

June 8, 2018

Chairman John Thune
Ranking Member Bill Nelson
U.S. Senate Committee on Commerce, Science, and Transportation
512 Dirksen Senate Building
Washington, D.C. 20510

Dear Chairman Thune, Ranking Member Nelson, and Members of the Committee:

Thank you for your questions for the record from the April 10, 2018 Hearing titled Facebook, Social Media Privacy, and the Use and Abuse of Data. Per your request, attached are the answers for the record for your questions.

Please note that we received over 2,000 questions from the Senate and House Committees before which we testified on April 10 and 11, 2018. We appreciate the extra time you gave us to respond to these questions. 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.

IMPORTANT -- PLEASE READ

DO NOT DETACH

United States Senate

Committee on Commerce, Science, and Transportation

Washington, D.C. 20510-6125

MEMORANDUM

Date of Hearing: April 10, 2018

Hearing: Facebook, Social Media Privacy, and the Use and Abuse of Data

Thank you for your recent testimony before the Senate Committee on Commerce, Science, and Transportation. The testimony you provided was greatly appreciated.

Attached are **post-hearing questions** pertaining to the above-mentioned hearing. As a courtesy, please submit a single document consolidating the posed questions followed by your answers for insertion in the printed hearing record. Your responses can be e-mailed to

[REDACTED]

Should the committee not receive your response within the time frame mentioned below or if the committee staffer assigned to the hearing is not notified of any delay, the committee reserves the right to print the posed questions in the formal hearing record noting your response was not received at the time the record was published.

Committee staffer assigned to the hearing: [REDACTED]

Phone: [REDACTED]

Date material should be returned: May 9, 2018.

Thank you for your assistance and, again, thank you for your testimony.

Questions from Chairman Thune

Question 1. In its April 2, 2018, response to the letter Sen. Wicker, Sen. Moran, and I sent you on March 19, 2018, Facebook committed to investigating all apps that potentially had access to the same type of data as Cambridge Analytica to identify other misuses of such data. Will you commit to having Facebook brief Commerce Committee staff on a periodic basis regarding the progress of these investigations and any future developments in Facebook's efforts to combat data misuse more generally?

We are in the process of investigating every app that had access to a large amount of information before we changed our Platform in 2014. The investigation process is in full swing, and it has two phases. First, a comprehensive review to identify every app that had access to this amount of Facebook data and to focus on apps that present reason for deeper investigation. And second, where we have concerns, we will conduct interviews, make requests for information (RFI)—which ask a series of detailed questions about the app and the data it has access to—and perform audits using expert firms that may include on-site inspections. We have large teams of internal and external experts working hard to investigate these apps as quickly as possible. To date thousands of apps have been investigated and around 200 apps have been suspended—pending a thorough investigation into whether they did in fact misuse any data. Where we find evidence that these or other apps did misuse data, we will ban them and let people know.

These apps relate to a handful of developers: Kogan, AIQ, Cube You, the Cambridge Psychometrics Center, and myPersonality, with many of the suspended apps being affiliated with the same entity. Many of these suspensions include apps that appear to be “test” apps that were never released to the public, and therefore would not have acquired significant user data, although our investigation into these apps is ongoing.

Additionally, we have suspended an additional 14 apps, which were installed by around one thousand people. They were all created after 2014, after we made changes to more tightly restrict our platform APIs to prevent abuse. However, these apps appear to be linked to AIQ, which was affiliated with Cambridge Analytica. So, we have suspended them while we investigate further. Any app that refuses to take part in or fails our audit will be banned.

We will commit to briefing your staff on future developments.

Question 2. Mr. Zuckerberg, as you know, Sen. Wicker, Sen. Moran, and I sent a letter to you on March 19th, requesting answers to several questions regarding Facebook's privacy practices. Facebook's general counsel sent a response letter on April 2nd that did not adequately answer some of the questions posed, saying that Facebook's review of the matter is ongoing. Will you commit to providing additional answers to our questions in writing in a timely manner as you learn more?

We responded to your questions to the best of our ability based on accessible data and information. Should additional or revised information related to the questions come to light, we respectfully request an opportunity to supplement or amend our response as needed.

Question 3. Mr. Zuckerberg, at the hearing you responded to over 20 questions from a number of Senators by saying that you would have to follow up at a later date. As you compile the promised information, please provide all such responses to these questions to Commerce Committee staff in addition to the Senator who posed the question.

Today we are submitting responses to the questions posed at the hearing requiring follow-up.

Question 4. Mr. Zuckerberg, given the concerns raised by a number of Senators that Facebook’s user agreement is too opaque to give users a real understanding of how their data may be used and how they can control their data privacy, do you intend to make any changes to the user agreement? If so, please summarize those changes and why you believe they will make the agreement more easily understood.

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 as people are using the Facebook service itself.

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, 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.

Improving people’s understanding of how digital services work is an industry-wide challenge that we are highly committed to addressing. That’s why, over the last 18 months, we’ve run a global 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 Paulo, 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.

Question 5. Mr. Zuckerberg, in the weeks since the revelations regarding Cambridge Analytica, the Committee has become aware that Facebook has surveyed users about whether they trust the company to safeguard their privacy. Please provide the Commerce Committee with the results of any such survey.

Privacy is at the core of everything we do, and our approach to privacy starts with our commitment to transparency and control. Our threefold approach to transparency includes, first, whenever possible, providing information on the data we collect and use and how people can control it in context and in our products. Second, we provide information about how we collect and use data in our user agreements and related educational materials. And third, we enable people to learn more about the specific data we have about them through interactive tools such as Download Your Information, which lets people download a file containing data that they may want to take to another service, and Access Your Information, a tool we are launching that will let people more easily access and manage their data on Facebook.

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. People can control the audience for their posts and the apps that can receive their data. They can see and delete the history of their activities on Facebook, and, if they no longer want to use Facebook, they can delete their account and the data associated with it. Of course, we recognize 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. 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 are launching a new settings menu that features core privacy settings in a single place. We are always working to help people understand and control how their data shapes their experience on Facebook.

Question 6. Mr. Zuckerberg, when did you personally become aware of Cambridge Analytica's breach of your policies in 2014-2015, and when did you personally become aware that Cambridge Analytica had not in fact deleted the data they obtained despite certifying otherwise?

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. (SCL)/Cambridge Analytica. As part of its investigation, Facebook contacted Kogan and Cambridge Analytica to investigate the allegations reflected in the reporting. Thereafter, Facebook obtained written certifications or confirmations from Kogan, GSR, and other third parties (including Cambridge Analytica and SCL) declaring that all such data they had obtained was accounted for and destroyed. In March 2018, Facebook received information from the media suggesting that the certification we received from SCL may not have been accurate and immediately banned SCL Group and Cambridge Analytica from purchasing advertising on our platform. Since then, Facebook has been actively investigating the issue, including pursuing a forensic audit of Cambridge Analytica, which is currently paused at the request of the UK Information Commissioner's Office (which is separately investigating Cambridge Analytica).

Mr. Zuckerberg did not become aware of allegations that Cambridge Analytica may not have deleted data about Facebook users obtained from Kogan's app until March of 2018, when these issues were raised in the media.

Question 7. On April 24, 2018, Facebook announced that it would institute an appeals process for posts that Facebook removes for violating its community standards. This process will initially only be available for posts that were removed for nudity/sexual activity, hate speech, or graphic violence. Why did Facebook decide to launch its appeals process for these categories? Prior to this new appeals process, did Facebook users have any recourse if their post was removed?

Prior to April 24, 2018, appeals generally were only available to people whose profiles, Pages, or Groups had been taken down, but we had not yet been able to implement an appeals process at the content level.

On April 24, we announced the launch of appeals for content that was removed for nudity/sexual activity, hate speech, and graphic violence. We focused on starting with these content violations initially based on feedback from our community.

We are working to extend this process further, by: supporting more violation types; giving people the opportunity to provide more context that could help us make the right decision; and making appeals available not just for content that was taken down, but also for content that was reported and left up.

Question 8. In your testimony, you discussed two typical business models employed by social media companies to make content available to users: an advertising-supported model and a subscription-based model. If Facebook were to shift from an advertising model to a subscription model, how much would consumers expect to pay in order to access Facebook content? Would you ever consider making such a shift? If not, why not?

Like many other free online services, we sell advertising space to third parties. Doing so enables us to offer our services to consumers for free. This is part of our mission to give people the power to build community and bring the world closer together.

Question 9. According to your testimony, Facebook has found that, while some users don't like advertisements, "people really don't like ads that aren't relevant" and the "overwhelming feedback we get from our community is that people would rather have us show relevant content." Can you elaborate on your basis for these statements about user preferences?

Part of Facebook's goal is to deliver the right content to the right people at the right time. This is just as true of posts and other content in users' News Feeds as it is for ads in their News Feed. And to choose the right ads Facebook listens to what feedback users provide. Users frequently provide feedback about what ads they want to see and don't want to see; they interact with ads positively (clicks, likes, comments, or shares) and negatively (by hiding the ad). Facebook takes all of this into consideration when selecting ads for its users.

In conjunction with this user feedback, Facebook has been working to better understand people's concerns with online ads. For example, Facebook has conducted multi-method, multi-market research surrounding ad blocking and personalization expectations among consumers.

And the take away from this has been that people don't like to see ads that are irrelevant to them or that disrupt or break their experience. Furthermore, people like to have control over the kinds of ads they see. For these reasons, Facebook seeks to provide users more relevant ads, as well as the tools to improve their control over which ads they see.

Question 10. You stated that “there is some discomfort ... with using information in making ads more relevant.” Why do you believe Facebook users feel this discomfort? Do you believe users would feel more comfortable if they had a clearer understanding of the relationship between their information, the relevance of the advertisements they are served, and Facebook’s ability to offer content without charging subscription fees?

We maintain our commitment to privacy by not telling advertisers who users are or selling people’s information to anyone. That has always been true. We think relevant advertising and privacy are not in conflict, and we’re committed to doing both well.

We believe targeted advertising creates value for people and advertisers who use Facebook. Being able to target ads to the people most likely to be interested in the products, service or causes being advertised enables businesses and other organizations to run effective campaigns at reasonable prices. This efficiency has particularly benefited small businesses, which make up the vast majority of the six million active advertisers on Facebook. That said, we are keenly aware of the concerns about the potential of our tools to be abused. That is why we are investing heavily in improving the security and integrity of our platform.

Separately, our core service involves personalizing all content, features and recommendations that people see on Facebook services. No two people have the same experience on Facebook or Instagram, and they come to our services because they expect everything they see to be relevant to them. If we were not able to personalize or select ads or other content based on relevance, this would fundamentally change the service we offer on Facebook—and it would no longer be Facebook.

We do not have a “business reason” to compromise the personal data of users; we have a business reason to protect that information. Our mission is to build community and bring the world closer together, but it is not enough to just connect people, we have to make sure those connections are positive. If people’s experiences are not positive—if we fail to maintain their trust—they will not use our services.

Question 11. Mr. Zuckerberg, how does Facebook determine whether and for how long to store user data or delete user data?

In general, when a user deletes their account, we delete things they have posted, such as their photos and status updates, and they won’t be able to recover that information later. (Information that others have shared about them isn’t part of their account and won’t be deleted.)

There are some limited exceptions to these policies: For instance, information can be accessed and preserved for an extended period when it is the subject of a legal request or obligation, governmental investigation, or investigations of possible violations of our terms or policies, or otherwise to prevent harm. We also retain information from accounts disabled for terms violations for at least a year to prevent repeat abuse or other term violations.

Question 12. Mr. Zuckerberg, you have discussed how a Facebook user can learn what data Facebook has collected about him or her. How can a non-user learn what data, if any, Facebook has collected about him or her?

If a person doesn't have a Facebook account but believes Facebook may have information about them, they can contact us to request a copy of their information. A contact form is available at <https://www.facebook.com/help/contact/180237885820953>. However, Facebook does not create profiles about or track web or app browser behavior of non-users.

Question 13. Does Facebook continue to track users who have turned off personalized ads? If so, why? Provide a list of uses Facebook makes of the data of users who have disabled personalized ads.

When people visit apps or websites that feature our technologies—like the Facebook Like or Comment button—our servers automatically log (i) standard browser or app records of the fact that a particular device or user visited the website or app (this connection to Facebook's servers occurs automatically when a person visits a website or app that contains our technologies, such as a Like button, and is an inherent function of Internet design); and (ii) any additional information the publisher of the app or website chooses to share with Facebook about the person's activities on that site (such as the fact that a purchase was made on the site). This is a standard feature of the Internet, and most websites and apps share this same information with multiple different third-parties whenever people visit their website or app. For example, the Senate Commerce Committee's website shares information with Google and its affiliate DoubleClick and with the analytics company Webtrends. This means that, when a person visits the Committee's website, it sends browser information about the visit to each one of those third parties. More information about how this works is available at <https://newsroom.fb.com/news/2018/04/data-off-facebook/>.

Question 14. Is Facebook's use of a user's data on the Facebook platform for targeted advertising a condition of using Facebook?

Users can't opt out of seeing ads altogether because selling ads is what keeps Facebook free, but they do have different options to control how their data can and can't be used to show them ads. They're all found in ad preferences, which allows users to turn off the use of all data collected from partners off Facebook to target ads.

Users can also decide which of their profile fields they want used for ad targeting in the Information section under "About you." Users can remove themselves from interests under "Your interests" and categories under "Your categories."

Question 15. Mr. Zuckerberg, on March 25, you took out several full-page ads in newspapers around the world in which you stated: "We're also investigating every single app that had access to large amounts of data before we fixed this," referring to your 2014 policy changes. You went on to say, "We expect there are others. And when we find them, we will ban them and tell everyone affected." How many other offending apps have you found so far? You mentioned, when you find offending apps, you will be notifying users. Please also provide a list of these apps to Congress.

See Response to Question 1.

Question 16. Mr. Zuckerberg, as you may know, Carol Davidsen, who in 2012 served as the Obama campaign’s director of data integration and media analytics, reportedly asserted that Facebook allowed the campaign to access users’ personal data “because they were on our side.” Did Facebook give preferential treatment to the Obama campaign with respect to data access in 2012? With respect to data access, did Facebook discriminate between the presidential campaigns in 2016?

Both the Obama and Romney campaigns had access to the same tools, and no campaign received any special treatment from Facebook. Likewise, we offered identical support to both the Trump and Clinton campaigns, and had teams assigned to both. Everyone had access to the same tools, which are the same tools that every campaign is offered.

Question 17. Since 2011, Facebook has been operating under a consent order issued by the Federal Trade Commission following agency charges that Facebook had deceived consumers by failing to keep privacy promises to them. You have indicated that—without prejudging the FTC’s decision to investigate the Cambridge Analytica incident—you do not believe the consent order is implicated in the current matter. Please explain why.

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. 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 apps they used. Facebook voluntarily changed this feature of the Platform in 2014, however.

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 ongoing review by an independent assessor (Sections IV and V). Facebook accurately represented the operation of its developer Platform and the circumstances under which people could share data (including friends data) with developers, honored the restrictions of all privacy settings that covered developer access to data, and implemented a comprehensive privacy program build on industry-leading controls and principles, which has undergone ongoing review by an independent assessor approved by the FTC.

Question 18. Initial media reports stated that 50 million Facebook users were impacted by the Cambridge Analytica incident, Facebook later reported that 87 million users were impacted. How did Facebook arrive at this number, and can we expect this number to rise?

Facebook users shared some data associated with approximately 87 million users with Kogan’s app, consisting of people who installed the app and the friends of those users whose settings permitted their data to be shared by their friends with apps. Facebook does not know how many of these users actually had data shared by Kogan with Cambridge Analytica, so this is a highly conservative estimate of the maximum number of users who could have been impacted. Several additional caveats apply to this figure:

- First, this figure does not include users who installed the app but have since deleted their Facebook account (since Facebook no longer has that information).
- Second, Facebook's counts of potentially affected friends of installers of the app are likely substantially higher than the "true" number of affected friends, because (a) the counts include any friend of any installer of the app during any time between when the app first became active on the Platform in November 2013 and when the app's access to friends data was limited in May 2015, even though the friend may not have been a friend when the app was actually installed by a relevant user; (b) the counts include any friend of any installer even if they changed their privacy settings during the relevant period to disallow sharing with apps installed by their friends (due to limited historical information about when or how users updated their settings), such that some of their data may not have been shared with the app; and (c) Facebook's counts include anyone who installed the app during its existence on Facebook's Platform, even if they installed the app at a time when its access to user data, including data from friends of installers, was more limited (due to limited historical information about when individual users installed the app).

In addition, it is worth noting that 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.

Question 19. Having discovered the improper data transfer to Cambridge Analytica in 2015, why did Facebook wait until 2018 to investigate or audit the data transfer to determine its full scope, including the type of data improperly transferred?

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, was accounted for and destroyed. Based on recent allegations, we have reopened our investigation into the veracity of these certifications and have hired a forensic auditor to conduct a forensic audit of Cambridge Analytica's systems. We are currently paused on the audit at the request of the UK Information Commissioner's Office request, which is conducting a regulatory investigation into Cambridge Analytica (based in the UK), and we hope to move forward with that audit soon.

Facebook banned Cambridge Analytica from our service. We understand that the company is now defunct.

Question 20. Mr. Zuckerberg, as you know, the Commerce Committee has been seeking to find a bipartisan path forward on net neutrality legislation. I believe bipartisan legislation is the best way to protect net neutrality and stop the partisan back-and-forth at the Federal Communications Commission over this issue. Will you commit to working with Congress to develop a bipartisan legislative solution to the issue of net neutrality?

Keeping the internet open for everyone is crucial. Not only does it promote innovation, but it lets people access information that can change their lives and gives voice to those who

might not otherwise be heard. For these reasons, Facebook supports net neutrality and is open to working with members of Congress and anyone else on a solution that will preserve strong net neutrality protections.

Questions from Senator Wicker

Question 1. Mr. Zuckerberg, during the hearing you confirmed that Facebook collects the call and text histories of its users that use Android phones. You also stated that Facebook only collects call and text histories if a consumer opts-in to this Facebook service.

Does Facebook collect the call and text history information of minors (13 to 17 years of age) that have Android phones and opt-in to this service?

If yes, does Facebook require parental consent for minors to be able to opt-in to this service?

How and in what manner does Facebook disclose to its users that it is collecting the call and text history information of those that opt-in to this service?

Call and text history logging is part of an opt-in feature that lets people import contact information to help them connect with people they know on Facebook and Messenger. We introduced the call and text history component of this feature for Android users several years ago, and currently offer it in Messenger and Facebook Lite, a lightweight version of Facebook, on Android.

Contact importers are fairly common among social apps and serve as a way to more easily find the people users want to connect with. They help users find and stay connected with the people they care about and provide them with a better experience across Facebook.

Before we receive call and text history from people, they specifically grant us permission to access this data on their device and separately agree to use the feature. If, at any time, they no longer wish to use this feature they can turn it off, and all previously shared call and text history shared via that app is deleted. People can also access information they previously imported through the Download Your Information tool.

We've reviewed this feature to confirm that Facebook does not collect the content of messages—and will delete all logs older than one year. In the future, people will only upload to our servers the information needed to offer this feature—not broader data such as the time of calls. We do allow people from 13 to 17 to opt into this service. However, we do take other steps to protect teens on Facebook and Messenger:


- We provide education before allowing teens to post publicly.
- We don't show search results based on specific profile data (high school, birthday/age, and hometown, or current city) of teens to unconnected adults when the adults search on Facebook.
- Unconnected adults can't message minors who are 13-17.
- We have age limits for advertisements. For example, ads for dating sites, financial services and other products or services are gated to users under 18. We've also helped many teenagers with information about bullying prevention campaigns and online safety tips, including creating a new website full of privacy and safety resources for teens: <https://www.facebook.com/safety/youth>

Question 2. Is the data Facebook collects from call and text histories of its users that have Android phones used for targeted advertising purposes?

No, Facebook does not use SMS history to target interest-based ads. Instead, call and text history logging is part of an opt-in feature for people using Messenger or Facebook Lite on Android. This helps Facebook users find and stay connected with the people they care about and provides them with a better experience across Facebook. This feature does not collect the content of users' calls or text messages.

