October 26, 2018

Chairman Richard Burr
Vice Chairman Mark Warner
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 September 5, 2018 Hearing titled Foreign Influence Operations' Use of Social Media Platforms. Per your request, attached are the answers for the record to your questions.

Please note that our work on many of the matters discussed by your questions is ongoing. We did our best to review and answer them in the available timeframe. We respectfully request an opportunity to supplement or amend our responses if needed.

Sincerely,

Facebook, Inc.

facebook

Address: 1601 Willow Road
Menlo Park, CA 94025

Questions for the Record
Senate Select Committee on Intelligence
Hearing on Foreign Influence
Operations Using Social Media
September 17, 2018

[From Chairman Burr]

1. **Aleksandr Kogan served as director of Global Science Research (GSR) where he used an app to harvest data from as many as 87 million Facebook users. Facebook has said publicly that Kogan claimed the data would only be used for academic purposes and then "lied to us" in passing the content to Cambridge Analytica.**

- **Did Facebook data scientists co-author academic papers with GSR co-founders Aleksandr Kogan and Joseph Chancellor?**
    - **If yes, how does this reconcile with Facebook's asserting a complete unawareness as to GSR's user data harvesting practices?**

Facebook was put in touch with Kogan (a researcher at the University of Cambridge) in late 2012 about a possible collaboration on research relating to the potential relationship between Facebook friendship ties and economic trade volumes between countries. Kogan collaborated with current and former Facebook employees on approximately 10 academic papers. As part of these collaborations, Kogan could only access fully anonymized, aggregated data. The anonymized, aggregated data provided to Kogan as part of the academic research collaboration were entirely separate from the data that GSR independently obtained from users through its App. We have not found evidence to suggest that the work Chancellor undertook at Facebook had any relationship to the work he performed when he was working with Kogan and Global Science Research Limited (GSR).

Facebook frequently partners with leading academic researchers to address topics pertaining to wellbeing, innovation, and other topics of public importance, following strict protocols to ensure personal information is safeguarded.

- **If Facebook found GSR to be in violation of its arrangement with Facebook, why did Facebook continue to employ former GSR co-founder Joseph Chancellor?**

Joseph Chancellor was a quantitative researcher on the User Experience Research team at Facebook, whose work focused on aspects of virtual reality. He is no longer employed by Facebook.

2. **On February 6, 2018, the day after the Senate Commerce and Judiciary hearing, Facebook terminated Joseph Chancellor's employment. What were the circumstances of his termination?**

- **What was the hire date (month and year) of Joseph Chancellor, co-founder of GSR?**

Joseph Chancellor's first day at Facebook was November 9, 2015. Chancellor's title was "Quantitative User Experience Researcher." On March 26, 2018, Joseph Chancellor was placed on (non-disciplinary) administrative leave. He is no longer employed at Facebook.

- **Were you or CEO Mark Zuckerberg aware of the hiring of Joseph Chancellor?**

Facebook has over 30,000 employees. Senior management does not participate in day-to-day hiring decisions.

[From Vice Chairman Warner]

3. On July 17th, in a podcast with Kara Swisher, Mark Zuckerberg said Facebook was "a long time away from doing anything" in China. On July 24th, the Washington Post reported that Facebook had registered a new subsidiary in China.

- **What is the current status of Facebook's engagement with China?**

- **Do you have existing plans for attempting to enter the Chinese market? If yes, please describe.**

- **Are there any current discussions underway within Facebook about entering China? If yes, please describe.**

Because Facebook has been blocked in China since 2009, we are not in a position to know exactly how the government would seek to apply its laws and regulations on content were we permitted to offer our service to Chinese users. Since 2013, Facebook has been a member of the Global Network Initiative (GNI), a multi-stakeholder digital rights initiative. As part of our membership, Facebook has committed to the freedom of expression and privacy standards set out in the GNI Principles—which are in turn based on the Universal Declaration of Human Rights and the United Nations Guiding Principles on Business and Human Rights—and we are independently assessed on our compliance with these standards on a biennial basis.

In keeping with these commitments, rigorous human rights due diligence and careful consideration of free expression and privacy implications would constitute important components of any decision on entering China. Facebook has been blocked in China since 2009, and no decisions have been made around the conditions under which any possible future service might be offered in China.

4. In responding to a question from Senator Rubio about potential engagement in China, you said that Facebook "would only operate in a country when we can do so in keeping with our values."

- **What do you consider Facebook's values to be?**

Facebook's mission is to give people the power to build community and bring the world closer together. We also recently announced new principles. Our principles are what we stand for, what we will fight to provide for people, and what kind of community we want to build. They are beliefs we hold deeply and that we already make real tradeoffs to pursue.

- **Give people a voice:** The one phrase in our mission that has never changed is "give people the power," and one of the ways we do that is by giving people a voice. This means we err on the side of free expression—even when that means defending the right of people we deeply disagree with to say things that are controversial or offensive. Of course, there are limits. We don't allow content that incites violence or attacks people, whether that's terrorism or bullying or hate.

- **Build connection and community:** Our services help people connect more, and when they're at their best, they also bring people closer together. That's why this year we reworked News Feed to prioritize meaningful social interactions over passive consumption. In order for a community to be cohesive, it must share enough common ground—so while we give everyone a voice, we must make sure misinformation doesn't spread virally and high quality, broadly trusted information is available to all.

- **Serve everyone:** Everyone deserves access to these tools. That's why we operate in countries where we might lose money, why we work on Internet.org to spread connectivity to people who can't even afford it, why our business model is ads—so our service can be free for everyone. And it's also why, when a country passes a law limiting voice or that conflicts with one of our other principles, we fight to make sure the service remains available for as many people as possible.

- **Keep people safe and protect privacy:** People try to use our services for good and bad, and we have a responsibility promote the good and prevent harm. That's why we have the initiatives on counterterrorism and self-harm. That's why we have more than doubled the number of people working on safety and security and now have over 20,000. And that's why, even though we care about giving people a voice, we take down a lot of content that is bullying, harassing, and attacking people.

- **Promote economic opportunity:** We talk a lot about the social aspect of community, but strong communities also provide people opportunity. Through our work helping small businesses grow, we aim to create more jobs and opportunity than any other company out there. Our services empower people, and our work supporting economic opportunity—whether it's through Marketplace, Pages, WhatsApp, Messenger or Instagram—is a fundamental part of what we stand for.

- **Which of those values will you weigh when considering potential engagement in China?**

We consider all of these values in evaluating our activities in all countries around the world. See Response to Question 3.

5. You have indicated your company's strong support for the Honest Ads Act. Thank you for your support and your efforts to largely abide by the terms of that legislation.

- **Do you support passage of the Honest Ads Act into law?**

Yes. We have taken proactive steps to require that advertisers clearly label all election-related and issue ads on Facebook and Instagram in the US—including a "Paid for by" disclosure

from the advertiser at the top of the ad. This will help people see who is paying for the ad—which is especially important when the Page name doesn't match the name of the company or person funding the ad. For more information, see https://newsroom.fb.com/news/2018/04/transparent-ads-and-pages/.

Our policy reflects language from existing laws as well as proposed laws. But we're not waiting for proposed legislation to pass before we act. We've been hearing calls for increased transparency around ads with political content for some time now. We've taken the first steps toward providing that transparency, and we hope others follow.

- **Have you seen evidence – in either the Russian context or any recent disruptions –that your new policies on ad transparency have helped stop foreign purchases of political ads on your platform?**

The policies and processes focused on transparency that we have implemented for advertisers on Facebook have created structural disincentives for bad actors to try to meddle and interfere in the electoral process. Our requirement that advertisers wanting to run ads with political or issue content in the US and certain other countries will need to verify their identity and location adds an important step to deterring some bad actors from running these types of ads.

The past few months have shown that bad actors have had to work harder to cover their tracks, in part due to the actions we've taken to help prevent abuse over the past year. We have removed many Pages and accounts from Facebook and Instagram because they were involved in coordinated inauthentic behavior, which is not allowed on Facebook. Since last fall, we have publicly announced more than 10 takedowns for inauthentic behavior.

But security is not something that's ever done. Determined and well-funded bad actors are persistent and constantly changing tactics. For these reasons, in addition to our implementing transparency measures in ads, Facebook has invested heavily in more people and better technology to help prevent bad actors misusing Facebook—as well as working much more closely with law enforcement and other tech companies to better understand the threats we face.

6. **Facebook has taken some steps to ensure transparency in political ads. One of the key disclosure provisions in the Honest Ads Act is a requirement to disclose "a description of the audience targeted by the advertisement." While your current ad archive reports certain information on the reach of the ad – including gender, state, and age – it does not appear that the archive reports on the ad purchaser's targeting criteria and its intended target.**

- **Does Facebook plan to disclose ad targeting data in the ad archive so users can see how the ad was specifically targeted?**

The archive displays general information about the amount spent on the ad, the number of people who saw it, plus aggregated, anonymized data on their age, gender and location.

- **Why or why not? If not, will you consider including targeting information in your transparency measures, similar to the Honest Ads Act requirements?**

We show information and demographic breakdown of people who actually saw the ad. We believe the actual breakdown of who saw a particular ad with political or issue content is more meaningful in understanding the ad's impact than its intended audience. However, we'll continue listening to feedback and working to improve our transparency tools.

7. **Under the terms of your 20-year consent decree with the FTC, Facebook was required to establish a "comprehensive privacy policy" to undertake, among other things, "the identification of reasonably foreseeable, material risks, both internal and external, that could result in [Facebook's] unauthorized collection, use, or disclosure of covered information and an assessment of the sufficiency of any safeguards in place to control these risks." The consent decree says this should extend to "product design, development, and research."**

- **Does Facebook believe its failure to identify and address the privacy concerns of allowing data access to third-party applications like Aleksandr Kogan's Global Science Research (GSR) is consistent with the "reasonably foreseeable" language of the FTC consent decree?**

Facebook has complied with the Consent Order. We furnished extensive information to the FTC regarding the ability for users to port their Facebook data (including friends' data that had been shared with them) with third-party apps on Facebook's platform as part of the FTC's investigation culminating in the July 27, 2012 Consent Order. The Consent Order memorializes the agreement between Facebook and the FTC and did not require Facebook to turn off the ability for people to port friends' data that had been shared with them on Facebook to third-party apps they used.

In addition, Facebook voluntarily limited the ability of people to port friends' data through platform in 2014, which operated as a further technical control to restrict the types of data available to developers on the public platform.

- **Why shouldn't the data breach brought about by the GSR/Cambridge Analytica episode constitute a breach of the consent decree with the FTC?**

As an initial matter, this was not a breach of Facebook's systems. In addition, we do not believe there was a violation of the FTC Consent Order. We furnished extensive information to the FTC regarding the ability for users to port their Facebook data (including friends' data that had been shared with them) with third-party apps on Facebook's platform, as part of the FTC's investigation culminating in the July 27, 2012 Consent Order and in several subsequent briefings and engagements with the FTC. The Consent Order memorializes the agreement between Facebook and the FTC and did not require Facebook to turn off or change the ability for people to port friends' data that had been shared with them on Facebook to apps they used. Facebook voluntarily changed this feature of its public developer platform in 2014, however.

- **Please describe the program Facebook established to comply with your consent agreement with the FTC?**

At Facebook, we make decisions about privacy through a cross-functional, cross-disciplinary effort overseen by the Chief Privacy Officer and our Privacy and Data Use organization that involves participants from departments across the company. This process is a collaborative approach to privacy that seeks to promote strong privacy protections and sound decision making at every stage of the product development process.

Our privacy program contains a number of controls in areas of privacy governance, data transparency, security, risk assessment, third-party developer access, and other areas of potential privacy risk.

Facebook undergoes ongoing privacy assessments to test the effectiveness of these controls pursuant to the July 27, 2012 Consent Order. These assessments are conducted by an independent third-party professional (PwC) pursuant to the procedures and standards generally accepted in the profession and required by the FTC, as set forth in the Consent Order. Facebook's privacy program and related controls are informed by GAPP principles, which are considered industry leading principles for protecting the privacy and security of personal information. Facebook provided the FTC with summaries of the controls and engaged extensively with the FTC regarding the structure of its privacy program. We monitor the privacy program and update the controls as necessary to reflect evolving risks. Facebook has submitted copies of each assessment to the FTC.

- **Did that program fail to flag the Cambridge Analytica sharing?**
    - **If yes, why?**

Our privacy program is a series of more than 40 controls that function to address privacy risk across our product and business operations. It does not function to flag specific incidents such as Cambridge Analytica, although it does contain several controls designed to ensure that third-party app developers obtain consent from people before accessing nonpublic user data through our platform and that developers adhere to our Terms and Data Policy.

8. Facebook learned of Cambridge Analytica's unauthorized access to its data in 2015.

- **Did it notify the FTC at that time? Why or why not?**

We furnished extensive information to the FTC regarding the ability for users to port their Facebook data (including friends' data that had been shared with them) with apps on Facebook's platform, as part of the FTC's investigation culminating in the July 27, 2012 Consent Order and in several subsequent briefings and engagements with the FTC. The Consent Order memorializes the agreement between Facebook and the FTC and did not require Facebook to turn off or change the ability for people to port friends' data that had been shared with them on Facebook to apps they used. Facebook voluntarily changed this feature of its public developer platform in 2014, however.

Instead, and among other things, the Consent Order obligates Facebook not to misrepresent the extent to which it maintains the privacy or security of covered information (Section I), not to materially exceed the restrictions of a privacy setting that applies to nonpublic user information without affirmative express consent (Section II), and to implement a

comprehensive privacy program that is subjected to assessments by an independent assessor (Sections IV and V).

The Consent Order does not contain ongoing reporting obligations to the FTC of the sort suggested in this question. Moreover, Kogan was authorized to access all data that he obtained through Facebook's platform by the people who authorized his app, and no data was shared with Kogan relating to friends who had enabled settings preventing their data from being shared with apps by their friends.

9. Regarding the data from Facebook that was passed from Aleksandr Kogan to Cambridge Analytica, are you aware of that data being passed to any entities outside of the United States or United Kingdom?

Kogan represented that, in addition to providing data to his Prosociality and Well-Being Laboratory at the University of Cambridge for the purposes of research, GSR provided some Facebook data to SCL Elections Ltd., Eunoia Technologies, and the Toronto Laboratory for Social Neuroscience at the University of Toronto. Our investigation is ongoing.

Facebook obtained written certifications from Kogan, GSR, and other third parties (including Cambridge Analytica and SCL) declaring that all data they had obtained, and any derivatives, were accounted for and destroyed. We are seeking to conduct a forensic audit of Cambridge Analytica's systems to confirm the veracity of these certifications, but the UK Information Commissioner's Office, which is conducting a regulatory investigation into Cambridge Analytica (based in the UK), has the only known copy of Cambridge Analytica's systems and will need to release that information for us to conduct this audit. We hope to move forward with that audit soon.

- **Are you aware whether anyone has used the Cambridge Analytica dataset to target advertising on Facebook during the 2016 presidential election or otherwise?**

  See Response to above Question.

10. Transparency on your platform is a significant concern for many of your users. Users should know what data you collect, how you collect that data, and how you monetize that data.

- **Is it a fair expectation for your users that they understand exactly how Facebook data is collected and what types of information you are collecting?**

Yes. We work hard to provide clear information to people about how their information is used and how they can control it. We agree that companies should provide clear and plain information about their use of data and strive to do this in our Data Policy, in in-product notices and education, and throughout our product—and we continuously work on improving this. We provide the same information about our data practices to users around the world and are required under many existing laws—including US laws (e.g., Section 5 of the FTC Act)—to describe our data practices in language that is fair and accurate.

11. Mr. Zuckerberg testified during his appearance before the Senate Commerce and Judiciary Committees, "I think everyone should have control over how their information is used."

- **Do you believe that is an accurate description of the control users on your platform exercise over their own information right now?**

Our approach to control is based on the belief that people should be able to choose who can see what they share and how their data shapes their experience on Facebook and should have control over all data collection and uses that are not necessary to provide and secure our service. We recognize, however, that controls are only useful if people know how to find and use them. That is why we continuously deliver in-product educational videos in people's News Feeds on important privacy topics like how to review and delete old posts and what it means to delete an account. We are also inviting people to take our Privacy Checkup—which prompts people to review key data controls—and we are sharing privacy tips in education campaigns off of Facebook, including through ads on other websites. To make our privacy controls easier to find, we launched a new settings menu that features core privacy settings in a single place.

