Research Center recently found that 65 percent of Americans identified an Internet-based source as their leading source of information about the 2016 election.**

    a. **Under current law, to what extent is Facebook responsible for providing a similar quality of disclosure to the public?**

Facebook requires advertisers to comply with all applicable laws and with our policies, including our authenticity policy. We are also implementing new verification and disclosure standards on Facebook that will bring greater transparency to political advertising on our platform in general and make it easier for us to enforce our policies. Starting with the 2018 U.S. federal elections this fall, we will require these advertisers to identify who is paying for the ads to Facebook and to the public. For election advertisers who do not self-identify, we are building automated tools that will help us identify these ads proactively.

2. **The *Honest Ads Act* requires disclosures and disclaimers for entities engaging in paid and intentional campaign activity, all in an effort to better understand and detect illegal activity, like foreign interference. Given the unique nature and accessibility of social media, bringing campaign finance law into the digital realm poses some challenges.**

    a. **What can Facebook and Congress do to ensure that free speech is not violated in this process?**

Facebook is a platform for people to express themselves freely and openly, and we are working to strike the right balance between enabling free expression and ensuring that our platform is safe. Our policies are aimed at encouraging expression and respectful dialogue. Facebook believes that transparency in political advertising and authentic dialogue promotes open and effective democracies. We are taking several steps to make it clear who is running election ads on Facebook and to identify and remove inauthentic accounts. To ensure that advertising is as transparent on other platforms as it is on Facebook, we support industry-wide standards that provide clear and consistent guidance to advertisers regarding their disclosure obligations.

3. **The public nature of television, radio, and print ensures that political advertisements are subject to scrutiny of the press, fact-checkers, and political opponents. Political advertisements sold on Facebook, however, can be targeted towards specific groups, and frequently bear no indication as to the purchaser of the advertisement.**

    a. **Does Facebook have plans to create a permanent online record where the public is able to monitor the content and ownership of certain advertisements?**

We support efforts to promote greater transparency in political advertising online and recently announced steps to make advertising on Facebook more transparent, increase requirements for authenticity, and strengthen our enforcement against ads that violate our policies. *See* https://newsroom.fb.com/news/2017/10/update-on-our-advertising-transparency-and-authenticity-efforts/. We'll require more thorough documentation from advertisers who want to run election-related ads. We expect these reforms to be in effect by the 2018 U.S. federal elections this fall and will progress from there to additional contests and elections in other countries and jurisdictions. As part of the documentation process, advertisers may be required to identify that they are running election-related advertising and verify both their entity and location. Once verified, these advertisers will have to include a disclosure in their election-related ads, which reads: "Paid for by." When users click on the disclosure, they will be able to see details about the advertiser, and we will maintain a searchable archive of information. Like other ads on Facebook, they will also be able to see an explanation of why they saw that particular ad. For political advertisers that do not proactively disclose themselves, we are building machine learning tools that will help us find them and require them to verify their identity.

4. Facebook's Terms of Service require that paid advertisers use authentic information in the purchase and promotion of their content.

a. What steps has Facebook taken to monitor and enforce this requirement?

Facebook builds and updates technical systems every day to better identify and remove inauthentic accounts, which also helps reduce the distribution of material that can be spread by accounts that violate our policies. Each day, we block millions of fake accounts at registration. Our systems examine thousands of account attributes and focus on detecting behaviors that are very difficult for bad actors to fake, including their connections to others on our platform. By constantly improving our techniques, we also aim to reduce the incentives for bad actors who rely on distribution to make their efforts worthwhile.

We are also implementing new verification and disclosure standards on Facebook that will bring greater transparency to political advertising on our platform in general and make it easier for us to enforce our policies. Starting with the 2018 U.S. federal elections this fall, we will require these advertisers to identify who is paying for the ads to Facebook and to the public. For election advertisers who do not self-identify, we are building automated tools that will help us identify these ads proactively.