Question 3. When a user uploads his or her contact list, Facebook collects the phone numbers of the user's contacts. Please provide all details regarding what Facebook does with the phone numbers of the users' contacts, including with whom Facebook shares those numbers, whether Facebook creates or updates profiles that associate these numbers with people's names, and how long Facebook stores those numbers.

Facebook allows people to upload, sync, and import their contacts, typically using permissions that are enabled by major operating systems like Apple's iOS and Google Android. When people use the contact upload tool, they see prompts explaining what data will be collected:





See Who's On Facebook

When you choose to find friends on Facebook, we'll use and securely store information about your contacts, including things like names and any nicknames; contact photo; phone numbers and other contact or related information you may have added like relation or profession; as well as data on your phone about those contacts. This helps Facebook make recommendations for you and others, and helps us provide a better service. You're always able to manage or delete contacts you share with Facebook. You can turn off contact uploading in settings.

You may have business and personal contacts in your phone. Please only send friend requests to people you know personally who would welcome the invite.

We use this information that people choose to share for a variety of purposes, including to provide, personalize, and improve our products; provide measurement, analytics, and other business services; promote safety and security; to communicate with people who use our services; and to research and innovate to promote the social good. We provide more information in our Data Policy about these uses as well. People can view and manage their contact uploads using our Contacts Uploading tools, available at <https://www.facebook.com/help/355489824655936>.

Question 4. There have been reports that Facebook can track a user's internet-browsing activity even after the user has logged off of the Facebook platform. Can you confirm whether or not this is true?

If yes, how does Facebook disclose to its users that it is engaging in this type of tracking or data collection activity when a user has logged off of the Facebook platform?

When people visit apps or websites that feature our technologies—like the Facebook Like or Comment button—our servers automatically log (i) standard browser or app records of the fact that a particular device or user visited the website or app (this connection to Facebook's servers occurs automatically when a person visits a website or app that contains our technologies, such as a Like button, and is an inherent function of Internet design); and (ii) any additional information the publisher of the app or website chooses to share with Facebook about the person's activities on that site (such as the fact that a purchase was made on the site). This is

a standard feature of the Internet, and most websites and apps share this same information with multiple different third-parties whenever people visit their website or app. For example, the Senate Commerce Committee’s website shares information with Google and its affiliate DoubleClick and with the analytics company Webtrends. This means that, when a person visits the Committee’s website, it sends browser information about their visit to each one of those third parties. More information about how this works is available at <https://newsroom.fb.com/news/2018/04/data-off-facebook/>.

When the person visiting a website featuring Facebook’s tools is not a registered Facebook user, Facebook does not have information identifying that individual, and it does not create profiles for that individual.

We use the browser and app logs that apps and websites send to us—described above—in the following ways for non-Facebook users. First, these logs are critical to protecting the security of Facebook and to detecting or preventing fake account access. For example, if a browser has visited hundreds of sites in the last five minutes, that’s a sign the device might be a bot, which would be an important signal of a potentially inauthentic account if that browser then attempted to register for an account. Second, we aggregate those logs to provide summaries and insights to websites and apps about how many people visit or use their product, or use specific features like our Like button—but without providing any information about a specific person. We do not create profiles for non-Facebook users, nor do we use browser and app logs for non-Facebook users to show targeted ads from our advertisers to them or otherwise seek to personalize the content they see. However, we may take the opportunity to show a general ad that is unrelated to the attributes of the person or an ad encouraging the non-user to sign up for Facebook.

When the individual is a Facebook user, we are also able to use this information to personalize their experiences on Facebook, whether or not they are logged out, but we will not target ads to users relying on this information unless the user allows this in their privacy settings. We do not sell or share this information with third-parties.

Question 5. Mr. Zuckerberg, if a user deletes his or her Facebook account, does Facebook still track that person on non-Facebook websites and applications?

Facebook does not create profiles or track website visits for people without a Facebook account. See Response to Question 4.

Question 6. Mr. Zuckerberg, you asserted in the hearing that “the expectations that people have” regarding use of data by ISPs are somewhat different than for edge platforms like yours. In fact, a survey by Peter D. Hart showed that 94% of consumers want their online data to be subject to a consistent level of privacy protection across the Internet and that ISPs and edge providers should be treated alike. Do you have any consumer survey data or empirical evidence to support your assertion that consumers expect or want different privacy protections for ISPs? If so, please provide the consumer survey data or empirical evidence that supports your assertion.

We believe that everyone should enjoy strong privacy protections, but we also realize that people have different expectations based on the context in which their information is provided. For instance, a person who orders shoes from a mail-order catalog would expect the retailer to know what is in the box that he is being sent. But the customer would not expect the post office

to know what he or she has purchased just because it is delivering the box. Because of this difference in expectations, the post office may need to do more to inform people if it intends to inspect packages it delivers and to give people control if it intends to use the information it learns in other ways.

Consistent with this difference, experts have observed, “The context in which broadband customers share private information with [Internet service] providers is specific and accompanied by cabined expectations: the customers share the information with [Internet service] providers to facilitate provision of a service for which they have contracted. The information is therefore most appropriately thought of as a loan to, rather than transferred to, broadband providers.”¹ In contrast, a group of leading academic experts led by Prof. Nick Feamster of Princeton University observed that people may have access to only one or a few ISPs and simply expect those ISPs to deliver their communications. Such a person has no choice about whether to send his or her traffic over an ISP’s network, whereas a “user may simply elect not to provide certain personal information or data to a social network, or even to not use the social network at all.”² Other experts have observed that edge providers’ collection of information is generally more expected because it is related to the services those companies provide.³

In our own services, Facebook needs to have a different understanding of a person’s data than an ISP would. For instance, when someone adds information to their profile or likes a Page on Facebook, we must have access to that information in order to display it and use it to personalize that person’s experience. People would not necessarily anticipate that other companies would have access to that information, which is why we do not sell people’s information to advertisers and are increasing our efforts to guard against misuse of people’s Facebook information by third parties. It is also why we provide people with the ability to turn off advertising based on the apps and websites they use outside of our service, and we are investing in enhanced transparency and control around this through our recent announcement of a new tool, Clear History, that we are building.

Although we have not reviewed the detailed survey by Mr. Hart to which the question refers, we understand that it focused on a different question than Mr. Zuckerberg’s testimony. Specifically, Mr. Hart’s survey asked people whether they believe that information should be subject to protection; this is different from asking whether people have different expectations about what information Facebook will receive when they put information on their Facebook profile, as compared to what information their Internet service provider will receive when they take the same action.

¹ Comments of New America Foundation, FCC 16-39, at 7.

² Comments of Nick Feamster, et al., FCC 16-39, at 3.

³ Paul R. Gaus, *Only the Good Regulations Die Young: Recognizing the Consumer Benefits of the FCC’s Now-Defunct Privacy Regulations*, 18 Minn. J. Law, Sci. & Tech. 713 (2017) (“Defining the internet consumer seems like a facile task, but it must incorporate how the person uses digital devices to connect to the internet and use content. In the context of ISPs, the digital consumer conforms to a traditional definition in that the consumer purchases ISP services to access the internet. In the space of edge providers, the digital consumer engages in traditional retail, watches content, interacts with others via social media, and performs a plethora of other activities that provide a telling summary about a person’s life.”).

Questions from Senator Blunt

Question 1. Does Facebook collect user data through cross-device tracking, and does this include off-line data (offline data defined as that which is not directly contributed by a user through usage of features of the Facebook app)?

Yes, Facebook's Data Policy specifically discloses that we associate information across different devices that people use to provide a consistent experience wherever they use Facebook.

Facebook's services inherently operate on a cross-device basis: understanding when people use our services across multiple devices helps us provide the same personalized experience wherever people use Facebook—for example, to ensure that people's News Feeds or profiles contains the same content whether they access our services on their mobile phone or in a desktop computer's web browser.

In support of those and other purposes, we collect information from and about the computers, phones, connected TVs and other web-connected devices our users use that integrate with our Products, and we combine this information across a user's different devices. For example, we use information collected about a person's use of our Products on their phone to better personalize the content (including ads) or features they see when they use our Products on another device, such as their laptop or tablet, or to measure whether they took an action in response to an ad we showed them on their phone or on a different device.

Information we obtain from these devices includes:

- Device attributes: information such as the operating system, hardware and software versions, battery level, signal strength, available storage space, browser type, app and file names and types, and plugins.
- Device operations: information about operations and behaviors performed on the device, such as whether a window is foregrounded or backgrounded, or mouse movements (which can help distinguish humans from bots).
- Identifiers: unique identifiers, device IDs, and other identifiers, such as from games, apps or accounts people use, and Family Device IDs (or other identifiers unique to Facebook Company Products associated with the same device or account).
- Device signals: Bluetooth signals, and information about nearby Wi-Fi access points, beacons, and cell towers.
- Data from device settings: information a user allows us to receive through device settings they turn on, such as access to their GPS location, camera, or photos.
- Network and connections: information such as the name of a user's mobile operator or ISP, language, time zone, mobile phone number, IP address, connection speed and, in some cases, information about other devices that are nearby or on their network, so we can do things like help them stream a video from their phone to their TV.

- Cookie data: data from cookies stored on a user's device, including cookie IDs and settings. More information is available at <https://www.facebook.com/policies/cookies/> and <https://help.instagram.com/1896641480634370?ref=ig>.

Advertisers, app developers, and publishers can send us information through Facebook Business Tools they use, including our social plug-ins (such as the Like button), Facebook Login, our APIs and SDKs, or the Facebook pixel. These partners provide information about a person's activities off Facebook—including information about their device, websites they visit, purchases they make, the ads they see, and how they use their services—whether or not they have a Facebook account or are logged into Facebook. For example, a game developer could use our API to tell us what games a person plays, or a business could tell us about a purchase a person made in its store. We also receive information about a person's online and offline actions and purchases from third-party data providers who have the rights to provide us with that person's information.

We use the information we have to deliver our Products, including to personalize features and content (including a person's News Feed, Instagram Feed, Instagram Stories, and ads) and make suggestions for a user (such as groups or events they may be interested in or topics they may want to follow) on and off our Products. To create personalized Products that are unique and relevant to them, we use their connections, preferences, interests and activities based on the data we collect and learn from them and others (including any data with special protections they choose to provide); how they use and interact with our Products; and the people, places, or things they're connected to and interested in on and off our Products.

For example, if people have shared their device locations with Facebook or checked into a specific restaurant, we can show them ads from an advertiser that wants to promote its services in their area or from the restaurant. We use location-related information—such as a person's current location, where they live, the places they like to go, and the businesses and people they're near—to provide, personalize and improve our Products, including ads, for them and others. Location-related information can be based on things like precise device location (if a user has allowed us to collect it), IP addresses, and information from their and others' use of Facebook Products (such as check-ins or events they attend). We store data until it is no longer necessary to provide our services and Facebook Products, or until a person's account is deleted—whichever comes first. This is a case-by-case determination that depends on things like the nature of the data, why it is collected and processed, and relevant legal or operational retention needs. 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 someone (information such as a person's name or email address that by itself can be used to contact them or identifies who they are) unless they give us permission. For example, we provide general demographic and interest information to advertisers (for example, that an ad was seen by a woman between the ages of 25 and 34 who lives in Madrid and likes software engineering) to help them better understand their audience. We also confirm which Facebook ads led users to make a purchase or take an action with an advertiser.

Question 2. Cross-device data collection allows for data and user profile meshing that the average users are likely not cognizant of. Last year, the Federal Trade Commission flagged

cross-device tracking as a possible concern, due to the fact that most companies do not explicitly discuss cross-device tracking in their privacy policies.

Does Facebook disclose its collection methods across each applicable device, and if so, do you offer your users choices about how cross-device activity is tracked?

See Response to Question 1.

Question 3. Are users required to resubmit their permissions for each separate device that utilizes the Facebook app, or are user permissions blanketed across devices?

Mobile operating systems like Google's Android and Apple's iOS have device-specific access controls implemented at the operating system level.

Question 4. Facebook has been criticized for previous versions of its mobile application on Android devices, and the manner in which permissions were bundled without the ability to grant or deny each permission individually. I understand that Facebook and Android have updated their platforms, allowing more latitude for users to review permissions individually.

What is the technical and commercial purpose of bundling permissions?

Android and other operating systems (like Apple's iOS) control the way device permissions work. Facebook can't, for example, request permissions in a way that's not permitted on an Android device. Accordingly, where permitted by the operating system, we generally ask for permission in-context—for example, requesting access to a device's camera roll when someone uses a feature that requires it. But for other permissions, on the Android operating system, we must list all of the permissions that various features might require at the point when a person installs the app, even if we do not intend to use those permissions until those features are accessed.

On our website, we explain more about permissions that we request and provide examples of how they are used. You can find this information at <https://www.facebook.com/help/210676372433246>.

Question 5. How does your company prioritize transparency and choice for users in the way that it collects and aggregates user data?

Our approach to transparency is threefold.

First, we provide information about the data we collect and use and how people can control it in context as people use Facebook. Research overwhelmingly demonstrates that in-product controls and education are the most meaningful to people and the most likely to be read and understood.

Second, we provide information about how we collect and use data in our user agreements and related educational materials. These materials include our Data Policy, which we updated recently to make it more detailed and easier to understand, and Privacy Basics, a series of short, interactive guides that answer some of the most common questions we receive about privacy.

Third, we enable people to learn more about the data we collect through interactive tools such as Download Your Information, which lets people download a file containing data that they may want to take to another service, and Access Your Information, a tool we've launched for people to more easily access and manage their data on Facebook.

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. People can control the audience for their posts and the apps that can receive their data. They can control the people, Pages, Groups, and Events they connect to, and how they see content from those connections in their News Feeds. They can provide feedback on every post they see on Facebook—feedback, for example, that they want to see less of a particular kind of post or fewer posts from a particular person or Page. They can see and delete the history of their activities on Facebook, and, if they no longer want to use Facebook, they can delete their account and the data associated with it.

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 are launching a new settings menu that features core privacy settings in a single place.

Questions from Senator Fischer

Question 1. Given ongoing user privacy concerns, American consumers are asking for a public dialogue about the *purposes* for which Facebook uses their personal data. However, a meaningful conversation cannot happen until users also understand the *sources* from which their data is gleaned, and the scope of the specific data—which characteristics, attributes, labels, or categories of data points—being collected and utilized. How many categories (i.e. attributes, factors, labels, or data points) does Facebook collect about particular users?

As explained in our Data Policy, we collect three basic categories of data about people:

- (1) data about things people do and share (and who they connect with) on our services;
- (2) data about the devices people use to access our services; and
- (3) data we receive from partners, including the websites and apps that use our business tools.

As far as the amount of data we collect about people, the answer depends on the person. People who have only recently signed up for Facebook have usually shared only a few things—such as name, contact information, age, and gender. Over time, as people use our products, we receive more data from them, and this data helps us provide more relevant content and services. That data will fall into the categories noted above, but the specific data we receive will, in large part, depend on how the person chooses to use Facebook. For example, some people use Facebook to share photos, so we receive and store photos for those people. Some people enjoy watching videos on Facebook; when they do, we receive information about the video they watched, and we can use that information to help show other videos in their News Feeds. Other people seldom or never watch videos, so we do not receive the same kind of information from them, and their News Feeds are likely to feature fewer videos.

The data we have about people also depends on how they have used our controls. For example, people who share photos can easily delete those photos. The same is true of any other kind of content that people post on our services. Through Facebook’s Activity Log tool, people can also control the information about their engagement—i.e., their likes, shares and comments—with other people’s posts. The use of these controls of course affects the data we have about people.

We recently announced improvements to our Download Your Information tool, as well as a new feature that makes it easier for people to see the information that’s in their account on Facebook. These recently-expanded tools for accessing your information will allow people to see their data, delete it, and easily download and export it.

Question 2. How many categories, as the term is described above, are used to construct the digital profiles that Facebook utilizes to direct ads to particular 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, 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 to edit or delete these 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>.

We use data about things people do on Facebook, such as the Pages they like, to associate “interests” with their accounts, and we enable advertisers to reach audiences—i.e., groups of people—that share those interests. For example, if a person has liked Pages about baseball, we might associate them with interests called “baseball” or “sports.”

We use data from devices (such as location data) to help advertisers reach people in particular areas. For example, if people have shared their device locations with Facebook or checked into a specific restaurant, we can show them ads from an advertiser that wants to promote its services in their area or from the restaurant.

We also help advertisers reach people who have given the advertiser their contact information or who have used the advertiser’s website or app. For example, advertisers can send us a hashed list of email addresses of people they would like to reach on Facebook. 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). The data we use to show ads to people depends on the data we have received from people. Again, for people who are new to Facebook, we may have minimal data that we can use. For people who have used our services for longer, we likely have more data, but the amount of data will depend on the nature of that use and how they have used our controls.

As noted above, in addition to general controls—such as Activity Log—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 as 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. People’s use of these controls will, of course, affect the data we use to show them ads.

Question 3. If a user opts out of directed advertising, does Facebook halt collection of all such data?

We give people a number of controls over the data we use to show them ads. These controls apply to our use of data to show people ads; they do not apply to the collection of data, because the same core data sets are used to ensure the safety and security of our platform and to provide our core service to our users. As noted above, people can see and control the advertising “interests” and “behaviors” we have associated with their accounts to show them ads. They can choose not to see ads from a particular advertiser or not to see ads based on their use of third-party websites and apps. They also can choose not to see ads off Facebook that are based on the interests we derive from their activities on Facebook.

Question 4. If a user opts out of directed advertising, does Facebook delete all such data that was previously stored? Alternatively, does Facebook instead simply stop utilization of that data for directed advertising purposes?

Our advertising controls apply only to the use of data for targeting and selecting ads. Using these controls does not result in deletion of data, because the same core data sets are used to ensure the safety and security of our platform and to provide our core service to our users. This is consistent with industry practice. For example, the Digital Advertising Alliance’s Self-Regulatory Principles set the industry standard for the collection and use of data for online behavioral advertising and related practices. Those principles require companies to offer controls over the use of data for advertising purposes. Companies are not required to stop collecting data from opted-out users or to delete previously collected data. Please note, however, that when a person removes an “interest” or “behavior” in Ad Preferences, that interest or behavior is permanently removed from the person’s ad profile; it will not be recreated even if the person subsequently engages in activities that otherwise would have resulted in the creation of the interest or behavior.

Question 5. When users download a copy of their Facebook data, as Facebook has recently enabled, is all ad targeting data included in that file?

Our Download Your Information or “DYI” tool is Facebook’s data portability tool and was launched many years ago to let people access 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 are also launching Access Your Information, a screenshot of which was included in our April 27, 2018 letter to you. This is a secure 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.

Questions from Senator Moran

Question 1. Cambridge Analytica had access to data on up to 87 million Facebook users because 270,000 individuals participated in a personality quiz that also exposed their friends' data. While I understand how the 270,000 individuals could have given their express consent, can you please walk me through how the many millions of friends could have given their "affirmative express consent" for their data to be shared with a third party as is required by the 2011 consent decree—when they were unaware that a friend of theirs was even participating in a personality quiz?

At the outset, we do not know what data Kogan may have shared with Cambridge Analytica. Our investigation into these matters is ongoing, and we are paused on investigating Cambridge Analytica directly (or conducting a forensic audit of its systems) due to the request of the UK Information Commissioner's Office, which is separately investigating Cambridge Analytica, a UK entity. The best information to date also suggests only US user data was shared by Kogan with Cambridge Analytica.

As was the practice of other online or mobile app platforms, at that time, people on Facebook were able to take their data and data their friends had shared with them off of Facebook to apps they authorized to obtain a broader range of experiences than were available on Facebook. But people could not share data for friends whose privacy settings did not permit their data to be shared by their friends with apps—and no data was shared with Kogan's app in violation of friends' settings. The 2011 consent decree requires Facebook to get affirmative express consent for materially expanding the audience of a user's existing privacy settings. No privacy settings were expanded or exceeded on Platform, and the consent order therefore does not apply here.

Approximately 300,000 Facebook users worldwide installed Kogan's app. For the majority of these users, the app requested consent to access the following data fields associated with the user and with the friends of the user: Public profile data, including name and gender; Birthdate; "Current city" in the "About" section of the user's profile, if provided; and Facebook Pages liked.