We are constantly improving and iterating on these controls and education to provide a better experience for people. We regularly provide people with notice through various channels about changes to our product, including improvements on privacy controls. We are always working to improve our controls and do not view this as something that is ever likely to be finished.

- **Do you feel that you've done enough to ensure users understand how and when their data is being collected and used?**

We believe that it's important to communicate with people about the information that we collect and how people can control it. This is why we work hard to provide this information to people in a variety of ways: in our Data Policy, and in Privacy Basics, which provides walkthroughs of the most common privacy questions we receive. Beyond simply disclosing our practices, we also think it's important to give people access to their own information, which we do through our Download Your Information and Access Your Information tools, Activity Log, and Ad Preferences, all of which are accessible through our Privacy Shortcuts tool. We also provide information about these topics in context as people are using the Facebook service itself.

Facebook seeks, as much as possible, to put controls and information in context within its service. While "up front" information like that contained in the terms of service are useful, research overwhelmingly demonstrates that in-product controls and education are the most meaningful to people and the most likely to be read and understood. On-demand controls are also important, and we recently redesigned our entire settings menu on mobile devices from top to bottom to make things easier to find. We also created a new Privacy Shortcuts menu where users can control their data in just a few taps, with clearer explanations of how our controls work. The experience is now clearer, more visual, and easy-to-find.

Improving people's understanding of how digital services work is an industry-wide challenge that we are highly committed to addressing. That's why we have run a series of design

workshops called "Design Jams," bringing together experts in design, privacy, law and computer science to work collaboratively on new and innovative approaches. These workshops have run in Paris, London, Dublin, Berlin, Sao Paolo, Hong Kong, and other cities, and included global regulators and policymakers. At these workshops, expert teams use "people centric design" methods to create innovative new design prototypes and experiences to improve transparency and education in digital services. These workshops inform Facebook's constantly-improving approach.

In recognition of the need for improved approaches to data transparency across all digital services, working with partners from academia, design, and industry we recently launched TTC Labs, a design innovation lab that seeks to improve user experiences around personal data. TTC Labs is an open platform for sharing and innovation and contains insights from leading experts in academia, design and law, in addition to prototype designs from the Design Jams, template services and open-source toolkits for people-centric design for transparency, trust and control of data. Working collaboratively, and based on open-source approaches, TTC Labs seeks to pioneer new and more people-centric best practices for people to understand how their data is used by digital services, in ways that they find easy to understand and control. Facebook is highly committed to improving people's experience of its own services as well as investing in new innovations and approaches to support improvements across the industry

- **What additional measures might you undertake to increase awareness of Facebook's collection and use of data?**

We believe that it's important to communicate with people about the information that we collect and how people can control it and we are always working to do better. We've heard loud and clear that privacy settings and other important tools were too hard to find and that we must do more to keep people informed. So, we've taken additional steps to put people more in control of their privacy. For instance, we redesigned our entire settings menu on mobile devices from top to bottom to make things easier to find. We also created a new Privacy Shortcuts in a menu where users can control their data in just a few taps, with clearer explanations of how our controls work. The experience is now clearer, more visual, and easy-to-find. Furthermore, we also updated our Terms of Service that include our commitments to everyone using Facebook. We explain the services we offer in language that's easier to read. We also updated our Data Policy to better spell out what data we collect and how we use it in Facebook, Instagram, Messenger, and other products.

Our Download Your Information or "DYI" tool is Facebook's data portability tool and was launched many years ago to let people access and download many types of information that we maintain about them. The data in DYI and in our Ads Preferences tool contain each of the interest categories that are used to show people ads, along with information about the advertisers that are currently running ads based on their use of an advertiser's website or app. People also can choose not to see ads from those advertisers. We recently announced expansions to Download Your Information, which, among other things, make it easier for people to see their data, delete it, and easily download and export it. More information is available at https://newsroom.fb.com/news/2018/04/new-privacy-protections/.

Responding to feedback that we should do more to provide information about websites and apps that send us information when people use them, we also announced plans to build Clear History. This new feature will enable users to see the websites and apps that send us information when they use them, disassociate this information from their account, and turn off Facebook's ability to store it associated with their account going forward.

Facebook allows people to view, manage, and remove the apps that they have logged into with Facebook through the App Dashboard. We recently prompted everyone to review their App Dashboard as a part of a Privacy Checkup, and we also provided an educational notice on Facebook to encourage people to review their settings. More information about how users can manage their app settings is available at https://www.facebook.com/help/218345114850283?helpref=about_content.

We have also introduced Access Your Information. This feature provides a new way for people to access and manage their information. Users can go here to delete anything from their timeline or profile that they no longer want on Facebook. They can also see their ad interests, as well as information about ads they've clicked on and advertisers who have provided us with information about them that influence the ads they see. From here, they can go to their ad settings to manage how this data is used to show them ads.

12. In 2016, a group of Princeton researchers revealed that Facebook was tracking users across nearly a third of the web, using sophisticated tracking techniques that were all but impossible for a user to evade.

- **Do you feel that the average Facebook user is fully aware of the amount of information that you are collecting?**

Our Download Your Information or "DYI" tool is Facebook's data portability tool and was launched many years ago to let people access and download many types of information that we maintain about them, with a focus on those types that a person may wish to use on another online service. The data in DYI includes each of the demographic and interests-based attributes we use to show or target people ads. Although we do not store this data within DYI, people can also use Ad Preferences to see which advertisers are currently running ads based on their use of an advertiser's website or app. People also can choose not to see ads from those advertisers.

We have also introduced Access Your Information. This feature provides a new way for people to access and manage their information. Users can go here to delete anything from their timeline or profile that they no longer want on Facebook. They can also see their ad interests, as well as information about ads they've clicked on and advertisers who have provided us with information about them that influence the ads they see. From here, they can go to their ad settings to manage how this data is used to show them ads.

Responding to feedback that we should do more to provide information about websites and apps that send us information when people use them, we announced plans to build Clear History. This new feature will enable users to see the websites and apps that send us information when they use them, disassociate this information from their account, and turn off Facebook's ability to store it associated with their account going forward. We are working with privacy

advocates, academics, policymakers, and regulators to get their input on our approach, including how we plan to remove identifying information and the rare cases where we need information for security purposes. We've already started a series of roundtables in cities around the world, and we're looking forward to doing more.

- **Do you think Facebook users have an understanding that their data can be collected by Facebook even when they are not on Facebook?**

Facebook does not create profiles for people without a Facebook account (whom we call "nonregistered users"). However, we do receive some information from devices and browsers that may be used by such non-registered users. For example, when people visit apps or websites that feature our technologies—such as the Facebook Like or Comment button—our servers automatically log standard browser or app records of the fact that a particular device visited the website or app. This connection to Facebook's servers occurs automatically when a device visits a website or app that contains our technologies, and is an inherent function of Internet design. Most websites and apps share this same information with multiple different third parties whenever people visit the website or app.

We also may receive additional information that the publisher of the app or website or other third party chooses to share with us, such as location information (which can be sent through our Places Graph). A developer that, for example, wants to highlight restaurants near a user of its app can send us information about a device's location along with the category "restaurants." The Places Graph will return a list of places in the "restaurant" category near the specified location, enabling the developer to show its users restaurants in their area. Facebook does not associate the information it receives through Places Graph with any person.

When a person visiting a website or using an app is a non-registered user, Facebook does not obtain information identifying that individual. We use the information we receive from these websites and apps to provide our services to the website or app, as well as for security and product improvement purposes. We require websites and apps to provide appropriate disclosures and obtain adequate consent from people when using our technologies.

We also may log basic information from the device of a non-registered user if that person visits a part of Facebook that does not require people to log in, such as a public Facebook Page. The information we log when people visit our websites or apps includes basic device and connection information—for example, device model, operating system, browser, IP address, and cookies or device identifiers. This is the same information that any provider of an online service would receive when a device visits its website.

Finally, Facebook may log certain information about devices on which Facebook apps are installed, including before people using those devices have registered for Facebook (such as when a user downloads a Facebook app, but has not yet created an account, or if the app is preloaded on a given device). This information includes information such as device model, operating system, IP address, app version, and device identifiers. We use this information in order to, for example, provide the right version of the app, help people who want to create accounts (for example, by optimizing the registration flow for the specific device), retrieve bug fixes, and measure and improve app performance.

13. "Dark patterns" are user interfaces that have been intentionally designed to sway users towards taking actions they would otherwise not take under effective, more informed consent questions. This is a particular challenge when users are pushed to generally agree to default options, which typically include more expansive data sharing than perhaps previously understood.

- **Do you believe Facebook engages in these types of dark pattern practices?**

We invest heavily in ensuring people understand the choices and controls we give them over their data. Our approach complies with the law, follows recommendations from privacy and design experts, and is designed to help people understand how the technology works and their choices.

To that end, the choices we gave people were written in both "short form" and "long form" notice to help people understand what they were saying yes or no to. We also encouraged people to review our updated Data Policy and Cookies Policy, providing a short summary of the key changes, as well as gave people the choice to accept our new Terms of Service to keep using Facebook. We are not aware of any other service going to such lengths to ensure that people understood what was being asked of them.

Improving people's understanding of how digital services work is an industry-wide challenge that we are highly committed to addressing. That's why we have run design workshops called "Design Jams," bringing together experts in design, privacy, law, and computer science to work collaboratively on new and innovative approaches. We ran these workshops in cities around the world and included global regulators and policymakers. At these workshops, expert teams use "people centric design" methods to create innovative new design prototypes and experiences to improve transparency and education in digital services. These workshops inform Facebook's constantly-improving approach.

In recognition of the need for improved approaches to data transparency across all digital services, working with partners from academia, design, and industry, we recently launched TTC Labs, a design innovation lab that seeks to improve user experiences around personal data. TTC Labs is an open platform for sharing and innovation and contains insights from leading experts in academia, design and law, in addition to prototype designs from the Design Jams, template services and open-source toolkits for people-centric design for transparency, trust and control of data. Working collaboratively, and based on open-source approaches, TTC Labs seeks to pioneer new and more people-centric best practices for people to understand how their data is used by digital services, in ways that they find easy to understand and control.

Facebook is highly committed to improving people's experience of its own services as well as investing in new innovations and approaches to support improvements across the industry.

14. We need to ensure that vulnerable users around the globe are able to maintain anonymity. We also need to ensure that fake accounts aren't attacking our democracy from St. Petersburg.

- **How might we think about requiring more authentication while still protecting privacy and protecting anonymity for individuals operating within oppressive regimes around the globe?**

Facebook was built for conversation and human connection. It's why we require that people using our service provide accurate information about who they are—whether it's an individual, a business or a nonprofit. However, we also recognize that while people want to connect, they may not want to share everything with everyone. This is why we provide people with controls that let them decide what information they want to share with whom.

Of course, there is always a balance to strike between protecting people's privacy and ensuring the integrity of our platform. We recently announced that people who manage Pages with large numbers of followers will need to be verified. Those who manage large Pages that do not clear the process will no longer be able to post. This will make it much harder for people to administer a Page using a fake account, which is strictly against our policies. We will also show people additional context about Pages to help people have more information to evaluate their content. For example, you can see whether a Page has changed its name.

15. Facebook actively helps political leaders and candidates develop their social media presence and following. Such assistance has worked with a wide assortment of political leaders, including Indian Prime Minister Narendra Modi, Filipino leader Rodrigo Duterte's campaign, the Alternative for Germany party in Germany, and many others.

- **How does Facebook determine with which candidates it is willing to work?**

We want all candidates, groups, and voters to use our platform to engage in elections. We want it to be easy for people to find, follow, and contact their elected representatives—and those running to represent them. We are focused on providing the same information to all elected officials and political campaigns via our revamped website at http://politics.fb.com/.

16. Some political advocacy from certain political organizations utilize what outside experts and observers might classify as hate speech, which Facebook's community standards currently ban.

- **Does Facebook apply different community standards for advertisers or political parties than it applies for regular users?**

Every day, people come to Facebook to share their stories, see the world through the eyes of others, and connect with friends and causes. The conversations that happen on Facebook reflect the diversity of a community of more than two billion people communicating across countries and cultures and in dozens of languages, posting everything from text to photos and videos.

We recognize how important it is for Facebook to be a place where people feel empowered to communicate, and we take our role in keeping abuse off our service seriously. That's why we have developed a set of Community Standards that outline what is and is not allowed on Facebook. Our Standards apply equally around the world to all types of content from all users—including advertisers and political parties. They're designed to be comprehensive—for example, content that might not be considered hate speech may still be removed for violating our bullying policies.

However, at times we will allow content that might otherwise violate our standards if we feel that it is newsworthy, significant, or important to the public interest. We do this only after weighing the public interest value of the content against the risk of real-world harm.

The goal of our Community Standards is to encourage expression and create a safe environment. We base our policies on input from our community and from experts in fields such as technology and public safety. We update our Community Standards regularly.

In addition, our Advertising Policies (https://business.facebook.com/policies/ads/prohibited_content) apply to all users who advertise on Facebook. Besides our Community Standards, there are additional restrictions placed on ads as well. Our ads policies prohibit certain content like illegal products and services, tobacco products, drugs and drug-related products, adult products and services, and adult content among other things. We also allow, but have restrictions on, certain content like alcohol, dating, state lotteries, and subscription services, among others.

17. Until 2014, reports suggest that Facebook allowed "friend permission," which meant that if one of your Facebook friends connected an authorized app to his Facebook account, the app could access not only that person's personal information, but also your personal information – and all of his other friends' personal information – regardless of his friends' privacy settings. According to press reporting, Facebook rightly changed that permission in 2014.

- **Is that accurate?**

In April 2014, we announced that we would more tightly restrict our public platform policies and APIs to prevent abuse. At that time, we made clear that existing apps would have a year to transition—at which point they would be forced (1) to migrate to the more restricted API and (2) be subject to Facebook's new review and approval protocols. The vast majority of companies were required to make the changes by May 2015, but we granted a small number of short term extensions to developers on our public platform.

- **While "friend permission" was in effect, how many third-party entities were authorized to collect friends' data?**

We are in the process of investigating apps that had access to a large amount of information before we changed our platform in 2014. The first phase of our investigation involves reviewing apps that had access to large amounts of Facebook data prior to the changes we made to our public platform in 2014, described above. A large team comprised of internal

and external experts is undertaking (1) a comprehensive review to identify every app that had access to this amount of Facebook data and (2) where we have concerns, we are conducting interviews, sending requests for information to developers, and/or performing audits to understand how data is stored and used by a developer. Where we find evidence that these or other apps did misuse data in violation of our policies, we will ban them and let people know.

- **Do you know what happened to that data and whether it was shared further?**

  See Response to above Question.

- **Do you have an estimate of the number of users (not just the 87 million users affected by the Cambridge Analytica episode) whose data has been shared in an unauthorized way by third-party applications?**

  See Response to above Question.

- **How is Facebook prepared to remedy the harms created by those episodes of unauthorized access?**

  See Response to above Question.

- **How does Facebook audit third-party applications to ensure that they are who they say they are?**

In general, on an ongoing basis, Facebook proactively reviews all apps seeking access to more than basic information through our public platform (and have rejected more than half of apps seeking such extended permissions). We also do a variety of manual and automated checks to ensure compliance with our policies and a positive experience for users. These include steps such as random checks of existing apps along with the regular and proactive monitoring of apps. We also respond to external or internal reports and investigate for potential app violations of our policies. When we find evidence of or receive allegations of violations, we investigate and, where appropriate, employ a number of measures, including restricting applications from our platform, preventing developers from building on our platform in the future, and taking legal action where appropriate.

- **Under Facebook's new policies, what information can app developers acquire about an app user?**

The App Review process introduced in 2014 required developers who create an app that asks for more than certain basic user information through our public platform to justify the data they are looking to collect and how they are going to use it. Facebook then reviews whether the developer has a legitimate need for the data in light of how the app functions. Only if approved following such review can the app ask for a user's permission to get their data. Facebook has rejected more than half of the apps submitted for App Review between April 2014 and April 2018.

We are further updating this process, so that the only data that an app can request through our public platform without App Review will include name, profile photo, and email address.

Requesting any other data will require approval from Facebook. We also no longer allow apps to ask for access to information like religious or political views, relationship status and details, custom friends lists, education and work history, fitness activity, book reading and music listening activity, news reading, video watch activity, and games activity. We will encourage people to manage the apps they use. We already show people what apps their accounts are connected to and allow them to control what data they have permitted those apps to use. But we are making it even easier for people to see what apps they use and the information they have shared with those apps.

- **What information can app developers acquire about that user's friends?**

  See Response to above Question.