HEARING BEFORE THE UNITED STATES SENATE SELECT COMMITTEE ON INTELLIGENCE

November 1, 2017

Testimony of Colin Stretch
General Counsel, Facebook

I. INTRODUCTION

Chairman Burr, Vice Chairman Warner, and distinguished members of the Committee, thank you for this opportunity to appear before you today. My name is Colin Stretch, and since July 2013, I've served as the General Counsel of Facebook. We appreciate this Committee's hard work to investigate Russian interference in the 2016 election.

At Facebook, our mission is to create technology that gives people the power to build community and bring the world closer together. We don't take for granted that each one of you uses Facebook to connect with your constituents, and that the people you represent expect authentic experiences when they come to our platform to share.

We also believe we have an important role to play in the democratic process—and a responsibility to protect it on our platform. That's why we take what's happened on Facebook so seriously. The foreign interference we saw is reprehensible and outrageous and opened a new battleground for our company, our industry, and our society. That foreign actors, hiding behind fake accounts, abused our platform and other internet services to try to sow division and discord—and to try to undermine our election process—is an assault on democracy, and it violates all of our values.

In our investigation, which continues to this day, we've found that these actors used fake accounts to place ads on Facebook and Instagram that reached millions of Americans over a two-year period, and that those ads were used to promote Pages, which in turn posted more content. People shared these posts, spreading them further. Many of these ads and posts are inflammatory. Some are downright offensive.

In aggregate, these ads and posts were a very small fraction of the overall content on Facebook—but any amount is too much. All of these accounts and Pages violated our policies, and we removed them.

Going forward, we're making some very significant investments—we're hiring more ad reviewers, doubling or more our security engineering efforts, putting in place tighter ad content restrictions, launching new tools to improve ad transparency, and requiring documentation from political ad buyers. We're building artificial intelligence to help locate more banned content, and bad actors. We're working more closely with industry to share information on how to identify and prevent threats so that we can all respond faster and more effectively. And we are expanding our efforts to work more closely with law enforcement.

I'm here today to share with you what we know so far about what happened—and what we're doing about it. At the outset, let me explain how our service works and why people choose to use it.

II. FIGHTING ELECTION INTERFERENCE ON FACEBOOK

A. Understanding what you see on Facebook

1. The News Feed Experience: A Personalized Collection of Stories. When people come to Facebook to share with their friends and discover new things, they see a personalized homepage we call News Feed. News Feed is a constantly updating, highly personalized list of stories, including status updates, photos, videos, links, and activity from the people and things you're connected to on Facebook. The goal of News Feed is to show people the stories that are most relevant to them. The average person has thousands of things on any given day that they could read in their News Feed, so we use personalized ranking to determine the order of stories we show them. Each person's News Feed is unique. It's shaped by the friends they add; the people, topics, and news sources they follow; the groups they join; and other signals like their past interactions. On average, a person in the US is served roughly 220 stories in News Feed each day. Over the time period in question, from 2015 to 2017, Americans using Facebook were exposed to, or "served," a total of over 33 trillion stories in their News Feeds.

2. Advertising and Pages as Sources of Stories in News Feed. News Feed is also a place where people see ads on Facebook. To advertise in News Feed, a person must first set up a Facebook account—using their real identity—and then create a Facebook Page. Facebook Pages represent a wide range of people, places, and things, including causes, that people are interested in. Any user may create a Page to express support for or interest in a topic, but only official representatives can create a Page on behalf of an organization, business, brand, or public figure. It is against our terms for Pages to contain false, misleading, fraudulent, or deceptive claims or content. Facebook marks some official Pages—such as for a public figure, media company, or brand—with a "verified" badge to let people know they're authentic. All Pages must comply with our Community Standards and ensure that all the stories they post or share respect our policies prohibiting hate speech, violence, and sexual content, among other restrictions. People can like or follow a Page to get updates, such as posts, photos, or videos, in their News Feed. The average person in the US likes 178 Pages. People do not necessarily see every update from each of the Pages they are connected to. Our News Feed ranking determines how relevant we think a story from a Page will be to each person. We make it easy for people to override our recommendations by giving them additional controls over whether they see a Page's updates higher in their News Feed or not at all. For context, from 2015 to 2017, people in the United States saw 11.1 trillion posts from Pages on Facebook.