For a small subset of users, it appears that the app also requested consent to access users' Facebook messages (fewer than 1,500 individuals, based on current information) and to posts that appeared in the user's News Feed or Timeline (approximately 100 individuals, based on current information)—but only for users who installed the app. For a small subset of users (fewer than 1,500 individuals, based on current information), it appears that the app also requested consent to access the hometowns that the users' friends had specified in the "About" section of their profiles. And for a handful of people (fewer than 10) who appear to be associated with Kogan/GSR, the app requested consent to email address and photos.

Question 2. According to Facebook's March 21st press release, one of the six changes that Facebook initially offered to "crack down on platform abuse" was to reward outside parties who find vulnerabilities through its bug bounty program. My subcommittee has held hearings and met with interested stakeholders on these types of data security solutions along with other cyber vulnerability disclosure programs. One concern I have regarding the utility of this approach is that vulnerability disclosure programs are normally geared to identify unauthorized access to data, not point out data sharing arrangements that likely

harm users but technically abide by the complex consent agreements Facebook pushes on their users. Could you please explain how Facebook's expansion of its bug bounty program will prevent future data sharing issues with its associated applications from occurring?

The Data Abuse Bounty will reward people with first-hand knowledge and proof of cases where a Facebook platform app collects and transfers people's data to another party to be sold, stolen or used for scams or political influence. We'll review all legitimate reports and respond as quickly as possible when we identify a credible threat to people's information. If we confirm data abuse, we will shut down the offending app and, if necessary, take legal action against the company selling or buying the data. We'll pay a bounty to the person who reported the issue, or allow them to donate their bounty to a charity, and we'll also alert those we believe to be affected. We also encourage our users to report to us content that they find concerning or that results in a bad experience, as well as other content that may violate our policies. We review these reports and take action on abuse, like removing content and disabling accounts.

Question 3. Facebook has confirmed alterations to its terms and conditions shifting more than 1.5 billion of its user from contracts with the international headquarters in Ireland to Facebook Inc. in the United States, thereby removing these users from the protections they would otherwise receive from the Europeans Union's General Data Protection Regulation (GDPR). With the recent scrutiny that Facebook has faced about its data collection, sharing, and security polices what is the justification for moving approximately 1.5 billion Facebook user away from the more stringent rules of the European Union's GDPR?

We will offer everyone who uses Facebook the same controls and settings, no matter where they live. However, the GDPR creates some specific requirements that do not apply in the rest of the world, for example the requirement to provide contact information for the EU Data Protection Officer or to specify legal bases for processing data. We are also looking to be more responsive to regional norms and legal frameworks going forward, and want to have the flexibility to work with local regulators, which is possible with this new model. At the same time, we are changing the provisions in our Facebook, Inc. terms in our user agreements outside the United States to allow people in other countries to file lawsuits against Facebook in their home country, rather than in courts in the US. This transition was part of a continued effort to be locally responsive in countries where people use our services.

Question 4. During your testimony, you noted that Facebook cooperates with law enforcement in two instances, where there is an "imminent threat of harm" or when law enforcement reaches out to the company with a "valid request for data." In December 2017, the Chicago Police Department announced that it had arrested fifty people who were utilizing Facebook private group features in order to communicate and facilitate illegal firearm and drug transactions. Several national news outlets reported that Facebook was not helpful in regards to this investigation and Chicago Police Superintendent Eddie Johnson was later quoted in response to media inquiries as saying "Quite frankly, they haven't been very friendly to law enforcement to prevent these things." What specific policies and procedures does Facebook currently have in place to aid law enforcement agencies in gaining access to relevant information that indicates a clear threat to public safety?

We recognize there are serious and evolving threats to public safety and that law enforcement has an important responsibility to keep people safe. Our legal and safety teams

work hard to respond to legitimate law enforcement requests while fulfilling our responsibility to protect people's privacy and security. We have a global team that strives to respond within minutes to emergency requests from law enforcement. In the second half of 2017, for example, we provided information in response to nearly 78% of the 1,808 requests for emergency disclosures that we received from US law enforcement agencies. Facebook also reaches out to law enforcement whenever we see a credible threat of imminent harm. We use automated and manual review and also rely on users to help by reporting violating accounts or content. We are also working with law enforcement and others to improve our ability to find users at risk of harming themselves or others. We also disclose information in response to law enforcement requests in accordance with our terms of service and applicable law. In the second half of 2017, for example, we disclosed data in response to 85% of law enforcement requests from agencies in the US. Facebook regularly produces a report on government requests to help people understand the nature and extent of these requests and the policies and processes in place to handle them.

In addition, we cooperated with the Chicago Police Department's investigation that led to the December 2017 arrests. We reached out immediately after we learned of the comments referenced in your question, and they issued follow-up statements indicating that we reached out and were planning to provide training. We followed up by training over 100 Chicago-area law enforcement officers in a working group hosted by the FBI and US Attorney's Office. We also met separately with the Chicago Police unit that conducted the investigation to make sure they understood Facebook's policies, how to submit requests to us, and how we could help them through additional training and support.

Question 5. What specifically qualifies as a "valid request for data," which is required to gain access to information?"

We disclose account records in accordance with our terms of service and applicable law, including the federal Stored Communications Act. In the United States, a valid subpoena issued in connection with an official criminal investigation is required to compel the disclosure of basic subscriber records. A court order issued under 18 U.S.C. § 2703(d) is required to compel the disclosure of certain records or other information pertaining to the account, not including contents of communications. A search warrant issued under the procedures described in the Federal Rules of Criminal Procedure or equivalent state warrant procedures upon a showing of probable cause is required to compel the disclosure of the stored contents of any account. Facebook may also voluntarily disclose information to law enforcement where we have a good faith reason to believe that the matter involves imminent risk of serious physical injury or death.

Question 6. How does Facebook determine what rises to an imminent threat of harm and does that determination change the threshold for deciding whether to respond to a law enforcement data request?

Facebook discloses account records in accordance with our terms of service and applicable law, including the federal Stored Communications Act. The law permits Facebook to voluntarily disclose information to law enforcement where we have a good faith reason to believe that the matter involves imminent risk of serious physical injury or death. Our law enforcement response team receives and responds to emergency data requests around the clock and from around the globe based on our timely and careful review of information submitted by law enforcement and any other relevant facts. We also rely on experience and input from law

enforcement, safety organizations, and industry to identify and respond to potential threats of harm.

Question 7. Facebook has made a big deal about users' ability to request and download the data that Facebook has compiled about the user. But that downloaded data does not include data such as the list of the websites Facebook users have visited that is collected by Facebook. Why is that the case, and when will Facebook make this information available to users? What other information about Facebook users is not available for download?

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 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, will 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, delete this information from their account, and turn off Facebook's ability to store it associated with their account going forward.

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.

Questions from Senator Sullivan

Question 1. In the hearing, the topics of anticompetitive consolidation and the enormous market capitalization of tech companies such as Facebook were frequently raised. Recent calculations value the four largest tech companies' capitalization at \$2.8 trillion dollars, which is a staggering 24% of the S&P 500 Top 50, close to the value of every stock traded on the Nasdaq in 2001, and to give a different perspective, approximately the same amount as France's current GDP. At what point, from an antitrust perspective, is Facebook simply too big? Would you say that your size inhibits the "next Facebook"?

In Silicon Valley and around the world, new social apps are emerging all the time. The average American uses eight different apps to communicate with their friends and stay in touch with people. There is a lot of choice, innovation, and activity in this space, with new competitors arising all the time. Facebook's top priority and core service is to build useful and engaging products that enable people to connect, discover and share through mobile devices and personal computers. Given its broad product offerings, Facebook faces numerous competitors, competing to attract, engage, and retain users, to attract and retain marketers, and to attract and retain developers who build compelling mobile and web applications. For instance, if you want to share a photo or video, you can choose between Facebook, DailyMotion, Snapchat, YouTube, Flickr, Twitter, Vimeo, Google Photos and Pinterest among many other services. Similarly, if you are looking to message someone, just to name a few, there's Apple's iMessage, Telegram, Skype, Line, Viber, WeChat, Snapchat and LinkedIn—as well as the traditional text messaging services your mobile phone carrier provides. Equally, companies also have more options than ever when it comes to advertising—from billboards, print and broadcast, to newer platforms like Facebook, Spotify, Twitter, Google, YouTube, Amazon, or Snapchat. Facebook represents a small part (in fact, just 6%) of this \$650 billion global advertising ecosystem and much of that has been achieved by helping small businesses—many of whom could never have previously afforded newspaper or TV ads—to cost-effectively reach a wider audience.

Question 2. Senator Peters asked if Facebook extracts audio from its users to enhance personal data profiles, to which you responded no—is that the case? There are countless anecdotes about this exact situation. Would you characterize these as coincidence or is targeted advertising just that effective?

To be crystal clear on this point: Facebook does not use users' phone's microphone or any other method to extract audio to inform ads or to determine what they see in their News Feed. Facebook show ads based on people's interests and other profile information—not what users are talking out loud about. Facebook only accesses users' microphone if the user has given our app permission and if they are actively using a specific feature that requires audio (like voice messaging features).

Question 3. As you are aware, children are increasingly active users of technology. Do you have concerns generally about children's increased use, in many cases that rises to the level of addiction, of electronics? And more specifically, since I'm very interested in the issue of individual privacy rights, what are your thoughts on the data footprint of children being collected?

We take the privacy, safety, and security of all those who use our platform very seriously and when it comes to minors (13 to 18 years old), we provide special protections and resources.

We also provide special protections for teens on Facebook and Messenger. We provide education before allowing teens to post publicly. We don't show search results based on specific profile data (high school, birthday/age, and hometown, or current city) of teens to unconnected adults when the adults search on Facebook. Unconnected adults can't message minors who are 13-17. We prohibit search engines off Facebook from indexing minors' profiles. And, we have age limits for advertisements. For example, ads for dating sites, financial services and other products or services are gated to users under 18.

We provide special resources to help ensure that they enjoy a safe and secure experience. For example, we recently announced the launch of our Youth Portal, which is available in 60 languages at facebook.com/safety/youth. This portal is a central place for teens that includes:

- **Education:** Information on how to get the most out of products like Pages, Groups, Events, and Profile, while staying safe. Plus, information on the types of data Facebook collects and how we use it.
- **Peer Voices:** First person accounts from teens around the world about how they are using technology in new and creative ways.
- **Ways to control your experience:** Tips on things like security, reporting content, and deciding who can see what teens share.
- **Advice:** Guidelines for how to safely get the most out of the internet.

Instagram also will be providing information to teens to show them where they can learn about all of the tools on Instagram to manage their privacy and stay safe online, including how to use the new Access and Download tools to understand what they have shared online and learn how to delete things they no longer want to share. We are also making this information available in formats specifically designed for young users, including video tutorials for our privacy and safety tools, and teen-friendly FAQs about the Instagram Terms of Use, Data Policy, safety features, and Community Guidelines.

Instagram has also launched new content on Instagram Together, including videos and FAQs about privacy controls; information on how to use safety features, including comment controls, blocking accounts, reporting abuse, spam, or troubling messages; information on responsible social media use; and FAQs about safety on Instagram. We will be reaching out to users under 18 on Instagram to encourage them to learn more on Instagram Together.

Further, we have content restrictions and reporting features for everyone, including minors. We have Community Standards that prohibit hate speech, bullying, intimidation, and other kinds of harmful behavior. We encourage people to report posts and rely on our team of content reviewers around the world to review reported content. Our reviewers are trained to look for violations and enforce our policies consistently and as objectively as possible. When reviewed by our team, we hide certain graphic content from users under 18 (and include a warning for adults). We are also working to improve our ability to get our community help in real time, especially in instances where someone is expressing thoughts of suicide or self-harm,

by expanding our use of proactive detection, working with safety experts and first-responders, and dedicating more reviewers from our Community Operations team.

In addition, with 9 out of 10 children under the age of 13 in the United States able to access a tablet or smartphone and 2 out of 3 with their own device, and parents seeking greater control over who connects with their children, the content they see and the time they spend online, we are committed to working with parents and families, as well as experts in child development, online safety and children's health and media, to ensure we are building better products for families.

That is why we're committed to both continued research and to building tools that promote meaningful interactions and help people manage their time on our platform.

Indeed, as we built Messenger Kids, we worked closely with leading child development experts, educators, and parents to inform our decisions. Our advisors include experts in the fields of child development, online safety, and children's media currently and formerly from organizations such as the Yale Center for Emotional Intelligence, Connect Safely, Center on Media and Child Health, Sesame Workshop and more. The app does not have ads or in app purchase and we recently added Sleep Mode which gives the parent the ability to set parameters on when the app can be used. Messenger Kids collects only a limited amount of information. Additionally, when a Messenger Kids user turns 13, which is the minimum age to join Facebook, they don't automatically get a Facebook account.

We recently launched a Parents Portal and Youth Portal, which are both focused on fostering conversations around online safety and giving parents and young people access to the information and resources they need to make informed decisions about their use of online technologies.

Question 4. I'm very proud to be a cosponsor of the recently passed SESTA legislation, which as you know, takes serious steps to hold websites and other institutions accountable that knowingly facilitate sex trafficking activity by closing loopholes in what was outdated federal communications law. As an active participant in the deliberations and negotiations throughout the process, I noticed that while Facebook ultimately supported the legislation, that was a stance that evolved significantly— can you explain Facebook's shifting views on this bill?

Facebook supports SESTA. We support the goal of the legislation of providing victims of sex trafficking with recourse in the courts against parties who directly support these illegal activities, but wanted to ensure that good actors were not penalized for their efforts to root out this type of harm online. We were very pleased to be able to work successfully with a bipartisan group of Senators on a bill that protects women and children from the harms of sex trafficking.

Facebook is committed to making our platform a safe place, especially for individuals who may be vulnerable. We have a long history of working successfully with governments to address a wide variety of threats to our platform, including child exploitation. When we learn of a situation involving physical abuse, child exploitation, or an imminent threat of harm to a person, we immediately report the situation to first responders or the National Center for Missing and Exploited Children (NCMEC).

Further, as part of official investigations, government officials sometimes request data about people who use Facebook. We have processes in place to handle these government requests, and we disclose account records in accordance with our terms of service and applicable law. We also have a global team that strives to respond within minutes to emergency requests from law enforcement.

Our relationship with NCMEC also extends to an effort that we launched in 2015 to send AMBER Alerts to the Facebook community to help find missing children. When police determine that a case qualifies for an AMBER Alert, the alert is issued by the NCMEC and distributed through the Facebook system with any available information, including a photograph of the missing child, a license plate number, and the names and descriptions of the child and suspected abductor. Law enforcement determines the range of the target area for each alert. We know the chances of finding a missing child increase when more people are on the lookout, especially in the critical first hours. Our goal is to help get these alerts out quickly to the people who are in the best position to help, and a number of missing children have been found through AMBER Alerts on Facebook.

Further, we work tirelessly to identify and report child exploitation images (CEI) to appropriate authorities. We identify CEI through a combination of automated and manual review. On the automated review side, we use image hashing to identify known CEI. On the manual review side, we provide in-depth training to content reviewers on how to identify possible CEI. Confirmed CEI is reported to the NCMEC, which then forwards this information to appropriate authorities. When we report content to the NCMEC, we preserve account information in accordance with applicable law, which can help further law enforcement investigations. We also reach out to law enforcement authorities in serious cases to ensure that our reports are received and acted upon.

Question 5. Were your terms of service for third party app developers violated by Cambridge Analytica? If not, have they ever been violated in the past and what were those situations and outcomes?

Cambridge Analytica signed certifications at our insistence declaring that they had deleted all copies of Facebook data and derivatives obtained from Kogan's app. In March 2018, we received reports that, contrary to the certification and confirmation we were given by SCL/Cambridge Analytica, not all data was deleted. We are moving aggressively to determine the accuracy of these claims. If true, this is an unacceptable violation of trust and a breach of the representations Cambridge Analytica made in the certifications.

Question 6. Can a user opt-out of Facebook collecting and compiling a user's web browsing history? If so, please provide the details regarding how a user opts out of this collection.

The Ad Preferences tool on Facebook shows people the advertisers whose ads the user might be seeing because they visited the advertisers' sites or apps. The person can remove any of these advertisers to stop seeing their ads.

In addition, the person can opt out of these types of ads entirely—so he or she never sees those ads on Facebook based on information we have received from other websites and apps.

We've also announced plans to build Clear History, a feature that will enable people to see the websites and apps that send us information when they use them, delete this information from their accounts, and turn off our ability to store it associated with their accounts going forward.

Apps and websites that use features such as the Like button or Facebook Analytics send us information to make their content and ads better. We also use this information to make user experience on Facebook better. If a user clears his or her history or uses the new setting, we'll remove identifying information so a history of the websites and apps the user used won't be associated with the user's account. We'll still provide apps and websites with aggregated analytics—for example, we can build reports when we're sent this information so we can tell developers if their apps are more popular with men or women in a certain age group. We can do this without storing the information in a way that's associated with the user's account, and as always, we don't tell advertisers who users are.

It will take a few months to build Clear History. We'll work 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 heard specific demands for controls like these at a session we held at our headquarters. We're looking forward to doing more.

Question 7. Finally, since you've recently spent some time in Alaska, I'm sure your travels gave you a sense for our ardent individualism and general skepticism about the benefits of conceding privacy in the name of security. How can my constituents be assured of their security online? Or more generally, what would you say should be their new expectation of privacy online?

We believe that everyone has the right to expect strong protections for their information, and that we also need to do our part to help keep our community safe, in a way that's consistent with people's privacy expectations. We've recently announced several steps to give people more control over their privacy, including a new Privacy Shortcuts tool that we're rolling out now to give people information about how to control their information, including choosing who can see what they post and adding protections like two-factor authentication to their account. People can learn more about how to protect their privacy in our updated Data Policy and in our Privacy Basics feature (<https://www.facebook.com/about/basics>).

Questions from Senator Cruz

Question 1. Please provide a wholly contained answer to each question. A question's answer should not cross-reference answers provided in other questions.

If a question asks for a yes or no answer, please provide a yes or no answer first and then provide subsequent explanation. If the answer to a yes or no question is sometimes yes and sometimes no, please state such first and then describe the circumstances giving rise to each answer.

If a question asks for a choice between two options, please begin by stating which option applies, or both, or neither, followed by any subsequent explanation.

If you disagree with the premise of a question, please answer the question as-written and then articulate both the premise about which you disagree and the basis for that disagreement.

If you lack a basis for knowing the answer to a question, please first describe what efforts you undertook as Chief Executive Officer of Facebook order to ascertain an answer to the question and then provide your tentative answer as a consequence of its reasonable investigation. If even a tentative answer is impossible at this time, please state what efforts you and Facebook intend to take to provide an answer in the future and give an estimate as to when the Committees shall receive that answer.

If it is impossible to answer a question without divulging confidential or privileged information, please clearly state the basis for confidentiality or privilege invoked and provide as extensive an answer as possible without breaching that confidentiality or privilege. For questions calling for answers requiring confidential information, please provide a complete answer in a sealed, confidential form. These materials will be kept confidential. For questions calling for privileged information, please describe the privileged relationship and identify the privileged documents or materials that, if disclosed, would fully answer the question.

If the answer to a question depends on one or more individuals' memory or beliefs and that individual or those individuals either do not recall relevant information or are not available to provide it, please state the names of those individuals, what efforts you undertook to obtain the unavailable information, and the names of other individuals who may have access to that information.

To the extent that an answer depends on an ambiguity in the question asked, please state the ambiguity you perceive in the question and provide multiple answers which articulate each possible reasonable interpretation of the question in the light of the ambiguity.

To the extent that a question inquires about you or Facebook's actions, omissions, or policies, the question also asks about any entities that you or Facebook owns or controls, including any subsidiaries and affiliates. If context suggests that a question may ask about Facebook as a service rather than as an entity, please answer the question as applied to both Facebook as a service as well as all of Facebook's affiliated entities or platforms.

Please attach a copy of each and every formal or informal policy, whether presently written or otherwise, regarding the moderation, promotion, evaluation, or alteration of users or content on Facebook. These include, for example, Facebook’s Terms of Service, its Community Guidelines, and similar policies.

Facebook’s Terms and Policies are available here: <https://www.facebook.com/policies>.
Facebook’s Community Standards are available at <https://www.facebook.com/communitystandards/>.