- **Do users have a way of tracking what data about them was shared with third-parties, including when this data is shared by their friends? Should they?**

Facebook allows people to view, manage, and remove the apps that they have logged into with Facebook through the App Dashboard. We recently prompted everyone to review their App Dashboard as a part of a Privacy Checkup, and we also provided an educational notice on Facebook to encourage people to review their settings. More information about how users can manage their app settings is available at https://www.facebook.com/help/218345114850283?helpref=about_content.

The categories of information that an app can access is clearly disclosed before the user consents to use an app on Facebook public platform. Users can view and edit the categories of information that apps they have used have access to through the App Dashboard.

18. Security researchers found that the applications included in the device manufacturer partnerships did not respect the privacy setting which prevents third-party access to data.

- **Did Facebook ever notify users that their data was being accessed in spite of this setting, and did you note this to the Federal Trade Commission?**

- **Approximately how many users did the applications that were granted this special access have, in total?**

Facebook's device integration partnerships are fundamentally different from the relationships that Facebook has with other developers that use our public platform to build third-party apps for consumers or businesses. The purpose of device integration partnerships was to build Facebook integrations for devices, operating systems, and other products where we and our partners wanted to offer people a way to receive Facebook or Facebook experiences, but where Facebook relied on a partner to build those experiences rather than doing so directly. By contrast, third-party app developers use the information they receive to build their own experiences.

For integration partners, people's privacy settings—namely the audience controls that people use to decide who can see the information they share on Facebook—applied whether people used a version of Facebook built by Facebook, or whether they used a version built by a

partner under an approved device integration. However, app settings that restricted information from being shared with third-party apps (including third-party apps used by friends) generally did not apply to integration partners, because the integrations they built were not third-party apps and instead offered core Facebook experiences.

Likewise, the obligations imposed by the FTC 2012 Consent Order on Facebook's use of service providers, such as these device integration partners, differ materially from those imposed on Facebook with respect to third parties. Indeed, the Consent Order excludes service providers from its definition of "third parties." Facebook's data policies—at least since 2010—have likewise informed users that Facebook works with other companies to provide its services in different contexts.

Finally, with respect to your question about the FTC, Facebook takes its obligations under the Consent Order very seriously, and discussed its device integration partnerships with the FTC both before and after the Consent Order was issued.

19. A major concern I had in 2013 with Facebook's widely reported "mood study" was the lack of informed consent by users.

- **Does Facebook provide for individualized, informed consent in all instances, including all cases where groups of users are exposed to novel interfaces or services not available to other users?**

In our Data Policy, we explain that we may use the information we have to conduct and support research in areas that may include general social welfare, technological advancement, public interest, health, and well-being. Researchers are subject to strict restrictions regarding data access and use as part of these collaborations.

Users do not have the ability to opt out of such research; however, we disclose our work with academic researchers in our Data Policy, and our work with academics is conducted subject to strict privacy and research protocols.

- **Does Facebook conduct user research into user comprehension of their options on Terms of Service consent screens or other locations where those Terms are located, and/or does it track the consent rates on those pages where Terms are shown and consent is requested?**

We do extensive research around our product and privacy features, including focus-groups and on platform surveys. Our research, consistent with extensive academic research, overwhelmingly demonstrates that in-product controls and education are the most meaningful to people and the most likely to be read and understood. On-demand controls are also important, and we recently redesigned our entire settings menu on mobile devices from top to bottom to make things easier to find. We also created a new Privacy Shortcuts, a menu where people can control their data in just a few taps, with clearer explanations of how our controls work. The experience is now clearer, more visual, and easy-to-find.

Improving people's understanding of how digital services work is an industry-wide challenge that we are highly committed to addressing. That's why we have run a series of design

workshops called "Design Jams," bringing together experts in design, privacy, law, and computer science to work collaboratively on new and innovative approaches. These workshops have run in cities around the world, and included global regulators and policymakers. At these workshops, expert teams use "people centric design" methods to create innovative new design prototypes and experiences to improve transparency and education in digital services. These workshops inform Facebook's constantly-improving approach.

In recognition of the need for improved approaches to data transparency across all digital services, working with partners from academia, design, and industry we recently launched TTC Labs, a design innovation lab that seeks to improve user experiences around personal data. TTC Labs is an open platform for sharing and innovation and contains insights from leading experts in academia, design, and law, in addition to prototype designs from the Design Jams, template services and open-source toolkits for people-centric design for transparency, trust and control of data. Working collaboratively, and based on open-source approaches, TTC Labs seeks to pioneer new and more people-centric best practices for people to understand how their data is used by digital services, in ways that they find easy to understand and control.

Facebook is highly committed to improving people's experience of its own services as well as investing in new innovations and approaches to support improvements across the industry.

20. Please provide the results of any user research Facebook has conducted into the user comprehension and consent rates of the set of consent screens created to comply with Europe's General Data Protection Regulations and released globally in 2018 (including the New Terms of Service, Data With Special Protections, Face Recognition, and Ads Based on Data from Partners, and Parental Consent screens).

In designing the GDPR roll out, like all product roll outs, we rely on design principles and research derived from numerous sources, including user research and academic research, to develop experiences that are engaging and useful for the broadest number of people. We also conducted cross-disciplinary workshops, called "Design Jams," with experts around the world to collect input on user interaction principles that would inform our work. We have learned from our work and other design research in the field that people are less likely to make informed or thoughtful decisions when bombarded with many different choices in succession. To avoid so-called "notice fatigue," we streamlined the number of data choices people are presented with as part of the GDPR roll out to 2-3 choices (depending on the user's existing settings), responding to early testing of a version with several additional choices, which the people who tested this version did not like. We also used a layered approach that gave people the information needed to make an informed choice on the first screen, while enabling ready access to deeper layers of information and settings for those interested in a particular topic. It's important to us that people have the information they need to make the privacy choices that are right for them. At this time we are not able to share specific information regarding user research and testing, but will continue to monitor how these and other privacy settings perform with users.

- **If there were multiple iterations of designs for any of the screens, please include the results for each iteration that was tested.**

Below are screenshots of the consent flows being provided in Europe:

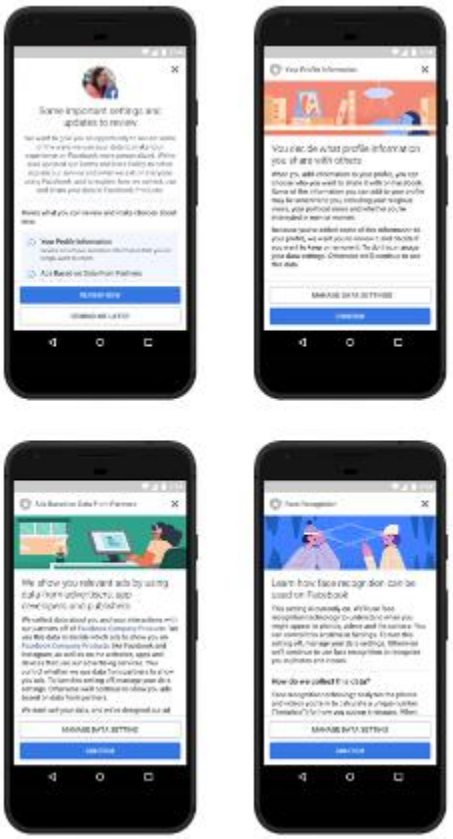

21. Facebook recently took some actions to address the horrific events unfolding in Myanmar by banning some of Myanmar's military leadership from the Facebook platform. However, the publication Wired reported that since at least May 2015, Facebook was aware of its platform's capacity to foment violence in Myanmar.

- **Is that accurate?**

- **Why didn't you take action earlier?**

- **How much are you investing in addressing the misinformation and violence prevention issues in Myanmar?**

- **What about in other parts of the world where similar threats are possible?**

We were too slow to respond to the concerns raised by civil society, academics and other groups in Myanmar. We don't want Facebook to be used to spread hatred and incite violence. This is true around the world, but it is especially true in Myanmar where our services can be used

to amplify hate or exacerbate harm against the Rohingya. There are challenges, which are unique to Myanmar, and we are focused on addressing them through a combination of people, technology, policies, and programs. One challenge is the fact that harmful content is not always reported to us, which means we can't rely on content reports and reviewers alone to solve the problem. That's why in the last year we have established a team of product, policy, and operations experts to roll out better reporting tools, a new policy to tackle misinformation that has the potential to contribute to offline harm, faster response times on reported content, and improved proactive detection of hate speech. There is more we need to do and we will continue to invest in Myanmar to do better.

22. Press reports have suggested that Russian trolls have targeted American military personnel and U.S. military veterans on Facebook with disinformation campaigns. In August 2017, the nonprofit Vietnam Veterans of America (VVA) discovered a Facebook page bearing its name, logo, and registered trademark that was not affiliated with the organization and whose posts linked to "vvets.eu"—a website anonymously registered through Netfinity JSC of Bulgaria. The page shared divisive political content, including posts about the NFL "Take a Knee" boycott controversies and the racially charged "Blue Lives Matter" movement. The page had nearly 200,000 followers by October 2017, according to VVA, but was not shut down when the organization first flagged it to a Facebook representative on August 23, 2017. It took months for Facebook to pull down this account.

- **Why did Facebook ultimately take action against this account? Why not earlier?**

We are aware that threat actors seek to leverage social media to target military personnel, including impersonating members of the public who are more likely to be considered trustworthy—such as members of the military, veterans, and other professionals. We recognize this and are working to combat impersonation in a variety of ways.

On October 24, we removed this Page after receiving a valid IP report claiming infringement from the rights owner.

- **Have you seen attempts to target U.S. military or U.S. veterans with disinformation?**

We are aware that threat actors seek to leverage social media to target military personnel. We have a threat intelligence team dedicated to countering these sorts of cybersecurity threats, and we are expanding that team along with other teams that work on safety and security at Facebook. The security features on Facebook that protect people from these threats are equally available to members of the military. For example, we suggest performing a security checkup, and we have systems that aim to prevent malicious files from being uploaded or shared on Facebook. In addition, we partnered with Blue Star Families and USAA to create an online safety guide specifically for service members and their families—and released a video PSA (https://www.facebook.com/FBMilVetCommunity/videos/1655416797877942/) to help people identify and report military scams. We regularly train and advise military officials on best practices for maintaining secure accounts and Pages, which include setting up two-factor authentication and managing Page Roles. And of course, military personnel, like all Facebook users, have the ability to control who sees their posts and other information.

- **What are you doing to ensure our military and our veterans are protected against this type of attack?**

See Response to above Question.

23. What is Facebook's current policy on the posting or promotion of hacked emails on 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. We do this only after weighing the public interest value of the content against the risk of real-world harm.

24. Europe has established new rules for data protection and privacy for European citizens (General Data Protection Regulation, or GDPR). These new rules include required data portability, the right to be forgotten online, a 72-hour data breach disclosure requirement, and first-party consent requirements.

- **How is Facebook complying with GDPR?**
    - **Are there protections that will flow to U.S. users as a result?**

As a part of our overall approach to privacy, we are providing the same tools for access, rectification, erasure, data portability, and others to people in the US (and globally) that we provide in the European Union under the GDPR. The controls and settings that Facebook is enabling as part of the GDPR include settings for controlling our use of facial recognition technology on Facebook and for controlling our ability to use data we collect off Facebook Company Products to show users relevant ads. We recently provided direct notice of these controls and our updated Terms to people around the world (including in the US), allowing them to choose whether or not to enable or disable these settings or to agree to our updated Terms. Many of these tools (like Download Your Information, which is Facebook's data portability tool; Ad Preferences; and Activity Log) have been available globally for many years.

The substantive protections in our user agreements offered by Facebook Ireland (where our European headquarters are located) and Facebook, Inc. are the same. However, there are certain aspects of our Facebook Ireland Data Policy that are specific to legal requirements in the GDPR—such as the requirement that we provide contact information for our EU Data Protection Officer or that we identify the "legal bases" we use for processing data under the GDPR. Likewise, our Facebook Ireland Terms and Data Policy address the lawful basis for transferring data outside the EU, based on legal instruments that are applicable only to the EU. And other provisions of the GDPR itself pertain to interactions between European regulators and other matters that are not relevant to people located outside of the EU.

We offered choice and obtained explicit consent through user engagement flows from people in Europe to three specific uses of data: facial recognition data (which previously was not enabled in Europe), special categories of data, and use of data we collect off Facebook Company Products to show users relevant ads. As noted above, we provided direct notice of these controls

and our updated Terms to people around the world (including in the US), allowing people to choose whether or not to enable or disable these settings or to agree to our updated Terms. Outside of Europe did not ask people to agree to facial recognition if they previously disabled it; in contrast, facial recognition was not previously available in Europe so more people there were asked. Also, we are not requiring people to complete those flows if they repeatedly indicate that they do not want to go through the experience. At the same time, the events of recent months have underscored how important it is to make sure people know how their information is used and what their choices are. So, we decided to communicate prominently on Facebook—through a full-screen message and a reminder to review at a later date. People can choose to dismiss or ignore these messages and continue using Facebook.

- **What lessons should we be learning from the European experiment with data protection?**

The GDPR is founded on core principles of accountability, transparency, and control, which are also central values we employ in designing our products. The controls and settings that Facebook is promoting as part of the GDPR are available to people around the world, including settings controlling our ability to use data we collect off Facebook Company Products to target ads. We provide the same tools for access, rectification, erasure, data portability, and others to people in the US and the rest of the world that we provide in Europe, and many of those tools (like our Download Your Information tool, Ad Preferences, and Activity Log) have been available globally for many years.

We support the GDPR's emphasis on transparency, choice and control, and its recognition that, while a consent requirement is appropriate in some cases (such as the processing of special category data), other legal frameworks may be appropriate in other circumstances, such as where a company has a "legitimate interest" in processing data, where processing data is necessary to perform a contract, or where data processing serves the broader public interest.

In this way, the GDPR provides strong protections for data that may be processed for different reasons and seeks to avoid over-burdening consumers with consent requests for every processing of data, which could increase what experts call "notice fatigue" and cause people to pay less attention to the privacy notices they receive.

- **Should we consider policy solutions like first-party consent?**

We support the GDPR's emphasis on transparency, choice and control, and its recognition that, while a consent requirement is appropriate in some cases (such as the processing of special category data), other legal frameworks may be appropriate in other circumstances, such as where a company has a "legitimate interest" in processing data, where processing data is necessary to perform a contract, or where data processing serves the broader public interest.

In this way, the GDPR provides strong protections for data that may be processed for different reasons and seeks to avoid over-burdening consumers with consent requests for every

processing of data, which could increase what experts call "notice fatigue" and cause people to pay less attention to the privacy notices they receive.

We support models for consent that ensure companies are able to design consent experiences that are intuitive and enhance people's ability to make an informed choice.

- **Why shouldn't companies be required to obtain explicit and informed consent before collecting or processing user data like in Europe?**

GDPR does not require consent for most uses of personal information, and instead, recognizes that many uses of data are necessary to provide a service or within a company's legitimate interests or contractual necessity. Similarly, the FTC's guidance recognizes that people's expectations vary based on the context in which their information was collected and based on their relationship with an organization that holds their data. Consistent with that distinction, the FTC agrees with the GDPR perspective that consent may be appropriate in some situations but is not suitable for every single processing of data.

Likewise, the GDPR does not differentiate between users and non-users, and indeed, many online or digital services around the world do not require registration or distinguish between "users" and "non-users" before collecting or logging data, such as browser logs of people who visit their website.

We agree that different levels of consent or notice are appropriate depending on the type of information or contemplated use at issue. We also support the GDPR's emphasis on transparency, choice and control, and its recognition that, while a consent requirement is appropriate in some cases (such as the processing of sensitive data), other legal frameworks may be appropriate in other circumstances, such as where a company has a "legitimate interest" in processing data, where processing data is necessary to perform a contract, or where data processing serves the broader public interest.

In this way, the GDPR provides strong protections for personal data that may be processed for different reasons and avoids over-burdening people who use our service with consent requests for every processing of data, which could increase what experts call "notice fatigue" and cause people to pay less attention to the privacy notices they receive.

We support models for consent that ensure companies are able to design consent experiences that are intuitive and enhance people's ability to make an informed choice.

25. Do you think Facebook might benefit from more independent insight into anonymized activity?

- **Isn't there a public interest in better understanding how your platform works and how users interact on social media?**

We are working with the broader community to identify and combat threats. One example is our partnership with the Atlantic Council's Digital Forensic Research Lab, which is providing us with real-time updates on emerging threats and disinformation campaigns around the world. They assisted in our work around the Mexico election, our recent takedown of a

financially motivated "like" farm in Brazil, and the accounts we recently disabled for coordinated inauthentic behavior here in the US.