3. Advertising to Promote Pages. Page administrators can create ads to promote their Page and show their posts to more people. The vast majority of our advertisers are small- and medium-sized businesses that use our self-service tools to create ads to reach their customers. Advertisers choose the audience they want to reach based on demographics, interests, behaviors or contact information. They can choose from different ad formats, upload images or video, and write the text they want people to see. Advertisers can serve ads on our platform for as little as $0.50 per day using a credit card or other payment method. By using these tools, advertisers agree to our

Self-Serve Ad Terms. Before ads appear on Facebook or Instagram, they go through our ad review process that includes automated checks of an ad's images, text, targeting and positioning, in addition to the content on the ad's landing page. People on Facebook can also report ads, find more information about why they are being shown a particular ad, and update their ad preferences to influence the type of ads they see.

B. Promoting Authentic Conversation

Our authenticity policy is the cornerstone of how we prevent abuse on our platform, and was the basis of our internal investigation and what we found.

From the beginning, we have always believed that Facebook is a place for authentic dialogue, and that the best way to ensure authenticity is to require people to use the names they are known by. Fake accounts undermine this objective, and are closely related to the creation and spread of inauthentic communication such as spam—as well as used to carry out disinformation campaigns like the one associated with the Internet Research Agency (IRA).

We build and update technical systems every day to better identify and remove inauthentic accounts, which also helps reduce the distribution of material that can be spread by accounts that violate our policies. Each day, we block millions of fake accounts at registration. Our systems examine thousands of account attributes and focus on detecting behaviors that are very difficult for bad actors to fake, including their connections to others on our platform. By constantly improving our techniques, we also aim to reduce the incentives for bad actors who rely on distribution to make their efforts worthwhile.

Protecting authenticity is an ongoing challenge. As our tools and security efforts evolve, so will the techniques of those who want to evade our authenticity requirements. As in other areas of cybersecurity, our security and operations teams need to continually adapt.

C. Protecting the Security of the 2016 Election and Learning Lessons Quickly

1. The Evolution of Facebook's Security Protections. From its earliest days, Facebook has always been focused on security. These efforts are continuous and involve regular contact with law enforcement authorities in the United States and around the world. Elections are particularly sensitive events for our security operations, and as the role our service plays in promoting political dialogue and debate has grown, so has the attention of our security team.

As your investigation has revealed, our country now faces a new type of national cyber-security threat—one that will require a new level of investment and cooperation across our society. At Facebook, we're prepared to do our part. At each step of this process, we have spoken out about threats to internet platforms, shared our findings, and provided information to investigators. As we learn more, we will continue to identify and implement improvements to our security systems, and work more closely with other technology companies to share information on how to identify and prevent threats and how to respond faster and more effectively.

2. Security Leading Up to the 2016 Election.

a. Fighting Hacking and Malware. For years, we had been aware of other types of activity that

appeared to come from Russian sources—largely traditional security threats such as attacking people's accounts or using social media platforms to spread stolen information. What we saw early in the 2016 campaign cycle followed this pattern. Our security team that focuses on threat intelligence—which investigates advanced security threats as part of our overall information security organization—was, from the outset, alert to the possibility of Russian activity. In several instances before November 8, 2016, this team detected and mitigated threats from actors with ties to Russia and reported them to US law enforcement officials. This included activity from a cluster of accounts we had assessed to belong to a group ("APT28") that the US government has publicly linked to Russian military intelligence services. This activity, which was aimed at employees of major US political parties, fell into the normal categories of offensive cyber activities we monitor for. We warned the targets who were at highest risk, and were later in contact with law enforcement authorities about this activity.

Later in the summer we also started to see a new kind of behavior from APT28-related accounts—namely, the creation of fake personas that were then used to seed stolen information to journalists. These fake personas were organized under the banner of an organization that called itself DC Leaks. This activity violated our policies, and we removed the DC Leaks accounts.

b. Understanding Fake Accounts and Fake News. After the election, when the public discussion of "fake news" rapidly accelerated, we continued to investigate and learn more about the new threat of using fake accounts to amplify divisive material and deceptively influence civic discourse. We shared what we learned with government officials and others in the tech industry. And in April 2017, we shared our findings with the public by publishing a white paper that described the activity we detected and the initial techniques we used to combat it.