Question 2. Yes or no: Are Facebook’s decisions to permit users access to its services or to permit content to remain displayed on its services, or the prominence or accessibility of that content, including its order, visibility, duration visible, inclusion in searches or order within search results, inclusion within “Trending” lists or analogous suggestions of content to users, determined in whole or part by Facebook’s corporate values, beliefs, priorities, or opinions?

The conversations that happen on Facebook reflect the diversity and free expression 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.

With regard to the order and visibility of content, a user’s News Feed is made up of stories from their friends, Pages they’ve chosen to follow and groups they’ve joined. *Ranking* is the process we use to organize all of those stories so that users can see the most relevant content at the top, every time they open Facebook. Ranking has four elements: the available *inventory* of stories; the *signals*, or data points that can inform ranking decisions; the *predictions* we make, including how likely we think they are to comment on a story, share with a friend, etc.; and a *relevancy score* for each story.

Misleading or harmful content on Facebook comes in many different forms, from annoyances like clickbait to hate speech and violent content. When we detect this kind of content in News Feed, there are three types of actions we take: remove it, reduce its spread, or inform people with additional context.

Our Community Standards and Ads Policies outline the content that is not allowed on the platform, such as hate speech, fake accounts, and praise, support, or representation of terrorism/terrorists. When we find things that violate these standards, we remove them. There are other types of problematic content that, although they don’t violate our policies, are still misleading or harmful and that our community has told us they don’t want to see on Facebook—things like clickbait or sensationalism. When we find examples of this kind of content, we reduce its spread in News Feed using ranking and, increasingly, we inform users with additional context so they can decide whether to read, trust, or share it.

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. Our policies are also rooted in the following principles:

- (1) Safety: People need to feel safe in order to build community. We are committed to removing content that encourages real-world harm, including (but not limited to)

physical, financial, and emotional injury.

- (2) Voice: Our mission is all about embracing diverse views. We err on the side of allowing content, even when some find it objectionable, unless removing that content can prevent a specific harm. Moreover, 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; and
- (3) Equity: Our community is global and diverse. Our policies may seem broad, but that is because we apply them consistently and fairly to a community that transcends regions, cultures, and languages. As a result, our Community Standards can sometimes appear less nuanced than we would like, leading to an outcome that is at odds with their underlying purpose. For that reason, in some cases, and when we are provided with additional context, we make a decision based on the spirit, rather than the letter, of the policy.

Question 3. Yes or no: Does Facebook promote, demote, or block users or content based on its assessment of the social value or social desirability of that content?

See Response to Question 2.

Question 4. Yes or no: Does Facebook promote, demote, or block users or content based on its assessment of that content's truth or falsity?

See Response to Question 2.

Question 5. Yes or no: Does Facebook promote, demote, or block users or content based on its assessment of the content's agreement or disagreement with Facebook's corporate values, beliefs, priorities, or opinions?

See Response to Question 2.

Question 6. Yes or no: Have Facebook's decisions to permit users access to its services or to permit content to remain displayed on its services, or the prominence or accessibility of that content, including its order, visibility, duration visible, inclusion in searches or order within search results, inclusion within "Trending" lists or analogous suggestions of content to users, ever been determined in whole or part by Facebook's corporate values, beliefs, priorities, or opinions?

See Response to Question 2.

Question 7. Yes or no: Has Facebook ever promoted, demoted, or blocked users or content based on its assessment of the social value or social desirability of that content?

See Response to Question 2.

Question 8. Yes or no: Has Facebook ever promoted, demoted, or blocked users or content based on its assessment of that content’s truth or falsity?

See Response to Question 2.

Question 9. Yes or no: Has Facebook ever promoted, demoted, or blocked users or content based on its assessment of the content’s agreement or disagreement with Facebook’s corporate values, beliefs, priorities, or opinions?

See Response to Question 2.

Question 10. Yes or no: Does Facebook employ its corporate values, beliefs, priorities, or opinions when deciding what content Facebook removes, republishes, moderates, promotes, or otherwise increases or decreases access to content?

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 around the world to all types of content. 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.

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. Our policies are also rooted in the following principles:

- (1) Safety: People need to feel safe in order to build community. We are committed to removing content that encourages real-world harm, including (but not limited to) physical, financial, and emotional injury.
- (2) Voice: Our mission is all about embracing diverse views. We err on the side of allowing content, even when some find it objectionable, unless removing that content can prevent a specific harm. Moreover, 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; and
- (3) Equity: Our community is global and diverse. Our policies may seem broad, but that is because we apply them consistently and fairly to a community that transcends regions, cultures, and languages. As a result, our Community Standards can sometimes appear less nuanced than we would like, leading to an outcome that is at odds with their underlying purpose. For that reason, in some cases, and when we are provided with additional context, we make a decision based on the spirit, rather than the letter, of the policy.

Question 11. Yes or no: Has Facebook ever employed its corporate values, beliefs, priorities, or opinions when deciding what content Facebook removes, republishes, moderates, promotes, or otherwise increases or decreases access to content?

See Response to Question 10.

Question 12. It has become a common position on colleges and universities that statements which a listener disagrees with severely either can constitute violence or can rise to the moral equivalent of violence. According to this position, statements may rise to the level of violence even without a threat, reasonable or otherwise, of imminent violence, the use of “fighting words,” or either a subjective intent or reasonably understood objective attempt to harass a listener.

See Response to Question 13.

Question 13. Yes or no: Does Facebook believe that speech neither advocating for physical violence against, threatening physical violence against, nor undertaken with either the subjective purpose or objective indicia of harassing a listener, may constitute violence?

Freedom of expression is one of our core values, and we believe that adding voices to the conversation creates a richer and more vibrant community. We want people to feel confident that our community welcomes all viewpoints and we are committed to designing our products to give all people a voice and foster the free flow of ideas and culture.

On the subject of credible violence, our Community Standards are explicit in what we don’t allow. We aim to prevent potential real-world harm that may be related to content on Facebook. We understand that people commonly express disdain or disagreement by threatening or calling for violence in facetious and non-serious ways. That’s why we try to consider the language, context and details in order to distinguish casual statements from content that constitutes a credible threat to public or personal safety. In determining whether a threat is credible, we may also consider additional information like a targeted person’s public visibility and vulnerability. We remove content, disable accounts, and work with law enforcement when we believe there is a genuine risk of physical harm or direct threats to public safety.

Question 14. Yes or no: Has Facebook ever believed that speech neither advocating for physical violence against, threatening physical violence against, nor undertaken with either the subjective purpose or objective indicia of harassing a listener, may constitute violence?

See Response to Question 13.

Question 15. Regardless of Facebook’s answer to Question 7, have any of Facebook’s policies ever required removal of content not described in Question 7 from Facebook? If so, what categories, and based on what policies?

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. Our policies are also rooted in the following principles:

- (1) Safety: People need to feel safe in order to build community. We are committed to removing content that encourages real-world harm, including (but not limited to) physical, financial, and emotional injury.
- (2) Voice: Our mission is all about embracing diverse views. We err on the side of allowing content, even when some find it objectionable, unless removing that content can prevent a specific harm. Moreover, 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; and
- (3) Equity: Our community is global and diverse. Our policies may seem broad, but that is because we apply them consistently and fairly to a community that transcends regions, cultures, and languages. As a result, our Community Standards can sometimes appear less nuanced than we would like, leading to an outcome that is at odds with their underlying purpose. For that reason, in some cases, and when we are provided with additional context, we make a decision based on the spirit, rather than the letter, of the policy.

Question 16. Yes or no: Does Facebook consider itself a publisher or speaker entitled to First Amendment protection when supervising its services, designing or implementing its policies, altering, reposting, promoting or demoting content, including through results displayed by a user search, their order or presence in a “Trending” list or similar suggestions to users regarding content?

Facebook does not create the content that users share on its Platform, although it does take steps to arrange, rank and distribute that content to those who are most likely to be interested in it, or to remove objectionable content from its service. These activities are protected functions under Communications Decency Act Section 230 and the First Amendment.

Question 17. Aside from content clearly marked as coming from Facebook or one of its officers or employees, under what circumstances does Facebook consider itself as acting as a First-Amendment-protected publisher or speaker in its moderation, maintenance, or supervision over its users or their content?

We are, first and foremost, a technology company. Facebook does not create or edit the content that users publish on our platform. While we seek to be a platform for a broad range of ideas, we do moderate content according to published community standards in order to keep users on the platform safe, to reduce objectionable content and to make sure users participate on the platform responsibly.

Question 18. Yes or no: Does Facebook provide access to its services on a viewpoint-neutral basis? For this question and its subparts, please construe “access to its services” and similar phrases broadly, including the position or order in which content is displayed on its services, the position or order in which users or content show up in searches (or

whether they show up at all), whether users or content are permitted to purchase advertisements (or be advertised), the rates charged for those advertisements, and so on.

We are committed to free expression and err on the side of allowing content. When we make a mistake, we work to make it right. And we are committed to constantly improving our efforts so we make as few mistakes as humanly possible.

Decisions about whether to remove content are based on whether the content violates our Community Standards.

Discussing controversial topics or espousing a debated point of view is not at odds with our Community Standards, the policies that outline what is and isn't allowed on Facebook. We believe that such discussion is important in helping bridge division and promote greater understanding.

We are committed to designing our products to give all people a voice and foster the free flow of ideas and culture. That said, when something crosses the line into hate speech, it has no place on Facebook, and we are committed to removing it from our platform any time we become aware of it.

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, sex, gender, gender identity, and serious disability or disease. We also provide some protections for immigration status. We define attack as violent or dehumanizing speech, statements of inferiority, and calls for exclusion or segregation. The detailed guidelines our reviewers use to assess whether content violates our hate speech policies are available here:

https://www.facebook.com/communitystandards/objectionable_content/hate_speech.

Question 19. Yes or no: Has Facebook ever discriminated among users on the basis of viewpoint when determining whether to permit a user to access its services? If so, please list each instance in which Facebook has done so.

See Response to Question 18.

Question 20. If so, does Facebook continue to do so today, or when did Facebook stop doing so?

See Response to Question 18.

Question 21. If so, what viewpoint(s) has Facebook discriminated against or in favor of? In what way(s) has Facebook done so?

See Response to Question 18.

Question 22. If so, does Facebook act only on viewpoints expressed on Facebook, or does it discriminate among users based on viewpoints expressed elsewhere? Has Facebook ever based its decision to permit or deny a user access to its services on viewpoints expressed off Facebook?

See Response to Question 18.

Question 23. Yes or no: Excluding content encouraging physical self-harm, threats of physical violence, terrorism, and other content relating to the credible and imminent physical harm of specific individuals, has Facebook ever discriminated among content on the basis of viewpoint in its services? If so, please list each instance in which Facebook has done so.

See Response to Question 18.

Question 24. Yes or no: Has Facebook ever discriminated against American users or content on the basis of an affiliation with a religion or political party? If so, please list each instance in which Facebook has done so and describe the group or affiliation against which (or in favor of which) Facebook was discriminating.

See Response to Question 18.

Question 25. Yes or no: Has Facebook ever discriminated against any American users or content on its services on the basis of partisan affiliation with the Republican or Democratic parties? This question includes advocacy for or against a party or specific candidate or official. If so, please list each instance and the party affiliation discriminated against.

See Response to Question 18.

Question 26. Yes or no: Has Facebook ever discriminated against any American users or content on its services on the basis of the user's or content's advocacy for a political position on any issue in local, State, or national politics? This question includes but is not limited to advocacy for or against abortion, gun control, consumption of marijuana, and net neutrality.

See Response to Question 18.

Question 27. Yes or no: Has Facebook ever discriminated against any American users or content on its services on the basis of the user's or content's religion, including advocacy for one or more tenets of that religion? If so, please list each such instance in which Facebook has done so and identify the religion, religious group, or tenet against which Facebook discriminated.

See Response to Question 18.

Question 28. Yes or no: Has Facebook ever discriminated between users in how their content is published, viewed, received, displayed in "trending" or similar lists, or otherwise in any function or feature, based on the user's political affinity, religion, religious tenets, ideological positions, or any ideological or philosophical position asserted? If so, please list each such incident as well as the basis on which Facebook discriminated against that user or content.

Being a platform for all ideas is a foundational principle of Facebook. We are committed to ensuring there is no bias in the work we do.

Suppressing content on the basis of political viewpoint or preventing people from seeing what matters most to them is directly contrary to Facebook's mission and our business objectives.

When allegations of political bias surfaced in relation to Facebook's Trending Topics feature, we immediately launched an investigation to determine if anyone violated the integrity of the feature or acted in ways that are inconsistent with Facebook's policies and mission. We spoke with current reviewers and their supervisors, as well as a cross-section of former reviewers; spoke with our contractor; reviewed our guidelines, training, and practices; examined the effectiveness of operational oversight designed to identify and correct mistakes and abuse; and analyzed data on the implementation of our guidelines by reviewers.

Ultimately, our investigation revealed no evidence of systematic political bias in the selection or prominence of stories included in the Trending Topics feature. In fact, our analysis indicated that the rates of approval of conservative and liberal topics are virtually identical in Trending Topics. Moreover, we were unable to substantiate any of the specific allegations of politically-motivated suppression of subjects or sources, as reported in the media. To the contrary, we confirmed that most of those subjects were in fact included as trending topics on multiple occasions, on dates and at intervals that would be expected given the volume of discussion around those topics on those dates.

Nonetheless, as part of our commitment to continually improve our products and to minimize risks where human judgment is involved, we are making a number of changes:

We have engaged an outside advisor, former Senator Jon Kyl, to advise the company on potential bias against conservative voices. We believe this external feedback will help us improve over time and ensure we can most effectively serve our diverse community and build trust in Facebook as a platform for all ideas.

We continue to expand our list of outside partner organizations to ensure we receive feedback on our content policies from a diverse set of viewpoints.

We have made our detailed reviewer guidelines public to help people understand how and why we make decisions about the content that is and is not allowed on Facebook.

We have launched an appeals process to enable people to contest content decisions with which they disagree.

We are instituting additional controls and oversight around the review team, including robust escalation procedures and updated reviewer training materials.

These improvements and safeguards are designed to ensure that Facebook remains a platform for all ideas and enables the broadest spectrum of free expression possible.

Question 29. Except for accidental instances, has Facebook ever removed, downgraded,

concealed, or otherwise censored content associated with any of the following? If yes, please describe the content that was removed, downgraded, concealed, or otherwise censored and the circumstances under which it was removed, downgraded, concealed, or otherwise censored.

Any individuals employed by Facebook?

Any elected official or candidate seeking elected office who self-identifies or is registered as a Democrat or a “Democratic Socialist”?

Any group who self-identifies as being part of the “Anti-Trump Resistance Movement”?

Any individuals employed by MSNBC?

Any individuals employed by CNN?

Any blogs that self-identify as “liberal” or “progressive”?

Any Facebook groups that self-identify as “liberal”, “progressive”, or being part of the “Anti-Trump Resistance Movement”?

Open Society Foundation?

Planned Parenthood?

Indivisible?

Sierra Club?

The American Civil Liberties Union?

The Anti-Defamation League?

The Council on American-Islamic Relations (CAIR)?

Emily’s List?

NARAL Pro-Choice America?

The National Association for the Advancement of Colored People (NAACP)?

NextGen Climate Action?

The Southern Poverty Law Center?

The Union of Concerned Scientists?

Everytown for Gun Safety?

Amnesty International?

Priorities USA Action?

Media Matters for America?

Human Rights Watch?

Every Voice?

NowThis?

The Women's March?

Organizing for America?

Organizing for Action?

When content that violates our policies is brought to our attention, we remove that content—regardless of who posted it. We have removed content posted by individuals and entities across the political spectrum.

On April 24, 2018, we published the detailed guidelines our reviewers use to make decisions about reported content on Facebook. These guidelines cover everything from nudity to graphic violence.

We published these guidelines because we believe that increased transparency will provide more clarity on where we draw lines on complex and continuously evolving issues, and we hope that sharing these details will prompt an open and honest dialogue about our decision making process that will help us improve—both in how we develop and enforce our standards. We recognize that our policies are only as good as the strength and accuracy of our enforcement—and our enforcement is not perfect. We make mistakes because our processes involve people, and people are not infallible. We are always working to improve.

We typically do not comment on specific cases of content removal for privacy reasons.

Question 30. In your testimony before the committees, you stated several times that Facebook prohibits content based on its status as “hate speech.” How have you and Facebook defined “hate speech” today and at any other stage in Facebook’s existence?

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, sex, gender, gender identity, and serious disability or disease. We also provide some protections for immigration status. We define attack as violent or dehumanizing speech, statements of inferiority, and calls for exclusion or segregation. The detailed guidelines our reviewers use to assess whether content violates our hate speech policies are available here:

https://www.facebook.com/communitystandards/objectionable_content/hate_speech.

Our Community Standards make an important distinction between targeting people and targeting particular beliefs or institutions. We believe that people should be able to share their views and discuss controversial ideas on Facebook.

Question 31. Did or does Facebook collaborate with or defer to any outside individuals or organizations in determining whether to classify a particular statement as “hate speech?” If so, please list the individuals and organizations.

Hate speech has no place on our platform. Our Community Standards prohibit attacks based on characteristics including race, ethnicity, religion, and national origin.

Facebook has partnerships with academics and experts who study organized hate groups and hate speech. These academics and experts share information with Facebook as to how organizations are adapting to social media and give feedback on how Facebook might better tackle these problems. We recently hosted several of these academics at Facebook for multiple days of observation and assessment, during which the academics attended substantive meetings on our content policies and the guidance we provide to our reviewers. Further, in the area of hate speech, there are very important academic projects that we follow closely. Timothy Garton Ash, for example, has created the Free Speech Debate to look at these issues on a cross-cultural basis. Susan Benesch established the Dangerous Speech Project, which investigates the connection between speech and violence. These projects show how much work is left to be done in defining the boundaries of speech online, which is why we will keep participating in this work to help inform our policies at Facebook. We are committed to continuing our dialogue with third parties to ensure we can have the widest possible expression of ideas, while preventing abuse of the platform.

Facebook works with organizations from across the political spectrum around changes to our content standards including hate speech. While we do not share individual pieces of content from users with these organizations out of concerns for user privacy, we do provide in-depth examples and explanations of what the policy changes would entail.

Question 32. Did or does Facebook collaborate with or defer to any outside individuals or organizations in determining whether a given speaker has committed acts of “hate speech” in the past? If so, please list the individuals and organizations.

In an effort to prevent and disrupt real-world harm, we do not allow any organizations or individuals that are engaged in organized hate to have a presence on Facebook. We also remove content that expresses support or praise for groups, leaders, or individuals involved in these activities.

In developing and iterating on our policies, including our policy specific to hate speech, we consult with outside academics and experts from across the political spectrum and around the world. We do not, however, defer to these individuals or organizations in making decisions about content on our platform. Content that violates our Community Standards is removed when we are made aware of it, and content that doesn’t violate is left on the platform.

Designating hate organizations and/or individuals is an extensive process that takes into account a number of different signals. We worked with academics and NGOs to establish this

process and regularly engage with them to understand whether we should refine it. Among the signals we consider are whether the individual or organization in question has called for or directly carried out violence against people based on protected characteristics.

Question 33. Did or does Facebook ban or otherwise limit the content of individuals or organizations who have spoken “hate speech” on its platform aside from the offending content? If so, under what circumstances?

See Response to Question 32.

Question 34. Yes or no: Did or does Facebook ban or otherwise limit the content of individuals or organizations on its platform based on hate speech or other behavior conducted outside of Facebook’s platform?

See Response to Question 32.

Question 35. Yes or no: Do you believe that “hate speech” is not protected under the First Amendment from government censorship?

The goal of our Community Standards is to encourage expression and create a safe community for our 2 billion users, more than 87% of whom are located outside the United States.

We err on the side of allowing content, even when some find it objectionable, unless removing that content prevents a specific harm.

We do not allow hate speech on Facebook because it creates an environment of intimidation and exclusion and in some cases may promote real-world violence.

Our current definition of hate speech is anything that directly attacks people based on what are known as their “protected characteristics”—race, ethnicity, national origin, religious affiliation, sexual orientation, sex, gender, gender identity, or serious disability or disease. However, our definition does allow for discussion around these characteristics as concepts in an effort to allow for and encourage expression and dialogue by our users.

There is no universally accepted answer for when something crosses the line.