Another example is that Facebook recently announced a new initiative to help provide independent, credible research about the role of social media in elections, as well as democracy more generally. It will be funded by the Laura and John Arnold Foundation, Democracy Fund, the William and Flora Hewlett Foundation, the John S. and James L. Knight Foundation, the Charles Koch Foundation, the Omidyar Network, and the Alfred P. Sloan Foundation. At the heart of this initiative will be a group of scholars who will:

- Define the research agenda;
- Solicit proposals for independent research on a range of different topics; and
- Manage a peer review process to select scholars who will receive funding for their research, as well as access to privacy-protected datasets from Facebook which they can analyze.

Facebook will not have any right to review or approve their research findings prior to publication. More information regarding the study is available at https://newsroom.fb.com/news/2018/04/new-elections-initiative/.

In addition, we regularly work with privacy experts outside the company, including academics, to understand how to improve privacy protections for people on Facebook and to support efforts to improve privacy protections for people overall. For example, we recently hosted a workshop for privacy academics to discuss research around online privacy and worked with academics as a part of recent privacy consultations that we have conducted at our headquarters and around the world.

26. The fact that Facebook failed to anticipate misuse is extremely troubling.

- **Why should we have confidence that you are any more prepared to handle issues of misuse now?**
- **How are you better protecting the users of your products?**
- **You have indicated that Facebook is now more fully addressing potential threats to new products *before* launching them.**
    - **Why was this not a part of Facebook's process previously?**

In the run-up to the 2016 elections, we were focused on the kinds of cybersecurity attacks typically used by nation states, for example phishing and malware attacks. And we were too slow to spot this type of information operations interference. Since then, we've made important changes to help prevent bad actors from using misinformation to undermine the democratic process.

Protecting a global community of more than 2 billion people involves a wide range of teams and functions, and our expectation is that those teams will grow across the board. For example, we have dedicated information security and related engineering teams.

Protecting the security of information on Facebook is at the core of how we operate. Security is built into every Facebook product, and we have dedicated teams focused on each aspect of data security. From encryption protocols for data privacy to machine learning for threat detection, Facebook's network is protected by a combination of advanced automated systems and teams with expertise across a wide range of security fields. Our security protections are regularly evaluated and tested by our own internal security experts and independent third parties. For the past 7 years, we have also run an open bug bounty program that encourages researchers from around the world to find and responsibly submit security issues to us so that we can fix them quickly and better protect the people who use our service.

We anticipate continuing to grow these teams by hiring a range of experts, including people with specific types of threat intelligence expertise.

This will never be a solved problem because we're up against determined, creative, and well-funded adversaries. But we are making steady progress. Here is a list of 10 important changes we have made:

- **Ads and Pages transparency.** Advertising should be transparent: users should be able to see all the ads an advertiser is currently running on Facebook, Instagram and Messenger. And for ads with political or issue content, we've created an archive that will hold ads with political or issue content for 7 years—including information about ad impressions and spend, as well as demographic data such as age, gender, and location. And people everywhere can see all the ads that Page is running on Facebook. We also announced in April that people who manage Pages with large numbers of followers will need to be verified. Those who manage large Pages that do not clear the process will no longer be able to post. This will make it much harder for people to administer a Page using a fake account, which is strictly against our policies. We will also show people additional context about Pages to help people have more information to evaluate their content. For example, you can see whether a Page has changed its name.

- **Verification and labeling.** Every advertiser will now need to confirm their ID and location before being able to run any ads with political or issue content in the US and certain other countries. All ads with political or issue content will also clearly state who paid for them.

- **Updating targeting.** We want ads on Facebook to be safe and civil. We thoroughly review the targeting criteria advertisers can use to ensure they are consistent with our principles. As a result, we removed nearly one-third of the targeting segments used by the IRA. We continue to allow some criteria that people may find controversial. But we do see businesses marketing things like historical books, documentaries, or television shows using them in legitimate ways.

- **Better technology.** Over the past year, we've gotten increasingly better at finding and disabling fake accounts. We now block millions of fake accounts each day as people try to create them—and before they've done any harm. This is thanks to improvements in machine learning and artificial intelligence, which can proactively identify suspicious behavior at a scale that was not possible before—without needing to look at the content itself.

- **Action to tackle fake news.** We are working hard to stop the spread of false news. We work with third-party fact-checking organizations to limit the spread of articles rated false. To reduce the spread of false news, we remove fake accounts and disrupt economic incentives for traffickers of misinformation. We also use various signals, including feedback from our community, to identify potential false news. In countries where we have partnerships with independent third-party fact-checkers, stories rated as false by those fact-checkers are shown lower in News Feed. If Pages or domains repeatedly create or share misinformation, we significantly reduce their distribution and remove their advertising rights. We also want to empower people to decide for themselves what to read, trust, and share. We promote news literacy and work to inform people with more context. For example, if third-party fact-checkers write articles about a news story, we show them immediately below the story in the Related Articles unit. We also notify people and Page Admins if they try to share a story, or have shared one in the past, that's been determined to be false. In addition to our own efforts, we're learning from academics, scaling our partnerships with third-party fact-checkers and talking to other organizations about how we can work together.

- **Significant investments in security.** As part of our larger company investment in the space, we have more than doubled the number of people working on safety and security and now have over 20,000. We expect these investments to impact our profitability. But the safety of people using Facebook needs to come before profit.

- **Industry collaboration.** Recently, we joined more than 60 global tech and security companies in signing a TechAccord pact to help improve security for everyone.

- **Information sharing and reporting channels.** In the 2017 German elections, we worked closely with the authorities there, including the Federal Office for Information Security (BSI). This gave them a dedicated reporting channel for security issues related to the federal elections.

- **Tracking 40+ elections.** We deployed new tools and teams to proactively identify threats in the run-up to specific elections. We first tested this effort during the Alabama Senate election, and have continued these efforts for elections around the globe, including the US midterms. Last year we used public service announcements to help inform people about fake news in 21 separate countries, including in advance of French, Kenyan and German elections.

- **Action against the Russia-based IRA.** In April, we removed 70 Facebook and 65 Instagram accounts—as well as 138 Facebook Pages—controlled by the IRA primarily targeted either at people living in Russia or Russian-speakers around the

world including from neighboring countries like Azerbaijan, Uzbekistan, and Ukraine. The IRA has repeatedly used complex networks of inauthentic accounts to deceive and manipulate people in the US, Europe, and Russia—and we don't want them on Facebook anywhere in the world. In July, we removed 32 Pages and accounts from Facebook and Instagram that were engaged in coordinated inauthentic behavior. These Pages had some links to previously removed IRA-affiliated accounts, but we were unable to determine whether this new cluster of activity was directly controlled by the IRA. Our security teams are continuing to monitor our platform for abuse in connection with future elections here and around the world.

27. At our most recent public hearing with experts on social media, all of our witnesses opined that Russian influence operations are *ongoing and currently using* several social media platforms, including Facebook.

- **Do you believe that the Russian-linked operatives continue to utilize Facebook for information operations to undermine our democracy?**

Facebook has conducted a broad search for evidence that Russian actors, not limited to the IRA or any other specific entity or organization, attempted to interfere in the 2016 election by using Facebook's advertising tools. We found coordinated activity that we now attribute to the IRA, despite efforts by these accounts to mask the provenance of their activity. We have used the best tools and analytical techniques that are available to us to identify the full extent of this malicious activity, and we continue to monitor our platform for abuse and to share and receive information from others in our industry about these threats.

In April, we removed 70 Facebook and 65 Instagram accounts—as well as 138 Facebook Pages—controlled by the IRA primarily targeted either at people living in Russia or Russian-speakers around the world including from neighboring countries like Azerbaijan, Uzbekistan, and Ukraine. The IRA has repeatedly used complex networks of inauthentic accounts to deceive and manipulate people in the US, Europe, and Russia—and we don't want them on Facebook anywhere in the world.

In July, we removed 32 Pages and accounts from Facebook and Instagram that were engaged in coordinated inauthentic behavior. These Pages had some links to previously removed IRA-affiliated accounts, but we were unable to determine whether this new cluster of activity was directly controlled by the IRA.

In August, we removed Pages, groups, and accounts that were linked to sources the US government had previously identified as Russian military intelligence services. This cluster was focused on politics in Syria and Ukraine. To date, we have not found activity by these accounts targeting the US. We are working with US law enforcement on this investigation.

Some state intelligence services, including Russia's, will use any medium available to conduct information operations. We continue to diligently search for their efforts to do so on our platform and will disrupt any that we find.

- **Have you seen non-IRA, Russian-linked activity on your platform conducting similar types of information operations?**

Our security teams are continuing to monitor our platform for abuse in connection with future elections here and around the world.

In August, we removed Pages, groups, and accounts that were linked to sources the US government had previously identified as Russian military intelligence services. This cluster was focused on politics in Syria and Ukraine. To date, we have not found activity by the accounts we removed in August targeting the US. We are working with US law enforcement on this investigation.

- **What percentage of Russian-linked activity do you think the IRA represents?**

Deciding when and how to publicly link suspicious activity to a specific organization, government, or individual is a challenge that governments and many companies face. Last year, we said the Russia-based Internet Research Agency (IRA) was behind much of the abuse we found around the 2016 election.

Since 2017 we've shut down Pages and accounts engaged in coordinated inauthentic behavior without saying that a specific group or country is responsible on several occasions.

Determining attribution to a specific organization or entity is challenging for a private sector company; it is especially hard without access to the type of information that governments can use to determine attribution. With the information available to us, we cannot accurately determine what percentage of Russian-linked activity the IRA represents.

- **Have you seen evidence of additional Russian-linked troll farms?**

Our security teams are continuing to monitor our platform for abuse in connection with future elections here and around the world. We have identified other actors engaged in disinformation activity, including false news campaigns run out of countries such as Macedonia and Armenia.

In August, we removed Pages, groups, and accounts that were linked to sources the US government had previously identified as Russian military intelligence services. This cluster was focused on politics in Syria and Ukraine. To date, we have not found activity by these accounts targeting the US. We are working with US law enforcement on this investigation.

- **Have you identified any troll farms backed by countries other than Russia?**

Our security teams are continuing to monitor our platform for abuse in connection with future elections here and around the world. We have identified other actors engaged in disinformation activity, including false news campaigns run out of countries such as Macedonia and Armenia.

- **Do you anticipate additional account take-downs in the weeks ahead?**

Last month, we removed 42 accounts and 11 Pages with a network we assessed to be involved in coordinated inauthentic behavior in Brazil. We also removed 15 Pages associated with coordinated inauthentic behavior ahead of the Belgian elections. On October 11, we removed 559 Pages and 251 accounts for violations of our spam policy and for coordinated inauthentic behavior. These Pages and accounts used fake profiles to drive users to ad-heavy websites in order to make money. As part of our efforts to protect elections, we are continually investigating potential threats, both targeting the United States and abroad. The pace of these investigations and take-downs is hard to predict, though we are committed to informing the public and law enforcement and government partners when we discover and disrupt these efforts. More information is available at http://newsroom.fb.com.

- **Will you commit to notifying the public should you identify other foreign influence operations?**

We have worked to notify people about foreign influence operations, broadly, starting with our white paper in April 2017, Information Operations on Facebook, and our disclosures about the IRA last fall. Since then, we have continued to publish updates on these issues in our Newsroom.

- **Will you alert users when they've been exposed to these types of operations?**

We have worked to notify people about foreign influence operations on a variety of occasions and will continue to do so as appropriate.

[From Senator Feinstein]

28. Over the last two months, Facebook has taken action against hundreds of foreign accounts conducting influence operations. However, it is concerning that in the context of the most recent examples from August 21st, action required input from the cybersecurity company FireEye – rather than Facebook finding the subject accounts exclusively through its own internal processes.

- **In the recent case of the Iranian-associated influence campaign, did an external company have to alert you to the activity; and if so, why?**

The investigation that led to the removal of 652 Pages, groups, and accounts originating in Iran in August was the result of a mixture of external assistance from FireEye, a cybersecurity firm that had identified a suspicious network of Facebook Pages and accounts on another online service, and our own internal work. While we are constantly monitoring for threats on our platform, some networks will invariably be discovered by industry partners who investigate these issues. This is precisely why we are so focused on working with academics, companies, and other experts to help identify threats.

- **What specific steps are you taking to enhance your ability to find and mitigate influence operations?**

In the run-up to the 2016 elections, we were focused on the kinds of cybersecurity attacks typically used by nation states, for example phishing and malware attacks. And we were too slow

to spot this type of information operations interference. Since then, we've made important changes to help prevent bad actors from using misinformation to undermine the democratic process.

Protecting a global community of more than 2 billion people involves a wide range of teams and functions, and our expectation is that those teams will grow across the board. For example, we have dedicated information security and related engineering teams that have grown in size and learned from investigating prior information operations on our platform.

Protecting the security of information on Facebook is at the core of how we operate. Security is built into every Facebook product, and we have dedicated teams focused on each aspect of data security. From encryption protocols for data privacy to machine learning for threat detection, Facebook's network is protected by a combination of advanced automated systems and teams with expertise across a wide range of security fields. Our security protections are regularly evaluated and tested by our own internal security experts and independent third parties. For the past 7 years, we have also run an open bug bounty program that encourages researchers from around the world to find and responsibly submit security issues to us so that we can fix them quickly and better protect the people who use our service.

We anticipate continuing to grow these teams by hiring a range of experts, including people with specific types of threat intelligence expertise.

This will never be a solved problem because we're up against determined, creative and well-funded adversaries. But we are making steady progress. Here is a list of 10 important changes we have made:

- **Ads and Pages transparency.** Advertising should be transparent: users should be able to see all the ads an advertiser is currently running on Facebook, Instagram and Messenger. And for ads with political or issue content, we've created an archive that will hold ads with political or issue content for 7 years—including information about ad impressions and spend, as well as demographic data such as age, gender, and location. And people everywhere can see all the ads that Page is running on Facebook. We also announced in April that people who manage Pages with large numbers of followers will need to be verified. Those who manage large Pages that do not clear the process will no longer be able to post. This will make it much harder for people to administer a Page using a fake account, which is strictly against our policies. We will also show people additional context about Pages to help people have more information to evaluate their content. For example, you can see whether a Page has changed its name.

- **Verification and labeling.** Every advertiser will now need to confirm their ID and location before being able to run any ads with political or issue content in the US and certain other countries. All ads with political or issue content will also clearly state who paid for them.

- **Updating targeting.** We want ads on Facebook to be safe and civil. We thoroughly review the targeting criteria advertisers can use to ensure they are consistent with our

principles. As a result, we removed nearly one-third of the targeting segments used by the IRA. We continue to allow some criteria that people may find controversial. But we do see businesses marketing things like historical books, documentaries, or television shows using them in legitimate ways.

- **Better technology.** Over the past year, we've gotten increasingly better at finding and disabling fake accounts. We now block millions of fake accounts each day as people try to create them—and before they've done any harm. This is thanks to improvements in machine learning and artificial intelligence, which can proactively identify suspicious behavior at a scale that was not possible before—without needing to look at the content itself.

- **Action to tackle fake news.** We are working hard to stop the spread of false news. We work with third-party fact-checking organizations to limit the spread of articles rated false. To reduce the spread of false news, we remove fake accounts and disrupt economic incentives for traffickers of misinformation. We also use various signals, including feedback from our community, to identify potential false news. In countries where we have partnerships with independent third-party fact-checkers, stories rated as false by those fact-checkers are shown lower in News Feed. If Pages or domains repeatedly create or share misinformation, we significantly reduce their distribution and remove their advertising rights. We also want to empower people to decide for themselves what to read, trust, and share. We promote news literacy and work to inform people with more context. For example, if third-party fact-checkers write articles about a news story, we show them immediately below the story in the Related Articles unit. We also notify people and Page Admins if they try to share a story, or have shared one in the past, that's been determined to be false. In addition to our own efforts, we're learning from academics, scaling our partnerships with third-party fact-checkers and talking to other organizations about how we can work together.

- **Significant investments in security.** As part of our larger company investment in the space, we have more than doubled the number of people working on safety and security and now have over 20,000. We expect these investments to impact our profitability. But the safety of people using Facebook needs to come before profit.

- **Industry collaboration.** Recently, we joined more than 60 global tech and security companies in signing a TechAccord pact to help improve security for everyone.