As with all security threats, we have also been applying what we learned in order to do better in the future. We use a variety of technologies and techniques to detect and shut down fake accounts, and in October 2016, for example, we disabled about 5.8 million fake accounts in the United States. At the time, our automated tooling did not yet reflect our knowledge of fake accounts focused on social or political issues. But we incorporated what we learned from the 2016 elections into our detection systems, and as a result of these improvements, we disabled more than 30,000 accounts in advance of the French election. This same technology helped us disable tens of thousands more accounts before the German elections in September. In other words, we believe that we're already doing better at detecting these forms of abuse, although we know that people who want to abuse our platform will get better too and so we must stay vigilant.

3. Investigating the Role of Ads and Foreign Interference. After the 2016 election, we learned from press accounts and statements by congressional leaders that Russian actors might have tried to interfere in the election by exploiting Facebook's ad tools. This is not something we had seen before, and so we started an investigation that continues to this day. We found that fake accounts associated with the IRA spent approximately $100,000 on more than 3,000 Facebook and Instagram ads between June 2015 and August 2017. Our analysis also showed that these accounts used these ads to promote the roughly 120 Facebook Pages they had set up, which in turn posted more than 80,000 pieces of content between January 2015 and August 2017. The Facebook accounts that appeared tied to the IRA violated our policies because they came from a

set of coordinated, inauthentic accounts. We shut these accounts down and began trying to understand how they misused our platform.

a. Advertising by Accounts Associated with the IRA. Below is an overview of what we've learned so far about the IRA's ads:

- **Impressions (an "impression" is how we count the number of times something is on screen, for example this can be the number of times something was on screen in a person's News Feed):**
  - 44% of total ad impressions were before the US election on November 8, 2016.
  - 56% of total ad impressions were after the election.
- **Reach (the number of people who saw a story at least once):**
  - We estimate 11.4 million people in the US saw at least one of these ads between 2015 and 2017.
- **Ads with zero impressions:**
  - Roughly 25% of the ads were never shown to anyone. That's because advertising auctions are designed so that ads reach people based on relevance, and certain ads may not reach anyone as a result.
- **Amount spent on ads:**
  - For 50% of the ads, less than $3 was spent.
  - For 99% of the ads, less than $1,000 was spent.
  - Many of the ads were paid for in Russian currency, though currency alone is a weak signal for suspicious activity.
- **Content of ads:**
  - Most of the ads appear to focus on divisive social and political messages across the ideological spectrum, touching on topics from LGBT matters to race issues to immigration to gun rights.
  - A number of the ads encourage people to follow Pages on these issues, which in turn produced posts on similarly charged subjects.

b. Content Posted by Pages Associated with the IRA. We estimate that roughly 29 million people were served content in their News Feeds directly from the IRA's 80,000 posts over the two years. Posts from these Pages were also shared, liked, and followed by people on Facebook, and, as a result, three times more people may have been exposed to a story that originated from the Russian operation. Our best estimate is that approximately 126 million people may have been served content from a Page associated with the IRA at some point during the two-year period. This equals about four-thousandths of one percent (0.004%) of content in News Feed, or approximately 1 out of 23,000 pieces of content.

Though the volume of these posts was a tiny fraction of the overall content on Facebook, **any amount is too much**. Those accounts and Pages violated Facebook's policies—which is why we

removed them, as we do with all fake or malicious activity we find. We also deleted roughly 170 Instagram accounts that posted about 120,000 pieces of content.

Our review of this activity is ongoing. Many of the ads and posts we've seen so far are deeply disturbing—seemingly intended to amplify societal divisions and pit groups of people against each other. They would be controversial even if they came from authentic accounts in the United States. But coming from foreign actors using fake accounts they are simply unacceptable.