Our approach to hate speech, like those of other platforms, has evolved over time and continues to change as we learn from our community, from experts in the field, and as technology provides us new tools to operate more quickly, more accurately and precisely at scale.

Question 36. Yes or no: Have you ever believed that “hate speech” is not protected under the First Amendment from government censorship?

See Response to Question 35.

Question 37. Yes or no: Does Facebook believe that “hate speech” is not protected under the First Amendment from government censorship?

See Response to Question 35.

Question 38. Yes or no: Has Facebook ever believed that “hate speech” is not protected under the First Amendment from government censorship?

See Response to Question 35.

Question 39. Yes or no: Does Facebook’s “hate speech” policy prohibit, exclude, remove, or censor content that, were Facebook a governmental entity, would be entitled to First Amendment protections?

See Response to Question 35.

Question 40. Facebook states on its website that, per its community standards, Facebook will remove hate speech, which it describes as “including content that directly attacks people based on their: race, ethnicity, national origin, religious affiliation, sexual orientation, sex, gender, or gender identity, or serious disabilities or diseases.” Yes or no: Does Facebook limit its definition of hate speech only to content that “directly attacks” people based on the aforementioned characteristics?

We define “attack” under our hate speech policy as violent or dehumanizing speech, statements of inferiority, and calls for exclusion or segregation. We allow discussion of issues related to characteristics like race, gender, ethnicity, and immigration status. We do not permit attacks against people based on these characteristics. Context matters in making what can be a difficult determination in some cases.

Specific details on the type of content that is prohibited under our hate speech policies are available here:

https://www.facebook.com/communitystandards/objectionable_content/hate_speech.

Question 41. What standard or procedure has Facebook applied now and in the past in determining whether content “directly attacks” an individual or group based on a protected characteristic under Facebook’s community standards?

See Response to Question 40.

Question 42. Yes or no: Has Facebook ever removed content for hate speech that did not directly attack a person on the basis of his or her race, ethnicity, national origin, religious affiliation, sexual orientation, sex, gender, or gender identity, or serious disabilities or diseases? If so, what criteria did Facebook use to determine that the content violated Facebook’s policy?

We define “attack” under our hate speech policy as violent or dehumanizing speech, statements of inferiority, and calls for exclusion or segregation.

Sometimes, it’s obvious that something is hate speech and should be removed—because it includes the direct incitement of violence against people possessing protected characteristics, or degrades or dehumanizes people. Sometimes, however, there isn’t a clear consensus—because

the words themselves are ambiguous, the intent behind them is unknown, or the context around them is unclear. Language also continues to evolve, and a word that was not a slur yesterday may become one today.

Here are some of the things we take into consideration when deciding what to leave on the site and what to remove.

- **Context:** Regional and linguistic context is often critical in deciding whether content constitutes hate speech, as is the need to take geopolitical events into account. In Myanmar, for example, the word “kalar” has benign historic roots, and is still used innocuously across many related Burmese words. The term can however also be used as an inflammatory slur, including as an attack by Buddhist nationalists against Muslims. We looked at the way the word’s use was evolving, and decided our policy should be to remove it as hate speech when used to attack a person or group, but not in the other harmless use cases.
- **Intent:** There are times someone might share something that would otherwise be considered hate speech but for non-hateful reasons, such as making a self-deprecating joke or quoting lyrics from a song. People often use satire and comedy to make a point about hate speech. In other cases, people may speak out against hatred by condemning someone else’s use of offensive language, which requires repeating the original offense. This is something we allow, even though it might seem questionable since it means some people may encounter material disturbing to them. But it also gives our community the chance to speak out against hateful ideas. We revised our Community Standards to encourage people to make it clear when they’re sharing something to condemn it, but sometimes their intent isn’t clear, and anti-hatred posts get removed in error.

On April 24, 2018, we announced the launch of appeals for content that was removed for hate speech. We recognize that we make enforcement errors on both sides of the equation—what to allow, and what to remove—and that our mistakes cause a great deal of concern for people, which is why we need to allow the option to request review of the decision and provide additional context that will help our team see the fuller picture as they review the post again. This type of feedback will allow us to continue improving our systems and processes so we can prevent similar mistakes in the future.

Question 43. Has Facebook ever removed content for hate speech that was posted by an individual employed by Facebook? If so, please describe each instance.

Our policies apply equally to all of our users. If a Facebook employee posted content that was reported to us and violated our policies, the content would be removed.

Question 44. Recording artist Taylor Swift recently released a cover of Earth, Wind & Fire’s “September.”

In response, Nathaniel Friedman, an author at GQ magazine, stated that “Taylor Swift’s cover of ‘September’ is hate speech.” Does Facebook agree?

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, sex, gender, gender identity, and serious disability or disease. We also provide some protections for immigration status. We define attack as violent or dehumanizing speech, statements of inferiority, and calls for exclusion or segregation. Our detailed hate speech policies are available at https://www.facebook.com/communitystandards/objectionable_content/hate_speech.

We generally do not assess whether content violates our policies (including our hate speech policy) unless it is part of our normal content review process. Context matters in making what can be a difficult determination in some cases. Sometimes, it's obvious that something is hate speech and should be removed—because it includes the direct incitement of violence against people possessing protected characteristics, or degrades or dehumanizes people. Sometimes, however, there isn't a clear consensus—because the words themselves are ambiguous, the intent behind them is unknown or the context around them is unclear. Language also continues to evolve, and a word that was not a slur yesterday may become one today.

Question 45. In response, Monique Judge, an author at The Root, stated that “Taylor Swift needs her * whooped.” Is this statement hate speech?**

See Response to Question 44.

Question 46. It was reported that Democratic D.C. Councilman Trayon White posted a video on his Facebook page blaming a recent snowstorm on wealthy Jewish families. According to USA Today, White said: “It just started snowing out of nowhere this morning, man. Y’all better pay attention to this climate control, man, this climate manipulation,” which White attributed to “the Rothschilds controlling the climate to create natural disasters they can pay for to own the cities, man.”

Yes or no: Does Facebook consider this video or this quote hate speech?

See Response to Question 44.

Yes or no: Did Facebook remove this video from its platform? If so, when? If not, why not?

See Response to Question 44.

Question 47. Multiple authors for the website Vox, including its founder, Ezra Klein, have described Charles Murray’s book, *The Bell Curve*, as “hate speech.” Similarly, the left-wing Southern Poverty Law Center perplexingly describes Murray as a “white nationalist,” largely relying on its depiction of *The Bell Curve*.

Does *The Bell Curve* qualify as “hate speech” for purposes of Facebook’s policies?

See Response to Question 44.

If so, what portions of *The Bell Curve* qualify as “hate speech?” Please provide quotations with page numbers for these portions.

See Response to Question 44.

If not, do Facebook’s content policies prohibit a false claim that someone has engaged in “hate speech?”

See Response to Question 44.

Question 48. What procedures or penalties does Facebook employ, if any, to discourage false claims that someone has engaged in hate speech?

See Response to Question 44.

Question 49. Are any portions of the Bible, quoted verbatim and with citation, subject to removal as:

“Hate speech?” If so, please list the quotations and under which translation Facebook considers the quote “hate speech.”

See Response to Question 44.

Harassment? If so, please list the quotations and under which translation Facebook considers the quote harassment.

We do not tolerate harassment on Facebook because we want people to feel safe to engage and connect with their community. Our harassment policy applies to both public and private individuals and includes behavior like repeatedly contacting a single user despite that person’s clear desire and action to prevent that contact and repeatedly contacting large numbers of people with no prior solicitation. It also applies to calls for death, serious disease or disability, or physical harm aimed at an individual or group of individuals in a message thread. Context and intent matter, however, and we allow people to share and re-share posts if it is clear that something was shared in order to condemn or draw attention to harassment. The detailed guidelines our reviewers use to assess whether content violates our hate speech policies are available at <https://www.facebook.com/communitystandards/safety/harassment>.

We released our updated Community Standards—which reflect the guidelines our reviewers use to evaluate content that is reported to us—in order to better demonstrate where we draw lines on complex and continuously evolving issues. We also simultaneously launched an appeals process for content that has been removed for nudity/sexual activity, hate speech, and graphic violence. With this launch, we are giving people an opportunity to request review of our decisions and provide additional context that will help our team see a more complete picture as they review the post again. This type of feedback allows us to continue improving our systems and processes so we can prevent similar mistakes in the future.

Question 50. On April 19, 2018, the California State Assembly voted in favor of a bill, AB 2943, which would make it an “unlawful business practice” to engage in any transaction for a good or service that seeks “to change an individual’s sexual orientation” The bill clarifies that this includes efforts to “change behaviors or gender expressions, or to eliminate or reduce sexual or romantic attractions or feelings toward individuals of the same sex.”

Multiple legal experts have observed that the bill’s language, reasonably interpreted, could be read to outlaw the sale and purchase of books, such as the Bible, the Torah, and the Koran, which advocate for traditional sexual ethics.

Yes or no: Does Facebook believe that books, such as the Bible, the Torah, and the Koran, which advocate for traditional sexual ethics, constitute hate speech?

See Response to Question 44.

Yes or no: Does Facebook consider any part of the Bible, the Torah, and/or the Koran hate speech? If so, what parts of the Bible, the Torah, and/or the Koran qualify? Please provide quotations with page numbers for each part identified as hate speech.

See Response to Question 44.

Yes or no: Does Facebook believe that the messages contained in books, such as the Bible, the Torah, and the Koran, which advocate for traditional sexual ethics (i.e. that sex should be had only within a marriage between one man and one woman), should be discouraged from public dissemination?

See Response to Question 44.

Yes or no: Does Facebook agree with the California State Assembly that goods or services that seek to change behaviors or gender expressions deserve to be discouraged, muted, or banned?

See Response to Question 44.

Yes or no: Does Facebook agree with the California State Assembly that goods or services that seek to eliminate or reduce sexual or romantic attractions or feelings toward individuals of the same sex deserve to be discouraged, muted, or banned?

See Response to Question 44.

Yes or no: In the event AB 2943 is fully enacted into law, will Facebook comply with its provisions by removing, denying, downgrading, concealing, or otherwise censoring content and advertisements restricted by the bill? If so, does Facebook intend to remove, deny, downgrade, conceal, or otherwise censor content and advertisements that pertain to the Bible, the Torah, the Koran, and other books which advance traditional sexual ethics.

See Response to Question 44.

Question 51. If an individual posted any of the following statements, standing alone and not directed to any Facebook user in particular, would that statement violate Facebook’s “hate speech” policy? To the extent that the decision would depend on additional facts, please describe whether the statement would prompt an investigation to determine whether it constitutes “hate speech,” and whether the decision would involve algorithmic or human decision making.

There are only two sexes or two genders, male and female.

Bathroom segregation based on sex is similar to segregation based on race.

God created man in his image, male and female.

Gender is a social construct.

A person's sex or gender are immutable characteristics.

Sex reassignment surgery is a form of bodily mutilation.

The abortion of an unborn child is murder.

It should be a crime to perform or facilitate an abortion.

It should be a crime to prevent someone from performing or obtaining an abortion.

No person of faith should be required to assist a same-sex wedding by providing goods or services to a same-sex marrying couple.

When an individual enters the marketplace, he gives up the right to choose whether to support a same-sex marriage.

Islam is a religion of peace.

Islam is a religion of war.

All white people are inherently racist.

All black people are inherently racist.

Black lives matter.

Blue lives matter.

All lives matter.

Donating to the NRA funds the murder of children, such as those slain in Parkland, Florida.

Donating to Planned Parenthood funds the murder of children, such as those dismembered by Kermit Gosnell.

Men should stop interrupting when women are talking.

Women should stop interrupting when men are talking.

DREAMers are Americans too and should be entitled to stay in this country.

Illegal aliens need to be sent back.

Religious beliefs are irrational and anti-science.

Non-believers have no path to eternal salvation.

aa) Affirmative Action policies discriminate on the basis of race and sex. bb) America is a “melting pot.”

See Response to Question 44.

Question 52. Facebook states on its website that per its community standards, “organizations and people dedicated to promoting hatred” against protected groups are not allowed a presence on Facebook.

What standards or policies does Facebook apply in determining whether a group violates this policy?

See Response to Question 32.

Yes or no: Does Facebook contract with or in any way rely upon an outside party to determine what organizations and people are dedicated to promoting hatred against protected groups? If yes, please list the outside parties.

See Response to Question 32.

Yes or no: Has Facebook ever referenced, used, consulted, or in any way relied upon the left-wing Southern Poverty Law Center’s list of designated hate groups in order to determine whether an organization or individual was dedicated to promoting hatred against protected groups?

See Response to Question 32.

Yes or no: Has Facebook ever denied an organization a presence on Facebook on account of the organization being dedicated to promoting hatred? If so, has Facebook ever reversed its decision to designate an organization a hate group under its community standards and reinstated the organization’s privilege to post and have a presence on Facebook?

See Response to Question 32.

Question 53. One group on Facebook, “TERMINATE the Republican Party,” has over 10,000 followers, one of which was James T. Hodgkinson. In June 2017, Hodgkinson opened fire on Republican members of Congress at a baseball practice, seriously wounding Rep. Steve Scalise, a congressional staffer, and two heroic police officers. Quotes from this group’s posts and comments include that “These people are all the same, criminals, rapists, racists, Republicans;” that, about Rep. Patrick McHenry, “who gives birth to sorry pieces of s* like him and allowed it to reach adulthood, truly needs a f*****g hammer to the head a few times;” and, referring to the President, “G*****n Russian roach traitor bastard**

... and his Republicanazi followers!” Each of these quotes took place long after Hodgkinson’s shooting, though similar quotes are available from before it as well.

Do these quotes constitute “hate speech?”

If so, why have they not been removed?

If not, why do they not?

If applied to Democrats, would the quotes above constitute “hate speech?”

How has Facebook changed its platform in response to Hodgkinson’s shooting? It has apparently not suspended or ended this group.

Does it concern Facebook that such rhetoric is being used in a group which had an attempted political assassin as a member?

Does Facebook permit threats of violence against the President?

Does Facebook permit threats of violence against members of Congress?

Does Facebook monitor its platforms for potential left-wing violence?

If so, what is Facebook doing to ensure that shooters like Hodgkinson do not coordinate using Facebook?

If so, what is Facebook doing to ensure that shooters like Hodgkinson do not use Facebook to incite violence against Republicans or conservatives?

If not, why is Facebook not doing so given that its platform was integral to at least one attempted political assassination?

The shooting at the Congressional baseball practice was a horrendous act. As a designated mass shooting, any praise for that conduct or the shooter is against Facebook policies. We also do not allow any Pages or accounts representing the shooter. If we are made aware of such comments, we would take them down.

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, sex, gender, gender identity, and serious disability or disease. We also provide some protections for immigration status. Political-party affiliation is not included in our list of protected characteristics. We define attack as violent or dehumanizing speech, statements of inferiority, and calls for exclusion or segregation. Our detailed hate speech policies are available at https://www.facebook.com/communitystandards/objectionable_content/hate_speech.

Our credible violence policies prohibit posting credible statements of intent to commit violence against any person, groups of people, or place (city or smaller). We assess credibility

based upon the information available to us and generally consider statements credible if the following are present:

- A target (person, group of people, or place) and:
 - Bounty/demand for payment, or
 - Mention or image of specific weapon, or
 - Sales offer or ask to purchase weapon, or
 - Spelled-out address or named building, or
- A target and 2 or more of the following details (can be 2 of the same detail):
 - Location
 - Timing
 - Method

We also prohibit calls for violence, statements advocating violence, or aspirational or conditional statements of violence targeting public individuals, provided those statements are credible, as defined above. Any calls for violence against heads of state, including the United States President, violate our policies.

There are times someone might share something that would otherwise be considered hate speech but for non-hateful reasons, such as making a self-deprecating joke or quoting lyrics from a song. People often use satire and comedy to make a point about hate speech. In other cases, people may speak out against hatred by condemning someone else's use of offensive language, which requires repeating the original offense. This is something we allow, even though it might seem questionable since it means some people may encounter material disturbing to them.

Question 54. In July 2012, Governor Mike Huckabee praised Chick-fil-A because of its support for traditional marriage and called on Christians to support Chick-fil-A in its position by purchasing its products. Facebook temporarily removed Governor Huckabee's post from its service before reinstating it.

Why was Governor Huckabee's post removed?

What Facebook rule was Governor Huckabee's post thought to have violated before it was reinstated?

Did Governor Huckabee's post violate Facebook's prohibition on "hate speech," either in 2012 or now?

Does a post opposing the Supreme Court's decision in *Obergefell v. Hodges* violate Facebook's prohibition on "hate speech?"

Does a post opposing legalized same-sex marriage violate Facebook’s prohibition on “hate speech?”

As of July 2012, had Facebook removed, downgraded, concealed, or otherwise censored any content created by a state Governor, member of the U.S. House of Representatives, member of the U.S. Senate, or the President on account of that individual’s support for same-sex marriage? If so, please include the removed content including identifying information indicating its author.

As of July 2012, had Facebook removed, downgraded, concealed, or otherwise censored any other content created by a state Governor, member of the U.S. House of Representatives, member of the U.S. Senate, or the President on account of that individual’s opposition to same-sex marriage? If so, please include the removed content including identifying information indicating its author.

Has, since July 2012, Facebook removed, downgraded, concealed, or otherwise censored any posts by a state Governor, member of the U.S. House of Representatives, member of the U.S. Senate, or the President on account of that individual’s (or that content’s) opposition to same-sex marriage? If so, please include the removed post identifying information indicating its author.

Has, since July 2012, Facebook removed, downgraded, concealed, or otherwise censored any posts by a state Governor, member of the U.S. House of Representatives, member of the U.S. Senate, or the President on account of that individual’s (or that content’s) support for same-sex marriage? If so, please include the removed post identifying information indicating its author.

Under what circumstances does Facebook remove, downgrade, conceal, or otherwise censor content that, though not threatening physical harm, promoting imminent physical self-harm, or advocating for terrorism, opposes same-sex marriage?

Under what circumstances does Facebook remove, downgrade, conceal, or otherwise censor content that, though not threatening physical harm, promoting imminent physical self-harm, or advocating for terrorism, supports same-sex marriage?

In July 2012, our automated systems incorrectly removed an event page entitled “Chick-fil-A Appreciation Day.” The page was restored within hours of coming to our attention. When we make mistakes on these important content decisions, we make every attempt to make it right as quickly as we can.

Our goal is to allow people to have as much expression as possible, including on the issue of same-sex marriage. We err on the side of allowing content, even when some find it objectionable, unless removing that content prevents a specific harm.

See also Response to Question 44.

Question 55. As described in the Washington Post, in October 2012, Facebook removed a post by a group called “Special Operations Speaks.” The post said: “Obama called the

SEALS and THEY got bin Laden. When the SEALs called Obama, they got denied,” a reference to the failure of the Executive Branch to provide military support to Americans under assault, and later killed, in Benghazi. Facebook first warned the group that the post violated its rules and then subsequently removed the post as a violation of “Facebook’s Statements of Rights and Responsibilities.” Facebook further suspended Special Operations Speaks for 24 hours following the removal. Facebook later admitted error and permitted the content to remain on its platform.

Why was Special Operations Speaks’ post removed?

What term of Facebook’s then-extant 2012 Statement of Rights and Responsibilities was Special Operations Speaks’ post thought to have violated before Facebook reversed its decision?

Yes or no: Did any member of the Obama Administration, including any administrative agency then-directed by an executive official appointed by the Obama administration, contact Facebook to request that the post be removed?

If so, whom?

What was Facebook’s response?

Yes or no: Did Facebook assure any government official or employee that this post would be removed? If so, whom?

Did Special Operations Speaks’ post violate Facebook’s prohibition on “hate speech,” either in 2012 or now?

As of October 2012, had Facebook removed, downgraded, concealed, or otherwise censored any other content created by a political action committee on the basis of that content’s disapproval of how the Obama administration handled the attack on U.S. diplomats and servicemen in Benghazi? If so, please include the removed content including identifying information about its author.

As of October 2012, had Facebook removed, downgraded, concealed, or otherwise censored any content created by a political action committee on the basis of that content’s approval of how the Obama administration handled the attack on U.S. diplomats and servicemen in Benghazi? If so, please include the removed content including identifying information about its author.