- **Information sharing and reporting channels.** In the 2017 German elections, we worked closely with the authorities there, including the Federal Office for Information Security (BSI). This gave them a dedicated reporting channel for security issues related to the federal elections.

- **Tracking 40+ elections.** We deployed new tools and teams to proactively identify threats in the run-up to specific elections. We first tested this effort during the Alabama Senate election, and have continued these efforts for elections around the globe, including the US midterms. Last year we used public service announcements

to help inform people about fake news in 21 separate countries, including in advance of French, Kenyan and German elections.

- **Action against the Russia-based IRA.** In April, we removed 70 Facebook and 65 Instagram accounts—as well as 138 Facebook Pages—controlled by the IRA primarily targeted either at people living in Russia or Russian-speakers around the world including from neighboring countries like Azerbaijan, Uzbekistan, and Ukraine. The IRA has repeatedly used complex networks of inauthentic accounts to deceive and manipulate people in the US, Europe, and Russia—and we don't want them on Facebook anywhere in the world. In July, we removed 32 Pages and accounts from Facebook and Instagram that were engaged in coordinated inauthentic behavior. These Pages had some links to previously removed IRA-affiliated accounts, but we were unable to determine whether this new cluster of activity was directly controlled by the IRA. Our security teams are continuing to monitor our platform for abuse in connection with future elections here and around the world.

29. In your statement for the record, you note that you "have more than doubled the number of people working on safety and security and now have over 20,000 people."

- **What is the number of employees Facebook has focused directly on foreign influence operations?**

We expect to have at least 250 people specifically dedicated to safeguarding election integrity on our platforms, and that number does not include the thousands of people who will contribute to this effort in some capacity. This type of abuse touches a number of different teams at Facebook. Thousands on our Business Integrity team will be working to better enforce our ad policies and to review more ads, and a significant number of engineers will build tools to identify ad and election abuse, and to enable us to follow through on our commitment to bring greater transparency to ads with political or issue content.

- **How many are Facebook employees and how many are contract employees?**

Our effort to make our platform safer and more secure is a holistic one that involves a continual evaluation of our personnel, processes, and policies, and we make changes as appropriate. To provide 24/7 coverage across dozens of languages and time zones and ensure that Facebook is a place where both expression and personal safety are protected and respected, our content review team includes a combination of employees, contractors, and vendor partners based in locations around the world. We partner with reputable vendors who are required to comply with specific obligations, including provisions for resiliency, support, transparency, and user privacy.

- **How does Facebook make prioritization decisions relative to detecting, investigating, and dealing with foreign influence operations?**

A large amount of our focus is dedicated to understanding coordinated efforts to manipulate users around democratic systems and processes, including our significant efforts around the integrity of elections. We also recognize the importance of ensuring that

conversations and interactions on Facebook are authentic at all times, so that people can trust the connections they make. We do not allow manipulation stemming from information operations on Facebook, and when we detect this behavior, we investigate and disrupt it as a matter of priority.

- **What was the protocol for bringing information operations to the attention of senior leadership at Facebook two years ago? What is the protocol today?**

Facebook has always had channels of communication for escalating matters to senior leadership. Today, we have a dedicated team of senior leaders across our company who coordinate the investigation and disruption of information operations on Facebook. When a potential information operation is discovered, that team ensures that appropriate senior leadership is informed.

30. Russia and other outside actors continue to weaponize social media platforms, Facebook included, to foment chaos and sow discord within the United States. At the Senate Intelligence Committee's August 1, 2018, open hearing, each witness assessed that Russian influence operations are *ongoing and currently using* several social media platforms, including Facebook.

- **Do you believe that the Russians continue to utilize your platform for information operations to undermine our democracy?**

Facebook has conducted a broad search for evidence that Russian actors, not limited to the IRA or any other specific entity or organization, attempted to interfere in the 2016 election by using Facebook's advertising tools. We found coordinated activity that we now attribute to the IRA, despite efforts by these accounts to mask the provenance of their activity. We have used the best tools and analytical techniques that are available to us to identify the full extent of this malicious activity, and we continue to monitor our platform for abuse and to share and receive information from others in our industry about these threats.

In April, we removed 70 Facebook and 65 Instagram accounts—as well as 138 Facebook Pages—controlled by the IRA primarily targeted either at people living in Russia or Russian-speakers around the world including from neighboring countries like Azerbaijan, Uzbekistan, and Ukraine. The IRA has repeatedly used complex networks of inauthentic accounts to deceive and manipulate people in the US, Europe, and Russia—and we don't want them on Facebook anywhere in the world.

In July, we removed 32 Pages and accounts from Facebook and Instagram that were engaged in coordinated inauthentic behavior. These Pages had some links to previously removed IRA-affiliated accounts, but we were unable to determine whether this new cluster of activity was directly controlled by the IRA.

In August, we removed Pages, groups, and accounts that were linked to sources the US government had previously identified as Russian military intelligence services. This cluster was focused on politics in Syria and Ukraine. To date, we have not found activity by these accounts targeting the US. We are working with US law enforcement on this investigation.

Some state intelligence services, including Russia's, will use any medium available to conduct information operations. We continue to diligently search for their efforts to do so on our platform will disrupt any that we find.

- **How many ongoing investigations does Facebook have underway?**

Our security teams are constantly monitoring for organized and emerging threats. While we do not publicly disclose the elements or number of these reviews for security reasons, factors include monitoring and assessing thousands of detailed attributes about accounts on Facebook, such as location information and connections to others on our platform. We are committed to keeping law enforcement apprised of our efforts and to bringing this information to the public as appropriate.

- **How many Russian-backed information operations is Facebook currently tracking? What are those operations focused on?**

See Response to above Questions. We are constantly monitoring for foreign information operations.

- **What percentage of Russian-linked activity do you think the Internet Research Agency represents?**

Deciding when and how to publicly link suspicious activity to a specific organization, government, or individual is a challenge that governments and many companies face. Last year, we said the Russia-based Internet Research Agency (IRA) was behind much of the abuse we found around the 2016 election.

But since then, we've shut down Pages and accounts engaged in coordinated inauthentic behavior without saying that a specific group or country is responsible on several occasions. Furthermore, the Russian government and intelligence services do not constrain themselves to information operations on social media. Russia's efforts to target democratic systems and processes target all levels of society, and rely just as heavily on traditional intelligence activities.

Determining attribution to a specific organization or entity is hard for a private sector company; it is especially hard to do so without access to the type of information that governments can use in determining attribution. With the information available to us, we cannot accurately determine what percentage of Russian-linked activity the IRA represents.

- **Do you anticipate additional account take-downs in the weeks ahead?**

Last month, we removed 42 accounts and 11 Pages with a network we assessed to be involved in coordinated inauthentic behavior in Brazil. We also removed 15 Pages associated with coordinated inauthentic behavior ahead of the Belgian elections. On October 11, we removed 559 Pages and 251 accounts for violations of our spam policy and for coordinated inauthentic behavior. These Pages and accounts used fake profiles to drive users to ad-heavy websites in order to make money. As part of our efforts to protect elections, we are continually investigating potential threats, both targeting the United States and abroad. The pace of these investigations and take-downs is hard to predict, though we are committed to informing the

public and law enforcement and government partners when we discover and disrupt these efforts. More information is available at http://newsroom.fb.com.

- **Do you commit to notifying the public should Facebook identify other foreign information operations?**

- **Will Facebook commit to institutionalizing the alerting of users who have been exposed to foreign information operations?**

We have worked to notify people about foreign influence operations on a variety of occasions and will continue to do so as appropriate.

31. As has been illustrated with the actions Facebook took in August, stopping Russian and Iranian-associated influence accounts requires close coordination between the government, social media companies, other private sector entities, and even the public. This construct has been useful in the past; in 2016, Facebook and other social media companies created a shared database of videos and images to counter online terrorist propaganda.

- **Do you believe there is a need for better information sharing between the social media companies?**

We agree that information sharing among companies and government is critical to combating constantly evolving cyber threats. 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 also have a history of working successfully with the DOJ, the FBI, and other law enforcement to address a wide variety of threats to our platform, and we look forward to continuing to work with law enforcement and government on these issues.

- **What is prohibiting your company from sharing more with your peers, government actors, and the public with respect to foreign information operations?**

We agree that information sharing among companies and government is critical to combating constantly evolving cyber threats. 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 also have a history of working successfully with the DOJ, the FBI, and other law enforcement to address a wide variety of threats to our platform, and we look forward to continuing to work with law enforcement and government on these issues. We'd be happy to discuss these issues further with your staff.

32. One of the major criticisms against this database countering extremist content is that there is little information about how it operates and how effective it is in preventing prohibited content from being uploaded again.

- **Have your companies agreed on a common standard for what constitutes prohibited extremist or terrorist content? If not, why not?**

- **Would a shared standard and the deployment of similar software used to detect spam and copyrighted material, facilitate the automated blocking of such content across all four platforms?**

- **In the interest of transparency, would you make this database open to the public or researchers to know which images are prohibited?**

At Facebook, we have deployed a variety of tools in the fight to find and remove content that violates our Community Standards, including artificial intelligence, specialized human review, and industry cooperation. Between January and March 2018, we took action on 1.9 million pieces of ISIS and al-Qaeda content, 99.5 percent of which we found and flagged with our technology.

At last year's EU Internet Forum, Facebook, Microsoft, Twitter, and YouTube declared our joint determination to curb the spread of terrorist content online. Over the past year, we have formalized this partnership with the launch of the Global Internet Forum to Counter Terrorism (GIFCT). The GIFCT is committed to working on technological solutions to help thwart terrorists' use of our services, including through a shared industry hash database, where companies can create "digital fingerprints" for terrorist content and share it with participating companies. The database, which became operational in the spring of 2017, now includes 13 companies that contribute to it and contains more than 88,000 hashes. It allows the thirteen member companies to use those hashes to identify and remove matching content—videos and images—that violate our respective policies or, in some cases, immediately take action on terrorist content. GIFCT also created an online resource for smaller tech companies to seek support and feedback. Each company has different policies, practices, and definitions as they relate to extremist and terrorist content. If content is removed from a company's platform for violating that platform's individual terrorism-related content policies, the company may choose to hash the content and include it in the database.

We are exploring ways to be more transparent about our efforts to combat terrorism without inadvertently further exploiting or disseminating terrorist content. A database of this kind explicitly holds content, in a hashed form, that violates not just our platform's guidelines but often US and other government's terrorist legislation. The content is often inherently disturbing and represents the worst of the worst in terms of terrorist content. We are very careful in this by-industry-for-industry effort to ensure we are not part of the further spreading of this content. We have discussed our GIFCT efforts and processes with many academics around the world, especially through the GIFCT Global Academic Network, which has 8 institutes from 7 countries on 4 continents that we consult with.

33. Will you commit to providing public access to a library of all ads that target users based on demographics? (What content, purchased by whom, targeting whom)? If not, why not?

We now require that advertisers clearly label all election-related and issue ads on Facebook and Instagram in the US—including a "Paid for by" disclosure from the advertiser at the top of the ad. This will help people see who is paying for the ad—which is especially

important when the Page name doesn't match the name of the company or person funding the ad. For more information, see https://newsroom.fb.com/news/2018/04/transparent-ads-and-pages/.

When people click on the label, they'll be taken to an archive with more information. For example, we'll provide the campaign budget associated with an individual ad and how many people saw it—including aggregated information about their age, location and gender. That same archive can be reached at https://www.facebook.com/politicalcontentads. People on Facebook visiting the archive can see and search ads we've identified with political or issue content that an advertiser has run in the US for up to 7 years.

Advertisers wanting to run ads with political or issue content in the US and certain other countries will need to verify their identity and location. More information is available at https://newsroom.fb.com/news/2018/04/transparent-ads-and-pages/. Enforcement of these new features and the Political Ads policy, available at https://www.facebook.com/policies/ads/restricted_content/political, began on May 24.

We're closely monitoring developments in Congress, including proposed legislation like the Honest Ads Act. Our policy reflects language from existing laws as well as proposed laws. But, we're not waiting. We've been hearing calls for increased transparency around ads with political content for some time now. We've taken the first steps toward providing that transparency, and we hope others follow.

[From Senator Wyden]

34. In July 2018, Facebook took down a fake account promoting a counter-protest against the United the Right demonstration in Washington, D.C.

- **Were there any advertisements, originating from fake or legitimate accounts, directing users to the pages associated with the fake account? If yes, what did Facebook do with regard to the accounts associated with those ads?**

In July 2018, we removed 32 Pages and accounts from Facebook and Instagram because they were involved in coordinated inauthentic behavior. This kind of behavior is not allowed on Facebook because we don't want people or organizations creating networks of accounts to mislead others about who they are, or what they're doing. We shared this information with US law enforcement agencies, Congress, other technology companies, and the Atlantic Council's Digital Forensic Research Lab, a research organization that helps us identify and analyze abuse on Facebook.

- In total, more than 290,000 accounts followed at least one of these Pages, the earliest of which was created in March 2017. The latest was created in May 2018.

- The most followed Facebook Pages were "Aztlan Warriors," "Black Elevation," "Mindful Being," and "Resisters." The remaining Pages had between zero and 10 followers, and the Instagram accounts had zero followers.

- There were more than 9,500 organic posts created by these accounts on Facebook, and one piece of content on Instagram.

- The 32 Pages and accounts ran about 150 ads for approximately $11,000 on Facebook and Instagram, paid for in US and Canadian dollars. The first ad was created in April 2017, and the last was created in June 2018.

- The Pages created about 30 events since May 2017. About half had fewer than 100 accounts interested in attending. The largest had approximately 4,700 accounts interested in attending, and 1,400 users said that they would attend.

We found this activity as part of our ongoing efforts to identify coordinated inauthentic behavior. Given these bad actors are now working harder to obscure their identities, we need to find every small mistake they make. It's why we're following up on thousands of leads, including information from law enforcement and lessons we learned from last year's IRA investigation. The IRA engaged with many legitimate Pages, so these leads sometimes turn up nothing. However, one of these leads did turn up something. One of the IRA accounts we disabled in 2017 shared a Facebook Event hosted by the "Resisters" Page. This Page also previously had an IRA account as one of its admins for only seven minutes. These discoveries helped us uncover the other inauthentic accounts we disabled.

The "Resisters" Page also created a Facebook Event for a protest on August 10 to 12 and enlisted support from real people. The Event—"No Unite the Right 2–DC"—was scheduled to protest an August 2018 "Unite the Right" event in Washington. Inauthentic admins of the "Resisters" Page connected with admins from five legitimate Pages to co-host the event. These legitimate Pages unwittingly helped build interest in "No Unite Right 2–DC" and posted information about transportation, materials, and locations so people could get to the protests.

We disabled the event on July 31, 2018 and reached out to the admins of the five other Pages to update them on what happened. We also informed the approximately 2,600 users interested in the event, and the more than 600 users who said they'd attend, about what happened.

35. Facebook's statement noted that the administrators of the fake account "connected with admins from five legitimate Pages to co-host the event," and that Facebook "reached out to the admins of the five other Pages to update them on what happened."

- **What is Facebook's policy in circumstances in which fake accounts have joined with legitimate, but unwitting American political actors in promoting events or causes?**

As discussed above, we disabled the "No Unite Right 2–DC" event on July 31, 2018 and reached out to the admins of the five other Pages to update them on what happened. We also informed the approximately 2,600 users interested in the event, and the more than 600 users who said they'd attend, about what happened. This is a challenging issue, and whenever we take action on inauthentic behavior on Facebook, we work to balance (a) enforcing against the inauthentic behavior; (b) preserving legitimate voices that may have unknowingly interacted

with inauthentic accounts; and (c) protecting the privacy of legitimate accounts that may have unknowingly interacted with inauthentic accounts.

36. Since the 2016 election, has any foreign government, or anyone that Facebook believes to be acting on the behalf of a foreign government, used Facebook to promote or amplify misleading or "hoax" content to users in the United States (for example, claims that a national tragedy did not occur or was perpetrated by our own government)?

- **If yes, please provide a detailed accounting of each case, including the suspected foreign entity, and the number of users that saw or interacted with the content (e.g. clicked or shared).**

Our security teams are constantly monitoring for foreign information operations. For example, in July, we removed 32 Pages and accounts from Facebook and Instagram that were engaged in coordinated inauthentic behavior. These Pages had some links to previously removed IRA-affiliated accounts, but we were unable to determine whether this new cluster of activity was directly controlled by the IRA.