That's why we've given the ads and posts to Congress—because we want to do our part to help investigators gain a deeper understanding of foreign efforts to interfere in the US political system and explain those activities to the public. These actions run counter to Facebook's mission of building community and everything we stand for. And we are determined to do everything we can to address this new threat.

D. Mobilizing to Address the New Threat

We are taking steps to enhance trust in the authenticity of activity on our platform, including increasing ads transparency, implementing a more robust ads review process, imposing tighter content restrictions, and exploring how to add additional authenticity safeguards.

1. Promoting Authenticity and Preventing Fake Accounts. We maintain a calendar of upcoming elections and use internal and external resources to best predict the threat level to each. We take preventative measures based on our information, including working with election officials where appropriate. Within this framework, we set up direct communication channels to escalate issues quickly. These efforts complement our civic engagement work, which includes voter education. In October 2017, for example, we launched a Canadian Election Integrity Initiative to help candidates guard against hackers and help educate voters on how to spot false news.

Going forward, we're also requiring political advertisers to provide more documentation to verify their identities and disclose when they're running election ads. Potential advertisers will have to confirm the business or organization they represent before they can buy ads. Their accounts and their ads will be marked as political, and they will have to show details, including who paid for the ads. We'll start doing this with federal elections in the US and then move onto other elections in the US and other countries**.** For political advertisers that don't proactively identify themselves, we're building machine learning tools that will help us find them and require them to verify their identity.

Authenticity is important for Pages as well as ads. We'll soon test ways for people to verify that the people and organizations behind political and issue-based Pages are who they say they are.

2. Partnering with Industry on Standards. We have been working with many others in the technology industry, including with Google and Twitter, on a range of elements related to this investigation. Our companies have a long history of working together on other issues such as child safety and counter-terrorism.

We are also reaching out to leaders in our industry and governments around the world to share information on bad actors and threats so that we can make sure they stay off all platforms. We

are trying to make this an industry standard practice.

3. Strengthening Our Advertising Policies. We know that some of you and other members of Congress are exploring new legislative approaches to political advertising—and that's a conversation we welcome. We are already working with some of you on how best to put new requirements into law. But we aren't waiting for legislation. Instead we're taking steps where we can on our own, to improve our own approach to transparency, ad review, and authenticity requirements.

a. Providing Transparency. We believe that when you see an ad, you should know who ran it to be able to understand what other ads they're running—which is why we show you the Page name for any ads that run in your News Feed.

To provide even greater transparency for people and accountability for advertisers, we're now building new tools that will allow you to see the other ads a Page is running as well—including ads that aren't targeted to you directly. We hope that this will establish a new standard for our industry in ad transparency. We try to catch material that shouldn't be on Facebook before it's even posted—but because this is not always possible, we also take action when people report ads that violate our policies. We're grateful to our community for this support, and hope that more transparency will mean more people can report violating ads.

b. Enforcing Our Policies. We rely on both automated and manual ad review, and we're now taking steps to strengthen both. Reviewing ads means assessing not just what's in an ad but also the context in which it was bought and the intended audience—so we're changing our ads review system to pay more attention to these signals. We're also adding more than 1,000 people to our global ads review teams over the next year and investing more in machine learning to better understand when to flag and take down ads. Enforcement is never perfect, but we will get better at finding and removing improper ads.

c. Restricting Ad Content. We hold people on Facebook to our Community Standards, and we hold advertisers to even stricter guidelines. Our ads policies already prohibit shocking content, direct threats and the promotion of the sale or use of weapons. Going forward, we are expanding these policies to prevent ads that use even more subtle expressions of violence.

III. CONCLUSION

Any attempt at deceptive interference using our platform is unacceptable, and runs counter to everything we are working toward. What happened in the 2016 election cycle was an affront to us, and, more importantly, to the people who come to Facebook every day to have authentic conversations and to share. We are committed to learning from these events, and to improving. We know we have a responsibility to do our part—and to do better. We look forward to working with everyone on this Committee, in the government, and across the tech industry and civil society, to address this important national security matter so that we can prevent similar abuse from happening again.