Has, since October 2012, Facebook removed, downgraded, concealed, or otherwise censored any posts by a political action committee on the basis of that content’s disapproval of how the Obama administration handled the attack on U.S. diplomats and servicemen in Benghazi? If so, please include the removed content including identifying information about its author.

Has, since October 2012, Facebook removed, downgraded, concealed, or otherwise censored any posts by a political action committee on the basis of that content’s

disapproval of how the Obama administration handled the attack on U.S. diplomats and servicemen in Benghazi? If so, please include the removed content including identifying information about its author.

Under what circumstances does Facebook remove, downgrade, conceal, or otherwise censor content that, though not threatening physical harm, promoting imminent physical self-harm, or advocating for terrorism, opposes the Obama Administration's handling of the attacks on U.S. diplomats and servicemen in Benghazi?

Under what circumstances does Facebook remove, downgrade, conceal, or otherwise censor content that, though not threatening physical harm, promoting imminent physical self-harm, supports the Obama Administration's handling of the attacks on U.S. diplomats and servicemen in Benghazi?

In this particular case, we removed the content as a violation of our standards. The content was deleted for 29 hours. However, we realized that we made a mistake, and we restored the content and apologized for the error.

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, sex, gender, gender identity, and serious disability or disease. We also provide some protections for immigration status. We define attack as violent or dehumanizing speech, statements of inferiority, and calls for exclusion or segregation. Our detailed hate speech policies are available at https://www.facebook.com/communitystandards/objectionable_content/hate_speech.

Our Community Standards prohibit hate speech and celebrating graphic violence and allow people to use Facebook to raise awareness of and condemn violence. Drawing that line requires complex and nuanced judgments, and we carefully review reports that we receive from the public, media, civil society, and governments. We remove content that violates our policies, regardless of who posted the content.

Question 56. In September 2017, Facebook deemed the videos of two African American Trump supporters, known as Diamond and Silk, as “dangerous.” In a company email, Facebook stated that the decision was final and “not appealable in any way.” Facebook then retracted this statement, explaining that the determination was inaccurate.

What about Diamond and Silk did Facebook initially determine to be “dangerous?”

What is Facebook's criteria for determining whether content that neither depicts nor advocates for violence as “dangerous?”

Aside from the illustration of or advocacy for violence, under what conditions is the discussion of non-classified speech “dangerous?”

Has Facebook implemented an appeals system by which users can challenge a determination of dangerousness?

How often does Facebook retract these determinations?

What is the internal review process for these types of determinations?

We mishandled communication with Diamond and Silk for months. Their frustration was understandable, and we apologized to them. The message they received on April 5, 2018 that characterized their Page as “dangerous” was incorrect and not reflective of the way we seek to communicate with our community and the people who run Pages on our platform.

As part of our commitment to continually improve our products and to minimize risks where human judgment is involved, we are making a number of changes:

- We have engaged an outside advisor, former Senator Jon Kyl, to advise the company on potential bias against conservative voices. We believe this external feedback will help us improve over time and ensure we can most effectively serve our diverse community.
- We continue to expand our list of outside organizations from across the political spectrum to provide feedback on potential changes to our content standards.
- We have made our detailed reviewer guidelines public to help people understand how and why we make decisions about the content that is and is not allowed on Facebook.
- We have launched an appeals process to enable people to contest content decisions with which they disagree. We recognize that we make enforcement errors on both sides of the equation—what to allow, and what to remove—and that our mistakes cause a great deal of concern for people, which is why we need to allow the option to request review of the decision and provide additional context that will help our team see the fuller picture as they review the post again. This type of feedback will allow us to continue improving our systems and processes so we can prevent similar mistakes in the future.

See also Response to Question 44.

Question 57. In October 2017, the social-media company Twitter refused to permit Representative Marsha Blackburn to pay to promote a campaign advertisement because Rep. Blackburn stated that she fought to stop the sale of children’s body parts. Twitter’s explanation was that Blackburn’s critique of “the sale of baby body parts” was an “inflammatory statement” that Twitter refused to advertise.

Does Representative Blackburn’s campaign advertisement (available readily on the internet) violate Facebook’s policies regarding acceptable advertisements?

Does Representative Blackburn’s campaign advertisement violate Facebook’s policies against “hate speech?”

Would the statement, standing alone, that Planned Parenthood sells baby body parts qualify as “hate speech?”

Would Facebook censor or otherwise downgrade or make unavailable the statement that Planned Parenthood sells baby body parts for any other reason?

As Facebook indicated publicly in October 2017, Representative Blackburn’s campaign advertisement, in which she mentioned “the sale of baby body parts” does not violate our Advertising Policies or our Community Standards.

We work to strike the right balance between enabling free expression around the globe and ensuring that our platform is safe. We currently define hate speech as anything that directly attacks people based on protected characteristics—race, ethnicity, national origin, religious affiliation, sexual orientation, sex, gender, gender identity, or serious disability or disease. We remove content that violates our policies, regardless of who posted the content, including the government.

Our policies allow content that may be controversial and at times even distasteful, but which does not cross the line into hate speech. This may include criticism of public figures, religions, professions, and political ideologies.

Question 58. Louis Farrakhan presently employs Facebook to reach numerous individuals. At present, he has over a million followers.

See Response to Question 59.

Question 59. On his Facebook page, Farrakhan links to an open letter of his which states: “We can now present to our people and the world a true, undeniable record of the relationship between Blacks and Jews from their own mouths and pens. These scholars, Rabbis and historians have given to us an undeniable record of Jewish anti-Black behavior, starting with the horror of the trans-Atlantic slave trade, plantation slavery, Jim Crow, sharecropping, the labor movement of the North and South, the unions and the misuse of our people that continues to this very moment.”

Does this statement violate Facebook’s policies against “hate speech?”

If so, why has this post been permitted to remain?

If not, why not?

On his Facebook page, Farrakhan links to a sermon in which he describes the “Synagogue of Satan” and its attempts to harm him.

Is the term “Synagogue of Satan” a violation of Facebook’s policies against “hate speech?”

If so, why has this post been permitted to remain?

If not, why not?

We are committed to designing our products to give all people a voice and foster the free flow of ideas and culture. That said, when something crosses the line into hate speech, it has no

place on Facebook, and we are committed to removing it from our platform any time we become aware of it.

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, sex, gender, gender identity, and serious disability or disease. We also provide some protections for immigration status. We define attack as violent or dehumanizing speech, statements of inferiority, and calls for exclusion or segregation. The detailed guidelines our reviewers use to assess whether content violates our hate speech policies are available at https://www.facebook.com/communitystandards/objectionable_content/hate_speech.

Question 60. In June 2013, Facebook blocked the following post written by Fox News Radio’s Todd Starnes for violating Facebook’s community standards, “I’m about as politically incorrect as you can get. I’m wearing an NRA ball cap, eating a Chick-fil-A sandwich, reading a Paula Deen cookbook and sipping a 20-ounce sweet tea while sitting in my Cracker Barrel rocking chair with the Gather Vocal Band singing ‘Jesus Saves’ on the stereo and a Gideon’s Bible in my pocket. Yes sir, I’m politically incorrect and happy as a June bug.” Although Facebook ultimately reversed its decision, for several hours, Todd Starnes could not access either his fan or person page.

Why was Todd Starnes’ post removed?

What Facebook rule was Todd Starnes’ post thought to have violated before it was reinstated?

Was any part of Starnes’ statement “hate speech?”

Was any part of Starnes’ statement considered harassment?

Yes or no: must posted content be “politically correct” to remain in accordance with Facebook’s community standards?

Is a statement that something is not “politically correct” a violation of Facebook’s standards?

The page where Todd Starnes posted the content was not unpublished. He was the administrator that made the post, and the action was taken on his profile. He posted the content at around 2 am on June 29, 2013, and it was restored shortly before 10 am the same day. During that time, he did not lose his ability to access either his profile or his page, just the post itself. When we reinstated the post, we sent him an apology the same day.

Our policies apply equally to individuals and entities across the political spectrum. We are committed to designing our products to give all people a voice and foster the free flow of ideas and culture. That said, when something crosses the line into hate speech, it has no place on Facebook, and we are committed to removing it from our platform any time we become aware of it.

We recognize that our policies are only as good as the strength and accuracy of our enforcement—and our enforcement is not perfect. We make mistakes because our processes involve people, and people are not infallible. We are always working to improve.

When we're made aware of incorrect content removals, we review them with team members so as to prevent similar mistakes in the future. We also audit the accuracy of reviewer decisions on an ongoing basis to coach them and follow up on improving, where errors are being made.

We hope that our recent decision to publicize our detailed Community Standards—which reflect our internal reviewer guidelines—and the introduction of appeals will aid in this process. By providing more clarity on what is and isn't allowed on Facebook, we hope that people will better understand how our policies apply to them. For some violation types, where people believe we have made a mistake, they can request review of our decisions, and we are working to extend this process further by supporting more violation types.

See also Response to Question 44.

Question 61. How many individuals at Facebook have the ability to moderate, remove, downgrade, conceal, or otherwise censor content, ban, suspend, warn, or otherwise discipline users, or approve, price, review, or refuse advertisements on the platform? This question includes individuals with the power to alter search results and similar mechanisms that suggest additional content to users in order to to promote or demote content, whether individually or routinely through an algorithm or by altering any of the platform's search functions. (Please include all employees, independent contractors, or others with such ability at Facebook.)

Into what divisions or groups are those individuals organized?

Who are the individuals responsible for supervising these individuals as their conduct relates to American citizens, nationals, businesses, and groups?

We understand from your April 10 testimony that Facebook has approximately 15,000 to 20,000 moderators. How many individuals have the responsibility to moderate, remove, downgrade, conceal, or otherwise censor content, ban, suspend, warn, or otherwise discipline users, or approve, price, review, or refuse advertisements as a primary or significant function of their role at Facebook? This question includes individuals with the power to alter search results and similar mechanisms that suggest additional content to users in order to to promote or demote content, whether individually or routinely through an algorithm or by altering any of the platform's search functions. (Going forward, we will refer to these individuals, with a primary or significant responsibility for reviewing content, users, or advertisements, as "moderators.")

Who are the individuals responsible for supervising these moderators as their conduct relates to American citizens, nationals, businesses, and groups?

How many moderators has Facebook had on its platform for each of the calendar years 2006 to 2018? Please provide approximations if exact numbers are impossible to obtain.

How many moderators does Facebook intend to retain for the years 2019 and 2020?

On average, how many pieces of content (e.g., a Facebook post, an Instagram photo, and so on) does a moderator remove a day?

On average, how many users does a moderator discipline a day?

On average, how many advertisements does a moderator approve, disapprove, price, consult on, review, or refuse a day?

Our content reviewers respond to millions of reports each week from people all over the world.

Our community of users helps us by reporting accounts or content that may violate our policies. Our content review teams around the world—which grew by 3,000 people last year—work 24 hours a day and in dozens of languages to review these reports. By the end of 2018, we will have doubled the number of people working on safety and security as compared to the beginning of the year—to a total of 20,000.

To help the Facebook community better understand our efforts to enforce the Community Standards, we recently published a Community Standards Enforcement Preliminary Report (<https://transparency.facebook.com/community-standards-enforcement>) describing the amount and types of content we take action against, as well as the amount of content that we flag for review proactively.

We are also committed to getting better at enforcing our advertising policies. We review many ads proactively using automated and manual tools, and reactively when people hide, block, or mark ads as offensive. We are taking aggressive steps to strengthen both our automated and our manual review. We are also expanding our global ads review teams and investing more in machine learning to better understand when to flag and take down ads, such as ads that offer employment or credit opportunity while including or excluding multicultural advertising segments. Enforcement is never perfect, but we will get better at finding and removing improper ads.

As to the questions regarding ranking and algorithmic changes, see Response to Question 66.

Question 62. What percentage of Facebook’s moderators:

Self-identify or are registered as Democrats?

Self-identify or are registered as Republicans?

Would identify themselves as “liberal?”

Would identify themselves as “conservative?”

Have donated to:

The Democratic Party?

A candidate running for office as a Democrat?

A cause primarily affiliated with or supported by the Democratic Party?

A cause primarily affiliated with or supported by liberal interest groups?

A political action committee primarily advocating for the Democratic Party, Democratic candidates or office-holders, or causes primarily supported by the Democratic Party?

The Republican Party?

A candidate running for office as a Republican?

A cause primarily affiliated with or supported by the Republican Party?

A cause primarily affiliated with or supported by conservative interest groups?

A political action committee primarily advocating for the Republican Party, Republican candidates or office-holders, or causes primarily supported by the Republican Party?

Worked on or volunteered for a Democratic campaign?

Worked on or volunteered for a Republican campaign?

Worked on, interned for, or volunteered for a Democratic legislator, State or federal?

Worked on, interned for, or volunteered for a Republican legislator, State or federal?

Worked on or interned for a Democratic administration or candidate?

Worked on or interned for a Republican administration or candidate?

We do not maintain statistics on these data points.

Question 63. What percentage of Facebook's employees:

Self-identify or are registered as Democrats?

Self-identify or are registered as Republicans?

Self-identify as "liberal?"

Self-identify as "conservative?"

Have donated to:

The Democratic National Committee, the Democratic Congressional Campaign Committee, or the Democratic Senatorial Campaign Committee?

A candidate running for office as a Democrat?

A cause primarily affiliated with or supported by the Democratic Party?

A cause primarily affiliated with or supported by liberal interest groups?

A political action committee primarily advocating for the Democratic Party, Democratic candidates or office-holders, or causes primarily supported by the Democratic Party?

The Republican National Committee, the National Republican Senate Committee, or the National Republican Congressional Committee?

A candidate running for office as a Republican?

A cause primarily affiliated with or supported by the Republican Party?

A cause primarily affiliated with or supported by conservative interest groups?

A political action committee primarily advocating for the Republican Party, Republican candidates or office-holders, or causes primarily supported by the Republican Party?

Worked on, interned for, or volunteered for a Democratic candidate campaigning for elected office or an elected Democratic official or candidate?

Worked on, interned for, or volunteered for a Republican campaigning for elected office or an elected Republican official or candidate?

We do not maintain statistics on these data points.

Question 63. What percentage of Facebook's management:

Self-identify or are registered as Democrats?

Self-identify or are registered as Republicans?

Self-identify as "liberal?"

Self-identify as "conservative?"

Have donated to:

The Democratic National Committee, the Democratic Congressional Campaign Committee, or the Democratic Senatorial Campaign Committee?

A candidate running for office as a Democrat?

A cause primarily affiliated with or supported by the Democratic Party?

A cause primarily affiliated with or supported by liberal interest groups?

A political action committee primarily advocating for the Democratic Party, Democratic candidates or office-holders, or causes primarily supported by the Democratic Party?

The Republican National Committee, the National Republican Senate Committee, or the National Republican Congressional Committee?

A candidate running for office as a Republican?

A cause primarily affiliated with or supported by the Republican Party?

A cause primarily affiliated with or supported by conservative interest groups?

A political action committee primarily advocating for the Republican Party, Republican candidates or office-holders, or causes primarily supported by the Republican Party?

Worked on, interned for, or volunteered for an elected Democratic official or candidate?

Worked on, interned for, or volunteered for an elected Republican official or candidate?

We do not maintain statistics on these data points.

Question 64. What percentage of Facebook's executives:

Self-identify or are registered as Democrats?

Self-identify or are registered as Republicans?

Self-identify as "liberal?"

Self-identify as "conservative?"

Have donated to:

The Democratic National Committee, the Democratic Congressional Campaign Committee, or the Democratic Senatorial Campaign Committee?

A candidate running for office as a Democrat?

A cause primarily affiliated with or supported by the Democratic Party?

A cause primarily affiliated with or supported by liberal interest groups?

A political action committee primarily advocating for the Democratic Party, Democratic candidates or office-holders, or causes primarily supported by the Democratic Party?

The Republican National Committee, the National Republican Senate Committee, or the National Republican Congressional Committee?

A candidate running for office as a Republican?

A cause primarily affiliated with or supported by the Republican Party?

A cause primarily affiliated with or supported by conservative interest groups?

A political action committee primarily advocating for the Republican Party, Republican candidates or office-holders, or causes primarily supported by the Republican Party?

Worked on, interned for, or volunteered for an elected Democratic official or candidate?

Worked on, interned for, or volunteered for an elected Republican official or candidate?

We do not maintain statistics on these data points.

Question 65. How many employees has Facebook hired that previously worked for 501(c)(3) or 501(c)(4) nonprofits? Please list the names of the 501(c)(3) and 501(c)(4) organizations employees have previously worked for and the number of employees for each.

We do not maintain statistics on these data points.

Question 66. Based on your testimony, we understand that Facebook conducts many of its editorial and moderating decisions using one or more algorithms.

What editorial and moderating functions do these algorithms undertake?

List and describe the factors that the algorithm evaluates and considers.

Describe what if any human oversight or auditing is in place to review the algorithm's functions.

Do any of the factors in these algorithms associated with promoting, demoting, flagging, removing, suggesting, or otherwise altering the visibility of content correlate strongly (defined as meeting any generally accepted threshold for strong correlation using any generally accepted bivariate or multivariate analysis technique, including, but not limited to, chi-square, ANOVA, MANCOVA, Probit, Logit, regression, etc.) with any of the following traits (if so, please list which factor and its correlation):

Self-identification with the Democratic Party?

Registration as a Democrat?

Self-identification as a liberal?

Self-identification with the Republican Party?

Registration as a Republican?

Self-identification as a conservative?

Do any of these factors correlate significantly (p greater than or equal to .05) with any of the following traits (if so, please list which factor and its correlation):

Self-identification with the Democratic Party?

Registration as a Democrat?

Self-identification as a liberal?

Self-identification with the Republican Party?

Registration as a Republican?

Self-identification as a conservative?

A user's News Feed is made up of stories from their friends, Pages they've chosen to follow and groups they've joined. Ranking is the process we use to organize all of those stories so that users can see the most relevant content at the top, every time they open Facebook. Ranking has four elements: the available inventory of stories; the signals, or data points that can inform ranking decisions; the predictions we make, including how likely we think a user is to comment on a story, share with a friend, etc.; and a relevancy score for each story.

News Feed considers thousands of signals to surface the content that's most relevant to each person who uses Facebook. Our employees don't determine the ranking of any specific piece of content. To help the community understand how News Feed works and how changes to News Feed affect their experience on Facebook, we publish a regularly-updated News Feed FYI blog (<https://newsroom.fb.com/news/category/inside-feed/>) where our team shares details of significant changes.

Question 67. What percentage of the individuals who design, code, implement, monitor, correct, or alter any of these algorithms:

Self-identify as Democrats?

Are registered as Democrats?

Self-identify as liberal?

Self-identify as Republicans?

Are registered as Republicans?

Self-identify as conservative?

We do not maintain statistics on these data points.

Question 68. In 2016, in response to complaints about “fake news” during the 2016 Presidential campaign and following President Trump’s election, Facebook procured the services of specific “fact-checking” outlets in order to flag certain stories or sources as disputed, challenged, or incorrect. Earlier this year, it additionally changed one or more of the algorithms that recommend websites to users, such as users’ news feeds.

On what basis did Facebook select the fact-checking organizations that it enlisted to identify incorrect assertions of fact?

Numerous sources have cited the presence of political bias in many “fact-checking” organizations; for example, according to one 2013 study by George Mason University’s Center for Media and Public Affairs, the site Politifact.com-- which Facebook employs to check facts on its platform--was between two and three times more likely to rate Republicans’ claims as false (32%) than Democrats’ claims (11%), and was between two and three times more likely to rate Democrats’ statements as mostly or entirely true (54%) compared to Republicans’ statements (18%). Indeed, the RealClearPolitics “Fact Check Review” notes that, in the last 120 days, approximately 1/6th of “facts” that Politifact.com claims to check aren’t facts at all, but mere opinions.

What steps does Facebook take to counteract liberal or left-wing bias by fact-checking outlets?

What steps does Facebook intend to take to bring political balance to its fact-checking review process?

What mechanisms for appealing a determination that a statement is false or otherwise disagreed-with does Facebook make available to entities that Politifact (or others) accuse(s) of lying?

If none exist, what mechanisms does Facebook intend to make available?

If none exist, to what extent will Facebook make its review of these claims publicly visible?