In August, we removed Pages, groups, and accounts that were linked to sources the US government had previously identified as Russian military intelligence services. This cluster was focused on politics in Syria and Ukraine. To date, we have not found activity by these accounts targeting the US. We are working with US law enforcement on this investigation. At the same time, we also removed a separate set of 652 Pages, groups, and accounts for coordinated inauthentic behavior that originated in Iran and targeted people across multiple internet services in the Middle East, Latin America, UK, and US.

More information is available at http://newsroom.fb.com.

37. Since the 2016 election, has any foreign government, or anyone that Facebook believes to be acting on the behalf of a foreign government, attempted to influence public opinion in the United States by using Facebook to coordinate with, or assist (e.g. by providing content, guidance, or other forms of support) individuals or groups known to promote "hoaxes" and misleading reports (such as those described in in the prior question)?

- **If yes, please provide a detailed accounting of each case, including the nature of the relationship, and whether the suspected foreign entity or its agent appears to have taken steps to mask their true identity or sponsor.**

See Response to Question 36.

38. What steps has Facebook taken to inform its users, the public, and the United States Government of each case listed in response to the two previous questions?

We have worked to notify people about foreign influence operations on a variety of occasions and will continue to do so as appropriate.

39. Facebook has confirmed that the Russian Government and its agents created fake organizations and personas to promote causes and issues in the United States during the 2016 presidential election. In July 2018, Facebook announced that an entity using tools and techniques that were similar to those used in 2016 by the Russian Internet Research Agency was attempting to manipulate public sentiment in the United States. In August 2018, Facebook announced that it had deactivated additional pages, groups and accounts linked to Russia and Iran that were spreading disinformation.

- **In addition to the cases listed above, has any foreign government, their agent, or an entity acting on the behalf of a foreign government, created content, groups, pages or accounts that masquerade as American for the purpose of influencing political debate or policymaking within the United States, not limited to elections?**

We are constantly monitoring for foreign information operations, including efforts to mislead users about the source of content or the location of other users. When we detect these networks, we investigate them and take them down. However, we generally do not discuss planned takedowns publicly to avoid compromising our investigation or alerting the actors.

40. Has any other foreign entity, even if it is not known to be acting on behalf of a foreign government, created content, groups, pages or accounts that masquerade as American for the purpose of influencing political debate or policymaking within the United States, not limited to elections?

- **If the answer to either of the previous two questions is yes, please provide a detailed accounting of each case, including the foreign government (if applicable), the issue, and the number of users that saw or interacted with the content (e.g. clicked or shared).**

See Response to Question 39.

41. What steps has Facebook taken to inform users, the public, and the United States Government of any cases that you have listed in response to the previous question?

See Response to Question 38.

42. Facebook, like several other major technology companies, warns users when it believes their accounts may have been targeted by foreign governments.

- **In each of the past five years, how many times has Facebook notified users located in the United States that their accounts were targeted by a foreign government?**
    - **Prior to being notified by Facebook, how many of these accounts had some form of two-factor authentication enabled on their accounts?**
    - **Prior to being notified by Facebook, how many of these accounts were secured with a two-factor authentication security key?**

- **In each of the past five years, how many times has Facebook notified users believed by Facebook to be elected officials or their staff in the United States that their accounts were targeted by a foreign government?**

    - **Prior to being notified by Facebook, how many of these accounts had some form of two-factor authentication enabled on their accounts?**

    - **Prior to being notified by Facebook, how many of these accounts were secured with a two-factor authentication security key?**

We do not maintain public statistics on this issue. For more information on two-factor authentication, see Response to Question 47.

This will never be a solved problem because we're up against determined, creative, and well-funded adversaries. But we are making steady progress. Here is a list of 10 important changes we have made:

- **Ads and Pages transparency.** Advertising should be transparent: users should be able to see all the ads an advertiser is currently running on Facebook, Instagram and Messenger. And for ads with political or issue content, we've created an archive that will hold ads with political or issue content for 7 years—including information about ad impressions and spend, as well as demographic data such as age, gender, and location. And people everywhere can see all the ads that Page is running on Facebook. We also announced in April that people who manage Pages with large numbers of followers will need to be verified. Those who manage large Pages that do not clear the process will no longer be able to post. This will make it much harder for people to administer a Page using a fake account, which is strictly against our policies. We will also show people additional context about Pages to help people have more information to evaluate their content. For example, you can see whether a Page has changed its name.

- **Verification and labeling.** Every advertiser will now need to confirm their ID and location before being able to run any ads with political or issue content in the US and certain other countries. All ads with political or issue content will also clearly state who paid for them.

- **Updating targeting.** We want ads on Facebook to be safe and civil. We thoroughly review the targeting criteria advertisers can use to ensure they are consistent with our principles. As a result, we removed nearly one-third of the targeting segments used by the IRA. We continue to allow some criteria that people may find controversial. But we do see businesses marketing things like historical books, documentaries, or television shows using them in legitimate ways.

- **Better technology.** Over the past year, we've gotten increasingly better at finding and disabling fake accounts. We now block millions of fake accounts each day as people try to create them—and before they've done any harm. This is thanks to improvements in machine learning and artificial intelligence, which can proactively

identify suspicious behavior at a scale that was not possible before—without needing to look at the content itself.

- **Action to tackle fake news.** We are working hard to stop the spread of false news. We work with third-party fact-checking organizations to limit the spread of articles rated false. To reduce the spread of false news, we remove fake accounts and disrupt economic incentives for traffickers of misinformation. We also use various signals, including feedback from our community, to identify potential false news. In countries where we have partnerships with independent third-party fact-checkers, stories rated as false by those fact-checkers are shown lower in News Feed. If Pages or domains repeatedly create or share misinformation, we significantly reduce their distribution and remove their advertising rights. We also want to empower people to decide for themselves what to read, trust, and share. We promote news literacy and work to inform people with more context. For example, if third-party fact-checkers write articles about a news story, we show them immediately below the story in the Related Articles unit. We also notify people and Page Admins if they try to share a story, or have shared one in the past, that's been determined to be false. In addition to our own efforts, we're learning from academics, scaling our partnerships with third-party fact-checkers and talking to other organizations about how we can work together.

- **Significant investments in security.** As part of our larger company investment in the space, we have more than doubled the number of people working on safety and security and now have over 20,000. We expect these investments to impact our profitability. But the safety of people using Facebook needs to come before profit.

- **Industry collaboration.** Recently, we joined more than 60 global tech and security companies in signing a TechAccord pact to help improve security for everyone.

- **Information sharing and reporting channels.** In the 2017 German elections, we worked closely with the authorities there, including the Federal Office for Information Security (BSI). This gave them a dedicated reporting channel for security issues related to the federal elections.

- **Tracking 40+ elections.** We deployed new tools and teams to proactively identify threats in the run-up to specific elections. We first tested this effort during the Alabama Senate election, and have continued these efforts for elections around the globe, including the US midterms. Last year we used public service announcements to help inform people about fake news in 21 separate countries, including in advance of French, Kenyan and German elections.

- **Action against the Russia-based IRA.** In April, we removed 70 Facebook and 65 Instagram accounts—as well as 138 Facebook Pages—controlled by the IRA primarily targeted either at people living in Russia or Russian-speakers around the world including from neighboring countries like Azerbaijan, Uzbekistan, and Ukraine. The IRA has repeatedly used complex networks of inauthentic accounts to deceive and manipulate people in the US, Europe, and Russia—and we don't want them on Facebook anywhere in the world. In July, we removed 32 Pages and

accounts from Facebook and Instagram that were engaged in coordinated inauthentic behavior. These Pages had some links to previously removed IRA-affiliated accounts, but we were unable to determine whether this new cluster of activity was directly controlled by the IRA. Our security teams are continuing to monitor our platform for abuse in connection with future elections here and around the world.

43. In each of the past five years, how many user accounts, if any, have been compromised, such that someone other than the user gained access to the user's non-public account data?

- **How many of these accounts had some form of two-factor authentication enabled on their accounts?**

- **How many of these accounts were secured with a two-factor authentication security key?**

We recently shared that we discovered a security issue affecting 30 million accounts. People's security is incredibly important, and we're sorry this happened. It's why we've taken immediate action to secure these accounts and let users know what happened.

Although two-factor authentication would not have mitigated this security attack, we believe strongly that two-factor authentication is a valuable tool for safeguarding an account. We enable it and promote it and we require it by default for groups that may be particular security targets, including anyone who wants to run ads related to politics or issues of national importance in the US and people who manage Pages with large audiences in the US. We provide training to candidates, government officials, advocacy groups, and others during live events on how to take common sense safety precautions, including turning on two-factor authentication. Please see Response to Question 47 for more information regarding two-factor authentication.

44. In each of the past five years, how many user accounts were compromised, such that someone other than the user gained access to the user's non-public account data, by adversaries that Facebook believes may be a foreign government or are working with a foreign government?

- **How many of these accounts had some form of two-factor authentication enabled on their accounts.**

- **How many of these accounts were secured with a two-factor authentication security key?**

We do not maintain public statistics on this issue.

45. Facebook provides the Custom Audiences tool to enable advertisers to micro-target individuals based on data about those users that they already possess.

- **Is Facebook aware of any advertisements targeted with Custom Audiences that appear to be designed to discourage any United States citizen from voting?**

Our policies prohibit—in both ads and organic content—misrepresentations of the dates, locations, and times for voting or voter registration. We also prohibit misrepresentation of who can vote, qualifications for voting, and what information and/or materials must be provided in order to vote. We remove this content when we become aware of it and ads that violate these policies are disapproved. Facebook is committed to transparency for all ads, including ads with political or issue content. Facebook believes that people should be able to easily understand why they are seeing ads, who paid for them, and what other ads those advertisers are running. As such, Facebook only allows authorized advertisers to run ads in the US about elections or issues that are being debated across the country. In order to be authorized by Facebook, advertisers need to confirm their identity and location. Furthermore, in the US, all political and issue ads include a disclosure, which reads: "Paid for by," and when users click on this disclosure they will be able to see more information about the ad and advertiser. Users will also be able to see an explanation of why they saw the particular ad.

  - **If yes, please provide a full accounting of each case, including the advertisement, what Facebook knows about the party that purchased the advertising, and the number of users that saw or interacted with the content (e.g. clicked).**

  See Response to above Question.

- **If the answer to the question above is yes, were these voter discouragement ads targeted at people of any particular race or ethnic group?**

  - **Were these voter discouragement ads predominantly targeted at people expected to vote for one party or the other?**

  See Response to above Question.

- **Has any foreign government, their agent, or other foreign entity ever used Custom Audiences to target individuals in the United States?**

  - **If yes, please provide a full accounting of each case, including the party that purchased the advertising, the foreign government sponsor (if applicable), and the number of users that saw or interacted with the content (e.g. clicked or shared).**

  See Response to above Question.

- **Has the Internet Research Agency ever used Custom Audiences to target users, in the United States or elsewhere, with advertisements?**

The targeting for the IRA ads that we have identified and provided to the Senate Committee on the Judiciary and the Senate Select Committee on Intelligence was relatively rudimentary, targeting very broad locations and interests, and for example, only used custom audiences in a very small percentage of its overall targeting and did not use Contact List Custom Audiences. In addition, all of the custom audiences used by the IRA were created based on user engagement with certain IRA Pages.

- **Does Facebook believe that any of the content created by the Russian Internet Research Agency was designed to discourage anyone from voting?**

We believe this is an assessment that can be made only by investigators with access to classified intelligence and information from all relevant companies and industries—and we want to do our part. Congress is best placed to use the information we and others provide to inform the public comprehensively and completely, which is why we provided IRA ads and content to the Senate Select Committee on Intelligence for review.

- **Can users opt out of being targeted with Custom Audiences?**
    - **If no, why not?**

We provide controls that specifically govern the use of data for ads. Through Ad Preferences, people see and control things like: (1) their "interests," which are keywords associated with a person based on activities such liking Pages and clicking ads; (2) their "behaviors" (which we also call "categories"), which generally reflect how, when and where they connect to Facebook; and (3) the advertisers that are currently showing them ads based on the person's contact information, based on the person's previous use of the advertiser's website or app, or based on a visit to the advertiser's store. People also can choose whether we use information about their activities on websites and apps off of Facebook to show them ads through Facebook, and whether we can use their Facebook advertising interests to show them ads off of Facebook.

Advertisers also bring us the customer information so they can reach those people on Facebook. These advertisers might have, for example, people's email addresses from purchases users made, or from some other data source. If we have matching email addresses, we can show those people ads from that advertiser (although we cannot see the email addresses which are sent to us in hashed form, and these are deleted as soon as we complete the match). In ad preferences people can see which advertisers with their contact information are currently running campaigns—and they can click the top right corner of any ad to hide all ads from that business.

- **Does Facebook have a policy of shutting down pages and accounts that seek to suppress voting, regardless of whether they are found to be inauthentic?**

    - **If yes, to what kind of content does Facebook apply that policy (e.g., content discouraging people from voting, content providing inaccurate information on how or when to vote, etc.)?**

As part of our ongoing efforts to prevent people from misusing Facebook during elections, we're broadening our policies against voter suppression—action that is designed to deter or prevent people from voting. These updates were designed to address new types of abuse that we're seeing online.

We already prohibit offers to buy or sell votes as well as misrepresentations about the dates, locations, times and qualifications for casting a ballot. We have been removing this type of content since 2016.

Last month, we extended this policy further and are expressly banning misrepresentations about how to vote, such as claims that you can vote using an online app, and statements about whether a vote will be counted (e.g. "If you voted in the primary, your vote in the general election won't count."). We've also recently introduced a new reporting option on Facebook so that people can let us know if they see voting information that may be incorrect, and have set up dedicated reporting channels for state election authorities so that they can do the same.

We recognize that some posts that are reported to us may require additional review. For example, we're unable to verify every claim about the conditions of polling places around the world (e.g. "Elementary School Flooded, Polling Location Closed"). In these cases, we will send content to our third-party fact-checkers for review. Content that is rated false will be ranked lower in News Feed, and accompanied by additional information written by our fact-checkers (what we call, Related Articles) on the same subject.

46. According to a British Member of Parliament, Britain's Information Commissioner's Office found evidence that data collected by Aleksandr Kogan was accessed from Russia and other countries.

- **Please list all entities or individuals outside the United States or the United Kingdom that Facebook is aware of that accessed or received any part of the user data originally obtained by Aleksandr Kogan.**

    - **Please explain what Facebook knows about each instance.**

Kogan represented that, in addition to providing data to his Prosociality and Well-Being Laboratory at the University of Cambridge for the purposes of research, GSR provided some Facebook data to SCL Elections Ltd., Eunoia Technologies, and the Toronto Laboratory for Social Neuroscience at the University of Toronto. Our investigation is ongoing.

Facebook obtained written certifications from Kogan, GSR, and other third parties (including Cambridge Analytica and SCL) declaring that all data they had obtained, and any derivatives, were accounted for and destroyed. We are seeking to conduct a forensic audit of

Cambridge Analytica's systems to confirm the veracity of these certifications, but the UK Information Commissioner's Office, which is conducting a regulatory investigation into Cambridge Analytica (based in the UK), has the only known copy of Cambridge Analytica's systems and will need to release that information for us to conduct this audit. We hope to move forward with that audit soon.

- **Is Facebook aware of any instances in which user data obtained by Kogan was subsequently used to target Facebook users, either during the 2016 Election, or at any other time?**

- **Please describe in detail all uses of user data obtained by Kogan of which Facebook is aware.**

- **What efforts have been made to ensure that user data obtained by Kogan has been completely deleted, and cannot be used in the future by any party, for any purpose?**

On December 11, 2015, *The Guardian* published an article reporting that Kogan and his company, GSR, may have passed information the app had obtained from Facebook users to SCL Elections Ltd./Cambridge Analytica. Kogan and his company violated Facebook's Platform Policies, which explicitly prohibited selling user data accessed from Facebook and from sharing any user data accessed from Facebook with any ad network, data broker or other advertising or monetization related service.

For this reason, Facebook immediately banned the app from our platform and investigated what happened and what further action we should take to enforce our Platform Policies. Facebook also contacted Kogan/GSR and demanded that they explain what data they collected, how they used it, and to whom they disclosed it. Facebook further insisted that Kogan and GSR, as well as other persons or entities to whom they had disclosed any such data, account for and irretrievably delete all such data and information.