Has Facebook ever labeled claims or articles by any of the following entities as false? If so, please identify which claims and when.

Huffington Post

Salon

Slate

ThinkProgress

Media Matters for America

ShareBlue

The Daily Kos

Vice

Vox

To reduce the spread of false news, one of the things we're doing is working with third-party fact checkers to let people know when they are sharing news stories (excluding satire and opinion) that have been disputed or debunked, and to limit the distribution of stories that have been flagged as misleading, sensational, or spammy. Third-party fact-checkers on Facebook are signatories to the non-partisan International Fact-Checking Network Code of Principles. Third-party fact-checkers investigate stories in a journalistic process meant to result in establishing the truth or falsity of the story.

In the United States, Facebook uses third-party fact-checking by the Associated Press, Factcheck.org, PolitiFact, Snopes, and the Weekly Standard Fact Check.

Publishers may reach out directly to the third-party fact-checking organizations if (1) they have corrected the rated content, or if (2) they believe the fact-checker's rating is inaccurate. To issue a correction, the publisher must correct the false content and clearly state that a correction was made directly on the story. To dispute a rating, the publisher must clearly indicate why the original rating was inaccurate. If a rating is successfully corrected or disputed, the demotion on the content will be lifted and the strike against the domain or Page will be removed. It may take a few days to see the distribution for the domain or Page recover. Additionally, any recovery will be affected by other false news strikes and related interventions (like demotions for clickbait). Corrections and disputes are processed at the fact-checker's discretion. Fact-checkers are asked to respond to requests in a reasonable time period—ideally one business day for a simple correction, and up to a few business days for more complex disputes.

We want Facebook to be a place where people can discover more news, information, and perspectives, and we are working to build products that help.

As to the questions regarding ranking and algorithmic changes, see Response to Question 66.

Question 69. TalkingPointsMemo

Does Facebook consider the basis for a fact-checker's determination that something is "false" when choosing to label it as such? For example, as numerous media outlets have noted, some fact-checking outlets concede that the factual statement a public figure has made is true, but then condemn it for lacking "context" or spin favorable to a left-wing politician.

If so, how does Facebook consider it?

If not, does Facebook intend to do so in the future? And if so, how? If not, why not?

When one of Facebook’s fact-checkers determines that a claim is false, how does Facebook determine what material to refer a user to in response? Please list all such sources and any method relied on for determining their priority.

Facebook’s 2018 alteration of its algorithm has had a noted and outsized impact on traffic to conservative websites while not having a similar effect on liberal websites. At least one study by the Western Journal estimated liberal publishers’ traffic from Facebook rose approximately 2% following the change, while conservative publishers’ traffic declined approximately 14%.

In what way(s) did Facebook change its content-screening or news- suggesting algorithms, or any other feature of its website which suggests content to users, in this 2018 instance?

Were any components of these changes intended to have a differential impact on conservative outlets versus liberal ones?

Were any components of these changes expected to have a differential impact on conservative outlets versus liberal ones?

Measured against pre-change traffic, how has the traffic of liberal publishers changed following this 2018 instance?

Measured against pre-change traffic, how has the traffic of conservative publishers changed following this 2018 instance?

Measured against pre-change traffic, how has this 2018 instance changed the traffic of the following publishers:

The Washington Post

The New York Times

The Washington Times

The New York Post

The New York Daily News

Fox News

National Review

The Daily Beast

Huffington Post

Buzzfeed

Newsweek

The Daily Wire

Vice

USA Today

Salon

Slate

Vox

The Daily Caller

The Blaze

PJ Media

The Washington Free Beacon

Reuters

The Associated Press

National Public Radio

Bloomberg

Does Facebook intend to do anything to reduce the differential effect on its recent algorithmic changes on conservative publishers?

If so, what?

If not, why not?

See Response to Question 68.

Question 70. Facebook's Help section explains that the posts that users see are influenced by their connections and activity on Facebook, including the number of comments, likes, and reactions a post receives and what kind of story it is. Some reporting suggests that Facebook's algorithm functions based on the content available (inventory), considerations about the content (signals), considerations about a person (predictions), and overall score.

How do Facebook employees determine how informative a post is or which interactions create a more meaningful experience?

Does a speaker's viewpoint determine in whole or part how informative or meaningful a post is?

Does a speaker's partisan affiliation determine in whole or part how informative or meaningful a post is?

Does a speaker's religious affiliation determine in whole or part how informative or meaningful a post is?

See Response to Question 66.

Question 71. Facebook is entitled to contribute money to federal and State elections both as a function of the First Amendment as well as of federal and State law. Including all of its subsidiaries, affiliates, as well as political action committees, partnerships, councils, groups, or entities organized with either a sole or significant purpose of electioneering, making political contributions to issue advocacy, candidates, or political parties, or of bundling or aggregating money for candidates or issue or party advocacy, whether disclosed by law or not, and during primary elections or general elections, how much money has Facebook contributed to:

All federal, State, and local candidates for office from 2008 to present?

All national party committees?

Of that amount, how much was to:

The Democratic National Committee?

The Democratic Senatorial Campaign Committee?

The Democratic Congressional Campaign Committee?

The Republican National Committee?

The National Republican Senate Committee?

The National Republican Congressional Committee?

All political action committees (or other groups outlined above in question 43) from 2008 to present?

All issue-advocacy campaigns, including initiatives, referenda, ballot measures, and other direct-democracy or similar lawmaking measures?

Candidates running for President:

In 2008?

How much of that money was to the Democratic candidate?

How much of that money was to the Republican candidate?

How much of that money was to other candidates?

In 2012?

How much of that money was to the Democratic candidate?

How much of that money was to the Republican candidate?

How much of that money was to other candidates?

In 2016?

How much of that money was to the Democratic candidate?

How much of that money was to the Republican candidate?

How much of that money was to other candidates?

Candidates running for the U.S. Senate: (for special or off-year elections going forward, please group donation amounts with the next nearest cycle)

In 2008?

How much of that money was to Democratic candidates?

How much of that money was to Republican candidates?

How much of that money was to other candidates?

In 2010?

How much of that money was to Democratic candidates?

How much of that money was to Republican candidates?

How much of that money was to other candidates?

In 2012?

How much of that money was to Democratic candidates?

How much of that money was to Republican candidates?

How much of that money was to other candidates?

In 2014?

How much of that money was to Democratic candidates?

How much of that money was to Republican candidates?

How much of that money was to other candidates?

In 2016?

How much of that money was to Democratic candidates?

How much of that money was to Republican candidates?

How much of that money was to other candidates?

In 2018?

How much of that money was to Democratic candidates?

How much of that money was to Republican candidates?

How much of that money was to other candidates?

Candidates running for the U.S. House of Representatives:

In 2008?

How much of that money was to Democratic candidates?

How much of that money was to Republican candidates?

How much of that money was to other candidates?

In 2010?

How much of that money was to Democratic candidates?

How much of that money was to Republican candidates?

How much of that money was to other candidates?

In 2012?

How much of that money was to Democratic candidates?

How much of that money was to Republican candidates?

How much of that money was to other candidates?

In 2014?

How much of that money was to Democratic candidates?

How much of that money was to Republican candidates?

How much of that money was to other candidates?

In 2016?

How much of that money was to Democratic candidates?

How much of that money was to Republican candidates?

How much of that money was to other candidates?

In 2018?

How much of that money was to Democratic candidates?

How much of that money was to Republican candidates?

How much of that money was to other candidates?

Candidates running for Governor:

In 2008?

How much of that money was to Democratic candidates?

How much of that money was to Republican candidates?

How much of that money was to other candidates?

In 2010?

How much of that money was to Democratic candidates?

How much of that money was to Republican candidates?

How much of that money was to other candidates?

In 2012?

How much of that money was to Democratic candidates?

How much of that money was to Republican candidates?

How much of that money was to other candidates?

In 2014?

How much of that money was to Democratic candidates?

How much of that money was to Republican candidates?

How much of that money was to other candidates?

In 2016?

How much of that money was to Democratic candidates?

How much of that money was to Republican candidates?

How much of that money was to other candidates?

In 2018?

How much of that money was to Democratic candidates?

How much of that money was to Republican candidates?

How much of that money was to other candidates?

Political action committees or other groups mentioned in question 43 that:

Contribute 75% or more of their money to Democratic candidates for office?

Contribute 75% or more of their money to Republican candidates for office?

Identify as liberal, progressive, or otherwise left-wing?

Identify as conservative or right-wing?

Facebook complies with all political contribution reporting requirements, and such reports are publicly available. For more information on Facebook's contributions, please see <https://newsroom.fb.com/news/h/facebook-political-engagement/>.

Question 72. How much has Facebook donated, either in the form of money or services (including free or discounted advertising or more prominent placements within the platform via searches and other suggested-content mechanisms), to the following not-for-profit organizations (or their affiliates or subsidiaries) in the last 10 years? (Please separate answers into cash and non-cash components.)

Planned Parenthood

NARAL

The Center for Reproductive Rights

The National Right to Life Committee

Americans United for Life

Everytown for Gun Safety

The Brady Campaign

The National Rifle Association

Gun Owners of America

Human Rights Campaign

Amnesty International

Lambda Legal

National Immigration Forum

Federation

GLAAD

ACLU

UnidosUS (formerly “La Raza” or the “National Council of La Raza”)

The Sierra Club

Greenpeace

The Heritage Foundation

The Cato Institute

The Institute for Justice

Southern Poverty Law Center

The Open Society Foundation(s)

Americans for Prosperity

We partner with various domestic and international non-governmental organizations, which span the political and ideological spectrum. We provide our partners with technical expertise, sponsorships, advertising credits, and trainings, among other support. Our partnerships are crucial to our mission of building community. More information about our partnerships is available at <https://newsroom.fb.com/news/h/facebook-political-engagement/>.

Question 73. Facebook sells advertisements to political candidates and organizations. Multiple sources report that Facebook charged different rates to the Hillary Clinton and Donald Trump campaigns during the 2016 election. For the following questions, to the extent that geographic or local-market concerns significantly explain disparate rates

between candidates, please explain how they do so and to what extent they do so, including calculations justifying that explanation.

Did Facebook charge the two campaigns different rates?

If so, on what basis?

If so, what rates did Facebook charge:

The Clinton Campaign?

The Trump Campaign?

If these campaigns purchased advertising rates on Facebook or its platforms, what rates did Facebook charge each of the following campaigns?

Barack Obama's 2008 campaign

John McCain's 2008 campaign

Barack Obama's 2012 campaign

Mitt Romney's 2012 campaign

On average, and among campaigns that purchased advertisements, what rates did Facebook charge:

Democrats running for Senate in 2008?

Republicans running for Senate in 2008?

Democrats running for the House of Representatives in 2008?

Republicans running for the House of Representatives in 2008?

Democrats running for Governor in 2008?

Republicans running for Governor in 2008?

Democrats running in State or local legislative races in 2008?

Republicans running in State or local legislative races in 2008?

Democrats running for Senate in 2010?

Republicans running for Senate in 2010?

Democrats running for the House of Representatives in 2010?

Republicans running for the House of Representatives in 2010?

Democrats running for Governor in 2010?

Republicans running for Governor in 2010?

Democrats running in State or local legislative races in 2010?

Republicans running in State or local legislative races in 2010?

Democrats running for Senate in 2012?

Republicans running for Senate in 2012?

Democrats running for the House of Representatives in 2012?

Republicans running for the House of Representatives in 2012?

Democrats running for Governor in 2012?

Republicans running for Governor in 2012?

Democrats running in State or local legislative races in 2014?

Republicans running in State or local legislative races in 2014?

Democrats running for Senate in 2014?

Republicans running for Senate in 2014?

Democrats running for the House of Representatives in 2014?

Republicans running for the House of Representatives in 2014?

Democrats running for Governor in 2014?

Republicans running for Governor in 2014?

Democrats running in State or local legislative races in 2014?

Republicans running in State or local legislative races in 2014?

Democrats running in State or local legislative races in 2016?

Republicans running in State or local legislative races in 2016?

Democrats running for Senate in 2016?

Republicans running for Senate in 2016?

Democrats running for the House of Representatives in 2016?

Republicans running for the House of Representatives in 2016?

Democrats running for Governor in 2016?

Republicans running for Governor in 2016?

Democrats running in State or local legislative races in 2016?

Republicans running in State or local legislative races in 2016?

Democrats running in State or local legislative races in 2018?

Republicans running in State or local legislative races in 2018?

Democrats running for Senate in 2018?

Republicans running for Senate in 2018?

Democrats running for the House of Representatives in 2018?

Republicans running for the House of Representatives in 2018?

Democrats running for Governor in 2018?

Republicans running for Governor in 2018?

Democrats running in State or local legislative races in 2018?

Republicans running in State or local legislative races in 2018?

Yes or no: does Facebook consider partisan affiliation in deciding whether to sell advertisements to a political candidate, political action committee, or other organization purchasing political advertisements?

Yes or no: does Facebook consider partisan affiliation in deciding at what rates to sell advertisements to a political candidate, political action committee, or other organization purchasing political advertisements?

Yes or no: does Facebook consider the likelihood of a candidate's ultimate electoral success (via polls or otherwise) in deciding whether to sell advertisements to a political candidate?

Yes or no: does Facebook consider the likelihood of a candidate's ultimate electoral success (via polls or otherwise) in deciding at what rates to sell advertisements to a political candidate?

Facebook offered identical support to both the Trump and Clinton campaigns, and had teams assigned to both. Everyone had access to the same tools, which are the same tools that every campaign is offered.

See also Response to Question 74.

Question 74. Please provide Facebook’s advertising rates for each U.S. Senate and U.S. House election for which Facebook quoted or sold advertisements to one or more candidates for the years 2008, 2010, 2012, 2014, 2016, and 2018. For elections not falling in those years or special elections, please provide and group these rates with the next sequential election cycle. Where Facebook offered or sold advertising to multiple candidates within the same race, please pair those quotes or prices together along with party affiliation.

People can run ads on Facebook, Instagram, and Audience Network on any budget. The exact cost associated with an ad being shown to someone is determined in Facebook’s ad auction.

Question 75. Yes or no: has Facebook ever provided at no cost advertising to political candidates, campaign committees, political action committees or similar groups, or issue-advocacy groups or campaigns, whether through outright advertising or by altering search rankings, trending topics, content rankings, or the position of content within any suggested content mechanism?

If so, please provide each instance in which Facebook has done so and indicate whether Facebook offered similar support to any other candidate or issue in that race or election.

If so, please indicate whether Facebook coordinated with that campaign, candidate, or issue in doing so, or if Facebook acted unilaterally.

Political candidates, campaign committees, political action committees and similar groups, as well as issue advocacy groups and campaigns can set up Facebook Pages for free and post free content via those Pages, in the same way that any Page creator may. To run ads on Facebook, a form of payment must be provided. The algorithms that set content rankings are not designed to promote any candidate or party.

Question 76. Please list and describe all mandatory trainings that Facebook employees are required to undergo and the topics involved in each, including any trainings on sexual harassment, unconscious bias, racial privilege, and inclusivity.

At Facebook, we treat any allegations of harassment, discrimination, or retaliation with the utmost seriousness, and we have invested significant time and resources into developing our policies and processes. We have made our policies and processes available publicly—not because we think we have all the answers, but because we believe that the more companies are open about their policies, the more we can all learn from one another. Our internal policies on sexual harassment and bullying are available on our Facebook People Practices website (<http://peoplepractices.fb.com/>), along with details of our investigation process and tips and resources we have found helpful in preparing our Respectful Workplace internal trainings. Our

philosophy on harassment, discrimination, and bullying is to go above and beyond what is required by law. Our policies prohibit intimidating, offensive, and sexual conduct even when that conduct might not meet the legal standard of harassment. Even if it's legally acceptable, it's not the kind of behavior we want in our workplace. In developing our policies, we were guided by six basic principles:

- First, develop training that sets the standard for respectful behavior at work, so people understand what's expected of them right from the start. In addition to prescribing mandatory harassment training, we wrote our own unconscious bias training program at Facebook, which is also available publicly on our People Practices website. Our training includes Sustainable Equity, a three-day course in the US about racial privilege and injustice, and Design for Inclusion, a multi-day course in the UK to educate on systemic inequity.
- Second, treat all claims—and the people who voice them—with seriousness, urgency, and respect. At Facebook, we make sure to have HR business partners available to support everyone on the team, not just senior leaders.
- Third, create an investigation process that protects employees from stigma or retaliation. Facebook has an investigations team made up of experienced HR professionals and lawyers trained to handle sensitive cases of sexual harassment and assault.
- Fourth, follow a process that is consistently applied in every case and is viewed by employees as providing fair procedures for both victims and those accused.
- Fifth, take swift and decisive action when it is determined that wrongdoing has occurred. We have a zero-tolerance policy, and that means that when we are able to determine that harassment has occurred, those responsible are fired. Unfortunately, in some cases investigations are inconclusive and come down to one person's word against another's. When we don't feel we can make a termination decision, we take other actions designed to help everyone feel safe, including changing people's roles and reporting lines.
- Sixth, make it clear that all employees are responsible for keeping the workplace safe—and anyone who is silent or looks the other way is complicit. There's no question that it is complicated and challenging to get this right. We are by no means perfect, and there will always be bad actors. Unlike law enforcement agencies, companies don't have access to forensic evidence and instead have to rely on reported conversations, written evidence, and the best judgment of investigators and legal experts. What we can do is be as transparent as possible, share best practices, and learn from one another—recognizing that policies will evolve as we gain experience. We don't have everything worked out at Facebook on these issues, but we will never stop striving to make sure we have a safe and respectful working environment for all our people.

We are also working to reduce unconscious bias. Our publicly available Managing Unconscious Bias class encourages our people to challenge and correct bias as soon as they see it—in others, and in themselves. We’ve also doubled down by adding two additional internal programs: Managing Inclusion, which trains managers to understand the issues that affect marginalized communities, and Be The Ally, which gives everyone the common language, tools, and space to practice supporting others.

Question 77. Please list and describe all optional recommended trainings that Facebook employees are required to undergo and the topics involved in each, including any trainings on sexual harassment, unconscious bias, racial privilege, and inclusivity.

See Response to Question 76.

Question 78. Do any of the materials Facebook uses in any of these trainings identify different preferences, values, goals, ideas, world-views, or abilities among individuals on the basis of the following? If so, please list each and include those materials.

Race

Sex

Sexual orientation

Place of origin

Diversity is core to our business at Facebook and we’re committed to building and maintaining a workforce as diverse and inclusive as the people and communities we serve. We have developed and implemented programs and groups to help build a more diverse and inclusive company, and to better engage and support employees from diverse backgrounds. We have a number of Facebook Resource Groups (FBRGs) that are run by our internal communities from different backgrounds, such as Asians and Pacific Islanders, African-Americans, People with Disabilities, those of faith, Latinos/Hispanics, LGBTQ, Veterans, and women. These FBRGs provide members with support, foster understanding between all people, and can coordinate programming to further support members. Examples of such programs include Women@ Leadership Day, Black@ Leadership Day, Latin@ Leadership Day, and Pride@ Leadership Day. Facebook also values and creates programming to support its Veterans and People with Disabilities through dedicated program managers and recruiters, mentoring programs and awareness campaigns to promote education and inclusion. These groups and programs are created to support and provide a more inclusive work experience for people from diverse backgrounds, with membership and participation open even to those who do not self-identify with these groups. For example, people who do not self-identify as Black are still members of Black@ and have attended Black@ Leadership Day, and there are male members of Women@ and men can attend Women@ Leadership Day. Facebook is also an Equal Opportunity Employer.

Question 79. Facebook acknowledges that it is located in a very liberal part of the country, and has suggested that it understands that many of its employees as well as the surrounding community share a particular (very liberal) culture.

Does Facebook have any training specifically aimed at discouraging political, ideological, or partisan bias in decision-making by its employees?

Does Facebook have any training specifically aimed at discouraging political, ideological, or partisan bias in hiring, retention, promotion, and firing of its employees?

Does Facebook have any training specifically aimed at discouraging political, ideological, or partisan bias in the monitoring and supervision of content, users, or advertisements on each of its platforms?

Our Community Standards are global and all reviewers use the same guidelines when making decisions.

They undergo extensive training when they join and, thereafter, are regularly trained and tested with specific examples on how to uphold the Community Standards and take the correct action on a piece of content. This training includes when policies are clarified, or as they evolve.