Facebook also contacted Cambridge Analytica to investigate the allegations reflected in the reporting. On January 18, 2016, Cambridge Analytica provided written confirmation to Facebook that it had deleted the data received from Kogan and that its server did not have any backups of that data. On June 11, 2016, Kogan signed certifications of deletion on behalf of himself and GSR. The certifications also purported to identify all of the individuals and entities that had received data from GSR (in addition to Kogan and his lab), listing the following: SCL, Eunoia Technologies (a company founded by Christopher Wylie), and a researcher at the Toronto Laboratory for Social Neuroscience at the University of Toronto. On July 7, 2016, a representative of the University of Toronto certified that it deleted any user data or user-derived data. On August 16, 2016, Eunoia (executed by Eunoia Founder Christopher Wylie) certified that it deleted any user and user-derived data. On September 6, 2016, counsel for SCL informed counsel for Facebook that SCL had permanently deleted all Facebook data and derivative data received from GSR and that this data had not been transferred or sold to any other entity. On April 3, 2017, Alexander Nix, on behalf of SCL, certified to Facebook, that SCL deleted the information that it received from GSR or Kogan.

Because all of these concerns relate to activity that took place off of Facebook and its systems, we have no way to confirm whether Cambridge Analytica may have retained Facebook data without conducting a forensic audit of its systems. Cambridge Analytica has agreed to submit to a forensic audit, but we have not commenced that yet due to a request from the UK Information Commissioner's Office, which is simultaneously investigating Cambridge Analytica (which is based in the UK). And even with an audit, it may not be possible to determine conclusively what data was shared with Cambridge Analytica or whether it retained data after the date it certified that data had been deleted.

Although our developer terms gave us the ability to audit Kogan's app, we did not have an agreement in place that would have allowed us to audit third parties that he may have shared data with. So we obligated him to obtain certifications of deletion from each of these parties, leveraging our rights as to Kogan, who was the developer of the app.

The existing evidence that we are able to access supports the conclusion that Kogan only provided SCL with data on Facebook users from the United States. While the accounts of Kogan and SCL conflict in some minor respects not relevant to this question, both have consistently maintained that Kogan never provided SCL with any data for Facebook users outside the United States. These consistent statements are supported by a publicly released contract between Kogan's company and SCL.

In March 2018, we learned from news reports that, contrary to the certifications given, not all of the Kogan data may have been deleted by Cambridge Analytica. We have no direct evidence of this and no way to confirm this directly without accessing Cambridge Analytica's systems and conducting a forensic audit. We have held off on audits of Cambridge Analytica and other parties that are being investigated by the UK Information Commissioner's Office at its request. Our investigation is ongoing.

47. For several years, Facebook has allowed its customers to protect their accounts from hacking through the use of two-factor authentication, including using physical security tokens as an enhanced form of two-factor authentication. However, two-factor authentication remains an opt-in feature for Facebook users.

- **Does Facebook require that its employees use two-factor authentication for their work accounts?**
    - **If yes, does Facebook require, like Google, that employees use a security key?**
- **Do you and Mr. Zuckerberg have two-factor authentication enabled for your personal Facebook and personal email accounts?**
    - **If yes, are you using security keys?**
- **What percentage of Facebook's U.S. customers have enabled any form of two-factor authentication?**
- **What percentage of Facebook's U.S. customers have enabled enhanced two-factor authentication using a security key?**

- **Facebook specially identifies the accounts of elected officials. What percentage of the Facebook accounts of elected officials in the United States currently have any form of two-factor authentication enabled?**
    - **What percentage are using a security key?**
- **What specific outreach, if any, has Facebook engaged in to encourage elected officials to enable two-factor authentication on their official and personal Facebook accounts?**
- **Facebook will place a blue verification badge on the accounts of brands, media organizations and public figures who have been verified as authentic by Facebook. Does Facebook currently require that verified accounts enable two-factor authentication?**

Two-factor authentication is a security feature that helps protect users' Facebook accounts and passwords. If a user sets up two-factor authentication, they are asked to enter a special login code or confirm their login attempt each time someone tries accessing Facebook from a computer or mobile device Facebook doesn't recognize. A user can also get alerts when someone tries logging in from a computer Facebook doesn't recognize. Two-factor authentication is an industry best practice for providing additional account security. We continue to encourage enabling two-factor authentication to add an extra layer of protection to Facebook accounts when people think it's appropriate.

Facebook requires employees to use two-factor authentication for their work accounts; they have the option to use security keys or Duo push notifications.

We are committed to helping people on our platform protect their accounts and take special steps to encourage people who may be more vulnerable to attack to enable two-factor authentication. This includes:

- Requiring two-factor authentication for anyone who wants to run ads related to politics or issues of national importance in the US.
- Requiring two-factor authentication for people who manage Pages with large audiences in the US.
- Sending notifications (on Facebook and via email) to people involved in politics, including the Page admins for elected officials, that encourage them to turn on two-factor authentication.
- Providing a Safety Guide for Page Admins which we delivered in person to every House and Senate office in September, which highlighted two-factor authentication.
- Highlighting how to use two-factor authentication on our website created especially for government officials and those involved in politics (http://politics.fb.com). On this website, turning on two-factor authentication is the very first step in our guide: https://politics.fb.com/learn-the-basics/.

- Training staff, candidates, government officials, advocacy groups, and others during live events how to take common sense safety precautions, including turning on two-factor authentication.

- Creating a video explainer on two-factor authentication specifically for those involved in politics, available at: https://politics.fb.com/learn-the-basics/#component-1-secure-your-account.

48. In June 2018, Facebook admitted to having entered into data sharing partnerships with device manufacturers, including Huawei. According to news reports, Facebook stated that user data made available through the Huawei partnership was stored on the smartphones of users, not on Huawei's servers, and the data was "controlled" by Facebook.

- **Has Facebook audited every version of Huawei's applications since the beginning of this partnership to ensure that there was never an instance in which user data was uploaded to Huawei's servers or was otherwise accessible by Huawei?**

Facebook, along with many other technology companies, has worked with Chinese device manufacturers to integrate services Facebook provides onto devices provided by those companies. Huawei, for instance, is the third largest mobile manufacturer in the world.

As previously noted, the purpose of the device integration partnerships Facebook had with partners like Huawei and other device manufacturers was not to share information with the partners (or to enable Facebook users to do so), but to provide limited rights to use APIs to build Facebook integrations and features into their devices and other products. Facebook's partnerships and engineering teams were involved in reviewing and approving the development of the device integrations like Huawei's, thereby ensuring oversight and involvement in the implementation of these APIs into Facebook-approved device integrations. There were likewise additional controls such as specifically-negotiated agreements with device integration partners (including Huawei), which again provided limited rights to use APIs to create the device integrations approved by Facebook, and not independent purposes determined by the partner.

Finally, we are not aware of any abuse of user data by Huawei (or other device integration partner), and Huawei has publicly confirmed that it has never collected or stored any Facebook user data on its servers.

- **As part of the device manufacturer partnerships revealed in June, did Facebook allow device manufacturers to bypass Facebook's normal user interface for obtaining permission from users to access their data and instead use custom prompts to obtain permission from users?**

Users were required to authorize the Facebook device integration on their device and log into Facebook just like they would if they logged into Facebook on the Facebook website or mobile app. These logins were often custom to the app and were approved by Facebook. Facebook's data policies, at least since 2010, have advised users that we work with other companies to provide our services in different contexts.

- **If yes:**

    a. **Who created the custom permission interfaces used by these applications, Facebook or the device manufacturer?**

The device integrations were designed by Facebook's partners and reviewed by Facebook, which had to approve implementations of the APIs. Typically, these apps were reviewed and approved by members of our partnerships and engineering teams.

    b. **Did Facebook disclose the existence of these custom permission interfaces to the Federal Trade Commission?**

Facebook has discussed its device integration partnerships with the FTC.

    c. **Were each of these custom permission screens reviewed by Facebook to ensure compliance with the Federal Trade Commission Consent Order?**

These device integrations were reviewed by Facebook, which had to approve the apps. Typically, these apps were reviewed and approved by members of our partnerships and engineering teams. The obligations imposed by the FTC 2012 Consent Order on Facebook's use of service providers, such as these device integration partners, differ materially from those imposed on Facebook with respect to third parties. Indeed, the Consent Order excludes service providers from its definition of "third parties." Facebook's data policies—at least since 2010—have likewise informed users that Facebook works with other companies to provide its services in different contexts.

    d. **Were any of these custom permission screens examined by Facebook's external auditors, as part of the biennial audits required by the Federal Trade Commission Consent Order?**

The independent firm's assessment process included an assessment of controls related to Facebook's device integration partners. Again, as noted above, the obligations imposed by the Consent Order on Facebook's service providers, such as these device integration partners, differ from those imposed on Facebook with respect to other third parties.

- **Did Facebook provide data on the friends of users as part of these partnerships, in addition to the users of the apps themselves?**

Facebook has previously identified device integrated partnerships with access to friends' data after that functionality was removed from Facebook's public platform in 2015. As discussed above, app settings that restricted friends' data from being shared with third-party apps that people's friends used generally did not apply to these integration partners, because they were not functioning as third-party apps, and instead were providing core Facebook experiences. Users' privacy settings did apply equally to integration partnerships, however.

- **Did Facebook ever permit companies that were part of these partnerships, including Huawei, to access data, either about a user or their friends, that the partner would**

otherwise be prevented from accessing because of Facebook privacy preferences configured by a user or their friends, as alleged by the New York Times in June 2018?

See Response to above Question.

- **Did Facebook ever notify its users that their data could be accessed by device manufacturers, regardless of how they had configured their Facebook privacy settings?**

As noted above, the relevant Facebook privacy controls and settings applied to information people shared with friends who used a partner's device integration.

In addition, users authorized Facebook device integrations by signing in on a device much like they would on Facebook's website and in the mobile apps we built. For example, users accessing Facebook on their Blackberry device would log into Facebook just like they would if they logged into Facebook on the http://www.facebook.com/ website (even though that version of Facebook was built by Blackberry under an agreement with Facebook). This is not unlike the experience people have when accessing their email account on a mobile device: in that case, the login experience may be facilitated by the device manufacturer (or other integration partner).

Finally, Facebook's Data Policies—since at least 2010—have informed users that we work with other companies to provide our services in different contexts.

- **Has Facebook ever disclosed to the Federal Trade Commission or its external auditor that Facebook's user privacy settings did not control device manufacturers' access to user data?**

As described above, Facebook privacy controls and settings applied to information people shared with friends who used a partner's device integration. Facebook has discussed its device integration partnerships and applicable settings with both the FTC and the independent firm that provides ongoing assessments under the consent order.

- **Prior to June of this year, was Facebook ever warned by its own employees or by any other entity about the partnerships, including that providing device manufacturers with access to user data that was not constrained by Facebook's user privacy settings might violate the terms of its 2011 Federal Trade Commission Consent Order?**
    - **If yes, please provide a copy of any documentation about this warning and the steps taken, if any, by Facebook in response.**

Please see the response to the prior question. Facebook's device integration partnerships did not violate the terms of the 2012 FTC Consent Order and honored users' privacy settings.

- **Approximately how many users did the devices that were granted this special access have, in total?**
    - **How many were users in the United States?**

As noted above, device integration partners differed significantly from third-party developers' building of independent third-party consumer apps on Facebook's developer platform. Device integration partnerships began in the early days of mobile when the demand for Facebook outpaced our ability to build versions of our product that work on every phone or operating system. The value of these device integration partnerships has diminished over time, as more people download the apps we build through app stores on iOS and Android. As a result, Facebook has wound down the majority of these arrangements.

- **If user data shared with these partners was only stored on user devices, are there circumstances such as a software bug or a user backing up their data to a cloud service where the data would have been sent to the partners' servers?**

Whether data was stored on the partner's server depended on the partner's infrastructure during the time when the device integrations were active. Some partners, such as Blackberry, offered client-server syncing that helped people back up their content to the partner's servers. Other partners did not. What is important to understand is that the purpose of these partnerships was not to share data directly with the partner, but to enable people to use Facebook and Facebook-like experiences on different devices and in different software.

- **What methods did Facebook employ to confirm that none of the data provided through these partnerships was accessed or used in an inappropriate way by the partners?**

The purpose of these device integration partnerships was not to share information with the partners (or to enable Facebook users to do so), but to provide limited rights to use APIs to build Facebook apps and features into their devices and other products. These device partners could only use data accessed through these APIs to provide the approved device integration, and only to support experiences specifically requested by the Facebook user. Partners were not allowed to use data received through the APIs for their own independent purposes, unless they separately obtained consent from the user.

Our partnerships and engineering teams were involved in reviewing and approving the development of the integrations with device manufacturers—thereby ensuring oversight and involvement in the implementation of these APIs into Facebook-approved device integrations. We monitored the usage patterns of our APIs for irregularities, and we are not aware of any violations of our agreements with these partners

- **Has Facebook ever audited any of its partners?**

See Response to above Question.

a. If yes, please describe the scope of each audit.

See Response to above Question.

- **Has Facebook ever been asked or advised by any U.S. government entities or officials not to share user data with Huawei or any other company with reported relationships with foreign intelligence services?**

Facebook maintains a dialogue with the US government on a range of cybersecurity issues.

  - **If yes, please describe each case.**

See Response to above Question.

- **Were there any other partner applications that were given special access which were not created by device manufacturers? If so, what were these applications, and why were they given access?**

As previously explained above, Facebook engaged companies to build integrations for a variety of devices, operating systems, and other products where Facebook and its partners wanted to offer people a way to receive Facebook or Facebook experiences. They included, for example, Facebook-branded apps, social networking service hubs, syncing integrations, and USSD services. As described in Facebook's Data Policies, Facebook also works with other types of partners in a variety of contexts which may involve access to user information depending on the nature of the partnership and agreement.

  - **Was data from these partnerships ever stored on these partners' servers?**

Whether data was stored on the partner's server depended on the partner's infrastructure during the time when the device integrations were active. Some partners, such as Blackberry, offered client-server syncing that helped people back up their content to the partner's servers. Other partners did not. What is important to understand is that the purpose of these partnerships was not to share data directly with the partner, but to enable people to use Facebook and Facebook-like experiences on different devices and in different software.

    a. If yes, which partners stored the data on their servers?

The nature of the partnership varied from partner to partner.

  - **About how many users did the devices that were granted this special access have, in total?**

As noted above, device integration partners differed significantly from third-party developers' building of consumer apps on Facebook's developer platform. Device integration partnerships began in the early days of mobile when the demand for Facebook outpaced our ability to build versions of our product that work on every phone or operating system. The value of these device integration partnerships has diminished over time, as more people download the apps we build through app stores on iOS and Android. As a result, Facebook has now wound down the majority of these arrangements.

- **How many were users in the United States?**

See Response to above Question.

- **What methods did Facebook employ to confirm that none of the data provided through these partnerships was accessed or used in an inappropriate way by the partners?**

The purpose of these device integration partnerships was not to share information with the partners (or to enable Facebook users to do so), but to provide limited rights to use APIs to build Facebook apps and features into their devices and other products. These device partners could only use data accessed through these APIs to provide the approved device integration, and only to support experiences specifically requested by the Facebook user. Partners were not allowed to use data received through the APIs for their own independent purposes, unless they separately obtained consent from the user.

Our partnerships and engineering teams were involved in reviewing and approving the development of the integrations with device manufacturers—thereby ensuring oversight and involvement in the implementation of these APIs into Facebook-approved device integrations. We monitored the usage patterns of our APIs for irregularities, and we are not aware of any violations of our agreements with these partners

49. The Knight First Amendment Institute at Columbia University recently sent a letter to Facebook stating that Facebook's terms of service impede important public-interest journalism and research focused on Facebook's platform, because Facebook's terms prohibit the use of the basic tools of digital journalism and research. The Institute proposed that Facebook amend its terms of service to create a "safe harbor" protecting digital journalism and research focused on the platform.

- **Is it true that Facebook's terms of service bar certain journalism and research focused on the platform?**

- **If you have concerns about the Knight Institute's proposal, are there modifications to the proposal that would address your concerns while safeguarding digital journalism and researched focused on Facebook's platform?**

We are committed to working with journalists, researchers, and others to promote efforts to conduct research about Facebook in the public interest. At the same time, we have a responsibility to protect the privacy of the information people share on Facebook—including protecting it from scraping or unauthorized access. These protections are important, in part, because it is challenging for us to guard against misuse of people's information after it leaves our servers.

We are in conversations with the Knight Institute to understand more about the work that they would like to do, and to evaluate whether there are ways for us to advance transparency while protecting the information that people choose to share on Facebook. We look forward to continuing that dialogue.

[From Senator Lankford]

50. Related to the subject of "deep fakes," what is your ability to verify the authenticity of videos on your platform? What are the specific actions you are taking to identify the authenticity of videos on your platform?

Deepfakes take a number of different forms—from manipulated videos of celebrities to manufactured statements by political figures. Much of this content runs afoul of our existing content policies. For example, a photoshopped video of a celebrity in which the celebrity is nude would violate our nudity policies. Further, we have automated systems that help us identify nude and pornographic photos and videos that have previously been removed for violating our Community Standards. Deepfakes also may be spread by inauthentic accounts, which violate our policies—in that case, the content posted by such accounts would also be removed.

As we do across our work on misinformation, we're working on both technical and human review solutions to tackle deepfakes. Last month, for example, we announced the expansion of fact-checking to photos and videos to all of our fact-checking partners around the world, including in the United States. This effort will help us identify and take action against more types of misinformation, including manipulated photos and videos, more quickly.

In connection with the launch of fact-checking photos and videos, we have built a machine learning model that uses various engagement signals, including feedback from people on Facebook, to identify potentially false content in photos and videos. We then send those photos and videos to fact-checkers for their review, or fact-checkers can surface content on their own. Many of our third-party fact-checking partners have expertise evaluating photos and videos and are trained in visual verification techniques, such as reverse image searching and analyzing image metadata, like when and where the photo or video was taken. Fact-checkers are able to assess the truth or falsity of a photo or video by combining these skills with other journalistic practices, like using research from experts, academics or government agencies.

We are paying close attention to how research develops and are interested in working with others in the industry to come up with solutions to deepfakes. We are also working closely with our Facebook AI Research lab to help identify this type of content. We are committed to working with our industry partners and with Congress to develop solutions to combat this issue.

51. What is the process you use to validate someone as a legitimate actor for the purposes of furnishing them information for micro-targeting of a specific demographic group?

We provide advertisers with reports about the kinds of people seeing their ads and how their ads are performing, but we don't share information that personally identifies people (information such as name or that by itself can be used to contact or identifies a person) unless we have permission from people.

Advertisers wanting to run ads with political or issue content in the US and certain other countries will need to verify their identity and location. More information is available at https://newsroom.fb.com/news/2018/04/transparent-ads-and-pages/. Enforcement of these new features and the Political Ads policy, available at https://www.facebook.com/policies/ads/restricted_content/political, began on May 24.

52. Recently, WhatsApp and Google partnered to allow WhatsApp users the ability to backup communications on cloud-based Google Drive, free of charge.

- **If users do not opt into this service, are all of their messages protected by end-to-end encryption? If any party to a messaged conversation elects to use this service, will the entirety of the communication be stored on the Cloud-based Google Drive?**

WhatsApp users can back up their chats and media—including chats and media they've received—to Google Drive or iCloud, so if they change phones or get a new one, their chats and media are transferrable. Starting November 12, 2018, WhatsApp backups will no longer count towards the Google Drive storage quota.

WhatsApp uses end-to-end encryption. WhatsApp backups are not protected by WhatsApp's end-to-end encryption while in Google Drive or iCloud. Please see https://faq.whatsapp.com/en/android/28000019 and https://faq.whatsapp.com/en/iphone/20888066 for more details.

[From Senator Harris]

53. How much revenue, in dollars, has Facebook earned from ads that ran alongside content created by fake Russian Facebook accounts and pages?

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.

On Facebook, advertisers who use targeted ads are able to use our robust people-based marketing to deliver ads to their audience. The focus on the individual is what powers both a person's News Feed and our people-based marketing. What a person sees in their feed is based on who they are, who they follow and their own interests, allowing advertisers to deliver ads based on relevancy to every user and not the context of the stories around it. Through our research we've found that people view stories—both ads and organic content—in their News Feed as distinct pieces of content, unaffiliated with each other. A person might see a post from a relative about a birthday party followed by an article about their local community, with a clear understanding that these pieces of content are not related. We are happy to meet with you or your staff to further discuss how Facebook ads work.

54. What is your definition of "organic content?"

All paid advertisements on Facebook bear a label that reads "Sponsored," which clearly distinguishes them from organic content on Facebook.

55. What percent of your content is not organic?

The majority of content on Facebook is organic.

56. Exactly how long did Facebook's training material (1) instruct reviewers to delete hate speech by targeting white men but not hate speech targeting Black children, and (2) suggest that Black children are not a protected class? Please be specific.

We define hate speech as a direct attack on people based on what we call protected characteristics—race, ethnicity, national origin, religious affiliation, sexual orientation, caste, sex, gender, gender identity, and serious disease or disability. We also provide some protections for immigration status.

We expanded protections under our hate speech policies such that we now remove violent speech directed at groups of people defined by protected characteristics, even if the basis for the attack may be ambiguous. Under the previous hate speech policy, a direct attack targeting women solely on the basis of gender, for example, would have been removed from Facebook, but the same content directed at a sub-group, like "female drivers," would have remained on the platform. We recognize that the distinction was overly narrow. As such, we no longer differentiate between the two forms of attack when it comes to violent hate speech. We made this policy change in August 2017.

We are constantly evaluating—and, where necessary, changing—our content policies to account for shifts in cultural and social norms around the world. We continue to explore how we can adopt a more granular approach to hate speech, both in how we draft our policies and the way we enforce on them.

57. When did Facebook adopt its current Community Standards? Please be specific.

On April 24, 2018, we published, for the first time, the internal guidelines we use to enforce those standards.

We published these internal guidelines for two reasons. First, the guidelines will help people understand where we draw the line on issues. Second, providing these details makes it easier for everyone, including experts in different fields, to give us feedback so that we can improve the guidelines and the decisions we make.

The Content Policy team at Facebook is responsible for developing our Community Standards. We have people in offices around the world, including subject matter experts on issues such as hate speech, child safety, and terrorism. Many of the people on the team have worked on the issues of expression and safety long before coming to Facebook. The team includes a former criminal prosecutor who worked on child safety and counterterrorism, a former rape crisis counselor, an academic who has spent her career studying hate organizations, a human rights lawyer, and a teacher, among other expertise. Every week, the Content Policy team seeks input from experts and organizations outside Facebook so we can better understand different perspectives on safety and expression, as well as the impact of our policies on different communities globally.

Based on feedback, as well as changes in social norms and language, our standards evolve over time. What has not changed—and will not change—are the underlying principles of safety, voice, and equity on which these standards are based. To start conversations and make connections people need to know they are safe. Facebook should be a place where people can

express their opinions freely, even if some people might find those opinions objectionable. This can be challenging given the global nature of our service, which is why equity is such an important principle: we aim to apply these standards consistently and fairly to all communities and cultures. We outline these principles explicitly in the preamble to the standards, and we bring them to life by sharing the rationale behind each individual policy.

58. Ms. Sandberg's written testimony notes that "One of the main ways we identify and stop foreign actors is by proactively detecting and removing fake accounts, since they are the source of much of the interference we see." It further states that Facebook disabled 1.27 billion fake accounts from October 2017 to March 2018.

- **How many of the 1.27 billion fake accounts were part of Russia's disinformation campaign? Please be specific.**

- **How many of the 1.27 billion fake accounts were part of other countries' disinformation campaigns? Please name each country and list how many accounts have been attributed to this country.**

Facebook has conducted a broad search for evidence that Russian actors, not limited to the IRA or any other specific entity or organization, attempted to interfere in the 2016 election by using Facebook's advertising tools. We found coordinated activity that we now attribute to the IRA, despite efforts by these accounts to mask the provenance of their activity. We have used the best tools and analytical techniques that are available to us to identify the full extent of this malicious activity, and we continue to monitor our platform for abuse and to share and receive information from others in our industry about these threats.

In April, we removed 70 Facebook and 65 Instagram accounts—as well as 138 Facebook Pages—controlled by the IRA primarily targeted either at people living in Russia or Russian-speakers around the world including from neighboring countries like Azerbaijan, Uzbekistan, and Ukraine. The IRA has repeatedly used complex networks of inauthentic accounts to deceive and manipulate people in the US, Europe, and Russia—and we don't want them on Facebook anywhere in the world.

In July, we removed 32 Pages and accounts from Facebook and Instagram that were engaged in coordinated inauthentic behavior. These Pages had some links to previously removed IRA-affiliated accounts, but we were unable to determine whether this new cluster of activity was directly controlled by the IRA.

In August, we removed Pages, groups, and accounts that were linked to sources the US government had previously identified as Russian military intelligence services. This cluster was focused on politics in Syria and Ukraine. To date, we have not found activity by these accounts targeting the US. We are working with US law enforcement on this investigation. At the same time, we also removed a separate set of 652 Pages, groups, and accounts for coordinated inauthentic behavior that originated in Iran and targeted people across multiple internet services in the Middle East, Latin America, UK, and US.

Some state intelligence services, including Russia's, will use any medium available to conduct information operations. We continue to diligently search for their efforts to do so on our platform and will disrupt any that we find.

Detecting and removing fake accounts does not require precise measurements by country. Generating confident breakdowns beyond estimates is complicated because fake accounts use various techniques to attempt to disguise their location, including redirecting their traffic from remote locations, using proxies, VPNs and botnets. Our approach has therefore focused instead on how these fake accounts are created and how they operate, no matter where the accounts are created.

In our recent Community Standards Enforcement Report (which can be found at https://transparency.facebook.com/community-standards-enforcement), we shared the following details about Q1 of 2018:

- We estimate that fake accounts represented approximately 3 percent to 4 percent of monthly active users (MAU) on Facebook;
- We disabled 583 million fake accounts; and
- 98.5 percent of fake accounts acted on were flagged by Facebook before users reported them.

- **In Facebook's estimation, how many more accounts are plausibly linked to Russia's disinformation campaign? If you cannot provide a specific number, please provide an estimate.**

- **In Facebook's estimation, how many more accounts are plausibly linked to other countries' disinformation campaigns? If you cannot provide a specific number, please provide an estimate.**

See Response above regarding our efforts to detect coordinated inauthentic behavior linked to Russia and other state-sponsored actors. Our security teams are continuing to monitor our platform for abuse in connection with future elections here and around the world.

- **What indicators does Facebook use when attempting to identify fake accounts with Russian origins? Please be specific and comprehensive.**

We are committed to finding and removing fake accounts. 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 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. 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.

- **How many of the fake accounts:**
    - **Claimed a location in the United States and used Cyrillic characters in the account profile or posts?**
    - **Claimed a location in the United States but accessed Facebook via a Russian IP address?**
    - **Used a virtual private network to access Facebook?**

See Response above regarding our efforts to detect coordinated inauthentic behavior linked to Russia and other state-sponsored actors. We are unable to provide a reliable breakdown of fake accounts by these criteria. Fake accounts use various techniques to attempt to disguise their location, including redirecting their traffic from remote locations, using proxies, VPNs and botnets.

59. Will Facebook commit to reporting in its quarterly filings with the Securities and Exchange Commission, and if not, why not:

- **The number of accounts Facebook suspends for being inauthentic?**
- **The national origins of those accounts?**
- **The total pieces of content generated by those fake accounts?**
- **The number of impressions generated by those fake accounts?**
- **The number of fake accounts deemed inauthentic for each of the reasons described in your Community Standards, including misrepresenting identity, misusing profiles, impersonating others, and engaging in inauthentic behavior?**

Stopping the abuse of fake accounts and malicious bot activity is a focus for many teams, some more directly and some in more of a supportive role. For example, we are expanding our threat intelligence team, and more broadly, we have more than doubled the number of people working on safety and security and now have over 20,000. We expect to have at least 250 people specifically dedicated to safeguarding election integrity on our platforms, and that number does not include the thousands of people who will contribute to this effort in some capacity. We also 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 publish information and metrics about fake accounts at https://transparency.facebook.com/community-standards-enforcement#fake-accounts and in our quarterly SEC filings.

We will refine our approach over time, and we also hope to release additional metrics in future reports.

60. There are machine learning techniques that can create entirely fake videos, called "deepfakes." These deepfakes often depict people saying things they never said or portray events that never occurred.

- **Are deepfakes a violation of Facebook's terms of use?**
- **What is Facebook doing to identify deepfakes on its platform and to alert users when they may be seeing deepfakes?**
- **How many deepfakes has Facebook identified on its platform to date?**
- **Can Facebook commit to:**
    - **assessing how foreign disinformation campaigns can use deepfakes;**
    - **developing a strategy to combat it; and,**
    - **reporting its findings and efforts to the committee by the end of the year?**

Deepfakes take a number of different forms—from manipulated videos of celebrities to politicians. Much of this content runs afoul of our existing content policies. For example, a photoshopped video of a celebrity in which the celebrity is nude would violate our nudity policies. Further, we have automated systems that help us identify nude and pornographic photos and videos that have previously been removed for violating our Community Standards. Deepfakes may be spread by inauthentic accounts, which violate our policies—in that case, the content posted by such accounts would also be removed.

As we do across our work on misinformation, we're working on both technical and human review solutions to tackle deepfakes. Last month, we announced the expansion of fact-checking to photos and videos to all of our fact-checking partners around the world, including in the United States. This will help us identify and take action against more types of misinformation, including manipulated photos and videos, more quickly.

In connection with this launch, we have built a machine learning model that uses various engagement signals, including feedback from people on Facebook, to identify potentially false

content. We then send those photos and videos to fact-checkers for their review, or fact-checkers can surface content on their own. Many of our third-party fact-checking partners have expertise evaluating photos and videos and are trained in visual verification techniques, such as reverse image searching and analyzing image metadata, like when and where the photo or video was taken. Fact-checkers are able to assess the truth or falsity of a photo or video by combining these skills with other journalistic practices, like using research from experts, academics or government agencies.

We are paying close attention to how research develops and are interested in working with others in the industry to come up with solutions to deepfakes. We are also working closely with our Facebook AI Research lab to help identify this type of content. We are committed to working with our industry partners and with Congress to develop solutions to combat this issue.

61. On July 16, 2017, Facebook filed for a patent called "Socioeconomic Group Classification Based on User Features." The company stated that the technology would use data such as a Facebook user's age, travel history, homeownership status, and internet usage to predict the Facebook user's socioeconomic status. According to the patent, the algorithm would classify Facebook users into three categories: working class, middle class, or upper class.

- **Has Facebook implemented this technology?**

- **Does Facebook categorize users into socioeconomic groups?**

Facebook has not implemented the technology referenced in the United States or used it with respect to Facebook users for classification in any of the categories described above. Every Facebook user can view specific interests and categories derived from their activity on and off Facebook in their Ads Preferences control.

- **Does Facebook allow its partners to categorize users into socioeconomic groups (e.g., through "partner categories")?**

"Partner Categories" were targeting options that were based on data provided by third-party data providers. We announced in April that we would stop offering Partner Categories and as of October 1, they are no longer available.

- **What is the complete set of categories Facebook has to characterize its users?**

The specific number of categories that are used to decide what ads a person will see vary from person to person, depending on the interests and information that they have shared on Facebook, how frequently they interact with ads and other content on Facebook, the controls and choices they have implemented and other factors. Any person can see each of the specific interests we maintain about them for advertising by visiting Ads Preferences, which lets people see what interests we use to choose ads for them—and they can add or delete interests. We also provide more detailed information about how we use data to decide what ads to show to people in our "About Facebook Ads" page, at https://www.facebook.com/ads/about.

- **What is the complete set of partner categories offered by Facebook's third-party data partners?**

"Partner Categories" were targeting options offered by third-party data providers. We announced in April that we would stop offering this kind of targeting and as of October 1, Partner Categories are no longer available.

Getback. What is Facebook's official stance on hate speech regarding legally defined unprotected classes, such as children? Have you removed the requirement that you will only protect with your hate speech policy those classes of people that have been designated as protected classes in a legal context? Is that no longer Facebook's policy?

We recognize how important it is for Facebook to be a place where people feel empowered to communicate, and we take our role in keeping abuse off our service seriously. That is why we have developed a set of Community Standards that outline what is and is not allowed on Facebook. These standards are comprehensive—for example, content that might not be considered hate speech may still be removed for violating our bullying policies.

Under Facebook's hate speech policy, we remove attacks on groups of people based on protected characteristics, which we define as race, ethnicity, national origin, religious affiliation, sexual orientation, caste, sex, gender, gender identity, and serious disease or disability. Our guidelines apply globally and are not based on any specific country's laws. We also provide some protections for immigration status.

As noted, Facebook's Community Standards also prohibit attacks on individuals under our bullying and harassment policy, and when the person being targeted is a minor, we have a lower threshold for removal in order to protect the child.