We seek to write actionable policies that clearly distinguish between violating and non-violating content and we seek to make the decision making process for reviewers as objective as possible.

Our reviewers are not working in an empty room. There are quality control mechanisms as well as management on site to help or seek guidance from if needed. When a reviewer isn't clear on the action to take based on the Community Standards, they can pass the content decision to another team for review.

We also audit the accuracy of reviewer decisions on an ongoing basis to coach them and follow up on improving, where errors are being made.

When we're made aware of incorrect content removals, we review them with our Community Operations team so as to prevent similar mistakes in the future.

We are introducing the right to appeal our decisions on individual posts so users can ask for a second opinion when they think we've made a mistake. As a first step, we are launching appeals for posts that were removed for nudity/sexual activity, hate speech or graphic violence. We are working to extend this process further, by supporting more violation types, giving people the opportunity to provide more context that could help us make the right decision, and making appeals available not just for content that was taken down, but also for content that was reported and left up. We believe giving people a voice in the process is another essential component of building a fair system.

Question 80. Please list the names of any third-party organizations or vendors that Facebook uses to facilitate its trainings.

We have a comprehensive training program that includes many hours of live instructor-led training, as well as hands-on practice for all of our reviewers.

All training materials are created in partnership with our policy team and in-market specialists or native speakers from the region.

After starting, reviewers are regularly trained and tested with specific examples on how to uphold the Community Standards and take the correct action on a report. Additional training happens continuously and when policies are clarified, or as they evolve.

Question 81. In the last five years, how many discrimination complaints has Facebook received from Christians? Please indicate how these complaints were resolved.

Decisions about content are made based on whether content violates our Community Standards. A user's personal characteristics do not influence the decisions we make, and Facebook does not track the religious beliefs or other personal characteristics of complainants.

Question 82. Yes or no: Does Facebook offer any compensation, amenities, trainings, or similar services to its employees on account of their race, sex, sexual orientation, or religious affiliation? If so, please list each and whether all other races, sexes, etc. are provided the same compensation, amenity, etc.

See Response to Question 78.

Question 83. In August 2017, Google fired James Damore for violating its code of conduct after Damore submitted an internal memo criticizing the company's hiring practices and arguing that the company's political bias created a negative work environment.

Yes or no: Does Facebook agree with Google's decision to fire James Damore?

Would an individual at Facebook have been fired for publishing a memorandum like Damore's? Assume no previous negative disciplinary history.

Does Facebook permit employees to believe that some portion of the career differences between men and women are the result of differing choices between the sexes?

Would a Facebook employee be disciplined for mentioning that opinion in a conversation to a willing participant?

Would a Facebook employee be disciplined for mentioning that opinion on his or her Facebook account?

Does Facebook permit employees to criticize its "diversity" efforts as being racist against whites or sexist against men?

Would a Facebook employee be disciplined for mentioning that opinion in a conversation to a willing participant?

Would a Facebook employee be disciplined for mentioning that opinion on his or her Facebook account?

We try to run our company in a way where people can express different opinions internally. We are not in a position to comment on the personnel decisions of another company or to engage in speculation about how we might respond in particular hypothetical circumstances.

Question 84. In October 2017, Prager University filed suit against Google and Youtube, alleging that the two companies illegally discriminated against Prager University because of its conservative political perspective. As evidence, Prager University pointed to the dozens of educational videos that Youtube either put in “restricted mode” or demonetized.

Yes or no: Does Facebook agree with YouTube/Google’s decision to restrict the following Prager University video, and if so, why?

The World’s Most Persecuted Minority: Christians?

Israel’s Legal Founding?

Are the Police Racist?

Why Did America Fight the Korean War?

What Should We Do About Guns?

Why America Must Lead?

The Most Important Question About Abortion?

Yes or no: Does Facebook agree with YouTube/Google’s decision to demonetize the following Prager University video, and if so, why?

Are The Police Racist?

Israel’s Legal Founding

The Most Important Question About Abortion?

Who’s More Pro-Choice: Europe or America?

Why Do People Become Islamic Extremists?

Is the Death Penalty Ever Moral?

Why Isn’t Communism as Hated as Nazism?

Radical Islam: The Most Dangerous Ideology?

Is Islam a Religion of Peace?

See Response to Question 44.

Question 85. Recently, Jack Dorsey, Twitter’s CEO, praised an article by two Democrats calling for a “new civil war” against the Republican Party, in which “the entire Republican Party, and the entire conservative movement that has controlled it for the past four decades” will be given a “final takedown that will cast them out” to the “political wilderness” “for a generation or two.”

Does you agree with the premise of this article? It is located here:

<https://medium.com/s/state-of-the-future/the-great-lesson-of-california-in-americas-new-civil-war-e52e2861f30>

Do you or Facebook believe it is appropriate for its platform or company to call for a “new civil war?”

Do you or Facebook believe it is appropriate for its platform or company to call for an end to one of the Nation’s two major political parties?

Do you or Facebook believe it is appropriate for its platform or company to call for an end to the conservative movement?

Do you or Facebook condemn Twitter for calling for an end to the Republican Party?

Do you or Facebook condemn Twitter for calling for an end to the conservative movement?

Do you or Facebook condemn Twitter for calling for a new American civil war?

We are not in a position to comment on the decisions of another company or on another company’s executive’s statements about a news articles.

We are committed to designing our products to give all people a voice and foster the free flow of ideas and culture. That said, when something crosses the line into hate speech, it has no place on Facebook, and we are committed to removing it from our platform any time we become aware of it.

Question 84. Does Facebook collect information regarding its users’:

Usage of non-Facebook apps?

Email?

Audio or ambient sound?

Telephone usage?

Text messaging?

iMessaging?

Physical location when the user is not using the Facebook app?

Spending?

As explained in our Data Policy, we collect three basic categories of data about people:

- (1) data about things people do and share (and who they connect with) on our services,
- (2) data about the devices people use to access our services, and
- (3) data we receive from partners, including the websites and apps that use our business tools.

As far as the amount of data we collect about people, the answer depends on the person. People who have only recently signed up for Facebook have usually shared only a few things—such as name, contact information, age, and gender. Over time, as people use our products, we receive more data from them, and this data helps us provide more relevant content and services. That data will fall into the categories noted above, but the specific data we receive will, in large part, depend on how the person chooses to use Facebook. For example, some people use Facebook to share photos, so we receive and store photos for those people. Some people enjoy watching videos on Facebook; when they do, we receive information about the video they watched, and we can use that information to help show other videos in their News Feeds. Other people seldom or never watch videos, so we do not receive the same kind of information from them, and their News Feeds are likely to feature fewer videos.

The data we have about people also depends on how they have used our controls. For example, people who share photos can easily delete those photos. The same is true of any other kind of content that people post on our services. Through Facebook’s Activity Log tool, people can also control the information about their engagement—i.e., their likes, shares and comments—with other people’s posts. The use of these controls of course affects the data we have about people.

Question 86. Does Facebook give its users the opportunity to opt out of Facebook collecting its users’ data while still using the service?

The Ad Preferences tool on Facebook shows people the advertisers whose ads the user might be seeing because they visited the advertisers’ sites or apps. The person can remove any of these advertisers to stop seeing their ads.

In addition, the person can opt out of these types of ads entirely—so he or she never sees those ads on Facebook based on information we have received from other websites and apps.

We’ve also announced plans to build Clear History, a feature that will enable people to see the websites and apps that send us information when they use them, delete this information from their accounts, and turn off our ability to store it associated with their accounts going forward.

Apps and websites that use features such as the Like button or Facebook Analytics send us information to make their content and ads better. We also use this information to make user experience on Facebook better.

If a user clears his or her history or uses the new setting, we'll remove identifying information so a history of the websites and apps the user used won't be associated with the user's account. We'll still provide apps and websites with aggregated analytics—for example, we can build reports when we're sent this information so we can tell developers if their apps are more popular with men or women in a certain age group. We can do this without storing the information in a way that's associated with the user's account, and as always, we don't tell advertisers who users are.

It will take a few months to build Clear History. We'll work 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 heard specific demands for controls like these at a session we held at our headquarters. We're looking forward to doing more.

Question 87. Yes or no: In preparation for the April 10, 2018 hearing, did Facebook, employees of Facebook, or independent contractors hired by Facebook examine the personal Facebook pages of the U.S. Senators scheduled to take part in the hearing?

If so, please identify the Facebook pages visited and the information sought.

If so, please identify the individuals who sought such information and what information they obtained.

If so, please identify all individuals who possessed or reviewed that information.

While Facebook employees regularly look at the public pages of members of Congress to track the issues that are important to them, we are confident that no employees accessed any private data on personal profiles to prepare for the hearing or the questions for the record.

Question 88. Yes or no: In preparation for the April 10, 2018 hearing, did Facebook, employees of Facebook, or independent contractors hired by Facebook examine the personal Facebook pages of U.S. Senators' family members?

If so, please identify the Facebook pages visited and the information sought.

If so, please identify the individuals who sought such information and what information they obtained.

If so, please identify all individuals who possessed or reviewed that information.

See Response to Question 87.

Question 89. Yes or no: In preparation for the April 10, 2018 hearing, did Facebook, employees of Facebook, or independent contractors hired by Facebook examine the personal Facebook pages of any Senate employees?

If so, please identify the Facebook pages visited and the information sought.

If so, please identify the individuals who sought such information and what information they obtained.

If so, please identify all individuals who possessed or reviewed that information.

See Response to Question 87.

Question 90. Yes or no: In responding to these or any other questions for the record arising from the April 10, 2018 hearing, did Facebook, employees of Facebook, or independent contractors hired by Facebook examine the personal Facebook pages of the U.S. Senators scheduled to take part in the hearing?

If so, please identify the Facebook pages visited and the information sought.

If so, please identify the individuals who sought such information and what information they obtained.

If so, please identify all individuals who possessed or reviewed that information.

See Response to Question 87.

Question 91. Yes or no: In responding to these or any other questions for the record arising from the April 10, 2018 hearing, did Facebook, employees of Facebook, or independent contractors hired by Facebook examine the personal Facebook pages of U.S. Senators' family members?

If so, please identify the Facebook pages visited and the information sought.

If so, please identify the individuals who sought such information and what information they obtained.

If so, please identify all individuals who possessed or reviewed that information.

See Response to Question 87.

Question 92. Yes or no: In responding to these or any other questions for the record arising from the April 10, 2018 hearing, did Facebook, employees of Facebook, or independent contractors hired by Facebook examine the personal Facebook pages of U.S. Senate employees?

If so, please identify the Facebook pages visited and the information sought.

If so, please identify the individuals who sought such information and what information they obtained.

If so, please identify all individuals who possessed or reviewed that information.

See Response to Question 87.

Question 93. Yes or no: Does Facebook collect data on individuals who are not registered Facebook users?

If so, does Facebook use this data as part of the advertising products it sells?

If so, does Facebook share or has Facebook ever shared this data with third parties?

Facebook does not create profiles for people who do not hold Facebook accounts.

When people visit apps or websites that feature our technologies—like the Facebook Like or Comment button—our servers automatically log (i) standard browser or app records of the fact that a particular device or user visited the website or app (this connection to Facebook’s servers occurs automatically when a person visits a website or app that contains our technologies, such as a Like button, and is an inherent function of Internet design); and (ii) any additional information the publisher of the app or website chooses to share with Facebook about the person’s activities on that site (such as the fact that a purchase was made on the site). This is a standard feature of the Internet, and most websites and apps share this same information with multiple different third-parties whenever people visit their website or app. For example, the Senate Commerce Committee’s website shares information with Google and its affiliate DoubleClick and with the analytics company Webtrends. This means that, when a person visits the Committee’s website, it sends browser information about their visit to each one of those third parties. More information about how this works is available at <https://newsroom.fb.com/news/2018/04/data-off-facebook/>.

When the person visiting a website featuring Facebook’s tools is not a registered Facebook user, Facebook does not have information identifying that individual, and it does not create profiles for this individual.

We use the browser and app logs that apps and websites send to us—described above—in the following ways for non-Facebook users. First, these logs are critical to protecting the security of Facebook and to detecting or preventing fake account access. For example, if a browser has visited hundreds of sites in the last five minutes, that’s a sign the device might be a bot, which would be an important signal of a potentially inauthentic account if that browser then attempted to register for an account. Second, we aggregate those logs to provide summaries and insights to websites and apps about how many people visit or use their product, or use specific features like our Like button—but without providing any information about a specific person. We do not create profiles for non-Facebook users, nor do we use browser and app logs for non-Facebook users to show targeted ads from our advertisers to them or otherwise seek to personalize the content they see. However, we may take the opportunity to show a general ad that is unrelated to the attributes of the person or an ad encouraging the non-user to sign up for Facebook.

Question 94. To the extent that Facebook collects and uses data from individuals who are not registered Facebook users, has Facebook gained consent from those individuals to collect and use their personal data?

Facebook does not create profiles about or track web or app browsing history for people who are not registered users of Facebook.

Question 95. To the extent that Facebook collects and uses data from individuals who are registered Facebook users, has Facebook obtained those individuals' informed consent on an opt-in basis prior to the acquisition of that data?

If so, please provide the basis for concluding that data was acquired on an informed consent basis.

If so, please provide the basis for concluding that users opted-in to Facebook's collection and commercialization of their data.

All users must expressly consent to Facebook's Terms and Data Policy when registering for Facebook. The Data Policy explains the kinds of information we collect, how we use this information, how we share this information, and how users can manage and delete information. After joining Facebook, people are presented with the opportunity to consent to additional data collection and uses, such as the use of location or the users' address book on their mobile device.

In response to your specific questions, depending on which Services a person uses, we collect different kinds of information from or about them. This is described in our Data Policy:

Things users and others do and provide.

- **Information and content users provide.** We collect the content, communications and other information users provide when they use our Products, including when they sign up for an account, create or share content, and message or communicate with others. This can include information in or about the content they provide (like metadata), such as the location of a photo or the date a file was created. It can also include what they see through features we provide, such as our camera, so they can do things like suggest masks and filters that users might like, or give them tips on using camera formats. Our systems automatically process content and communications users and others provide to analyze context and what's in them for the purposes described below.
 - Data with special protections: Users can choose to provide information in their Facebook profile fields or Life Events about their religious views, political views, who they are "interested in," or their health. This and other information (such as racial or ethnic origin, philosophical beliefs or trade union membership) could be subject to special protections under the laws of a user's country.
- **Networks and connections.** We collect information about the people, Pages, accounts, hashtags and groups users are connected to and how users interact with them across our Products, such as people users communicate with the most or groups they are part of. We also collect contact information if users choose to upload, sync or import it from a device (such as an address book or call log or SMS log history), which we use for things like helping users and others find people they may know and for the other purposes listed below.

- **Users' usage.** We collect information about how users use our Products, such as the types of content they view or engage with; the features they use; the actions they take; the people or accounts they interact with; and the time, frequency and duration of their activities. For example, we log when users are using and have last used our Products, and what posts, videos and other content users view on our Products. We also collect information about how users use features like our camera.
- **Information about transactions made on our Products.** If users use our Products for purchases or other financial transactions (such as when they make a purchase in a game or make a donation), we collect information about the purchase or transaction. This includes payment information, such as their credit or debit card number and other card information; other account and authentication information; and billing, shipping and contact details.
- **Things others do and information they provide about users.** We also receive and analyze content, communications and information that other people provide when they use our Products. This can include information about users, such as when others share or comment on a photo of them, send a message to them, or upload, sync or import their contact information.

Device Information

- As described below, we collect information from and about the computers, phones, connected TVs and other web-connected devices users use that integrate with our Products, and we combine this information across different devices users use. For example, we use information collected about users' use of our Products on their phone to better personalize the content (including ads) or features they see when they use our Products on another device, such as their laptop or tablet, or to measure whether they took an action in response to an ad we showed them on their phone on a different device.
- Information we obtain from these devices includes:
 - Device attributes: information such as the operating system, hardware and software versions, battery level, signal strength, available storage space, browser type, app and file names and types, and plugins.
 - Device operations: information about operations and behaviors performed on the device, such as whether a window is foregrounded or backgrounded, or mouse movements (which can help distinguish humans from bots).
 - Identifiers: unique identifiers, device IDs, and other identifiers, such as from games, apps or accounts users use, and Family Device IDs (or other identifiers unique to Facebook Company Products associated with the same device or account).

- Device signals: Bluetooth signals, and information about nearby Wi-Fi access points, beacons, and cell towers.
- Data from device settings: information users allow us to receive through device settings they turn on, such as access to their GPS location, camera or photos.
- Network and connections: information such as the name of users' mobile operator or ISP, language, time zone, mobile phone number, IP address, connection speed and, in some cases, information about other devices that are nearby or on their network, so we can do things like help users stream a video from their phone to their TV.
- Cookie data: data from cookies stored on a user's device, including cookie IDs and settings. Learn more about how we use cookies in the Facebook Cookies Policy (<https://www.facebook.com/policies/cookies/>) and Instagram Cookies Policy (<https://www.instagram.com/legal/cookies/>)

Information from Partners

- Advertisers, app developers, and publishers can send us information through Facebook Business Tools they use, including our social plug-ins (such as the Like button), Facebook Login, our APIs and SDKs, or the Facebook pixel. These partners provide information about users' activities off Facebook—including information about their device, websites they visit, purchases they make, the ads they see, and how they use their services—whether or not they have a Facebook account or are logged into Facebook. For example, a game developer could use our API to tell us what games a user plays, or a business could tell us about a purchase a user made in its store. We also receive information about users' online and offline actions and purchases from third-party data providers who have the rights to provide us with users' information.
- Partners receive users' data when users visit or use their services or through third parties they work with. We require each of these partners to have lawful rights to collect, use and share users' data before providing any data to us.

Question 96. Yes or no: Does Facebook give non-Facebook users a reasonable opportunity to learn what information has been collected about them by Facebook? If yes, please describe how.

Yes. If a person doesn't have a Facebook account but believes Facebook may have information about them, they can contact us to request a copy of their information. A contact form is available at <https://www.facebook.com/help/contact/180237885820953>.

However, Facebook does not create profiles about or track web or app browser behavior of non-users.

Question 97. During the April 10, 2018 joint committee hearing, you stated, "Every piece of content that you share on Facebook, you own and you have complete control over who sees

it and—and how you share it, and you can remove it at any time.” To corroborate that statement, you cited multiple mechanisms provided by Facebook that allow users to locate, edit, download, and delete information collected about them by Facebook.

Yes or no: Does Facebook offer non-Facebook users the same opportunities to control and edit any data collected about them by Facebook?

A user owns the information they share on Facebook. This means they decide what they share and who they share it with on Facebook, and they can change their mind. We believe everyone deserves good privacy controls. We require websites and apps who use our tools to tell users they’re collecting and sharing their information with us, and to get users’ permission to do so. However, non-Facebook users cannot post content on Facebook. Accordingly, there are not corresponding controls for non-Facebook users.

Facebook’s “Privacy Basics” on deleting posts states “Hiding lets you keep your post but no one else will be able to see it when they view your Timeline. Note that it might still show up in search results and other places on Facebook.”

How does an individual have “complete control” over their data if a post that has been hidden still shows up “in search results and other places on Facebook?”

A user can delete any post they have made. If they do so, it will not appear in search results and in other places on Facebook. The language you refer to appears in a feature that allows people to hide—not delete—content from their personal timeline. That is, a person can choose to delete a post that they have made from Facebook entirely, or they can choose to hide a post from their timeline even though it may be visible in other places on Facebook.

Does Facebook give users an opportunity delete their content or information from these “other places” or search results?

Yes. See above.

Does Facebook give non-users an opportunity to delete content containing or relating to them from these “other places” or search results?

Since this passage refers to content created by Facebook users and whether it’s visible on their timeline, this does not apply to non-users. See the Responses to the sub-questions above and below.

If a Facebook user deletes a post will it show up in search results and other places on Facebook? If so, please describe the other places on Facebook in which a deleted post may appear.

In general, when a user deletes their account, we delete things they have posted, such as their photos and status updates, and they won’t be able to recover that information later. (Information that others have shared about them isn’t part of their account and won’t be deleted.)

There are some limited exceptions to these policies: For instance, information can be accessed and preserved for an extended period when it is the subject of a legal request or obligation, governmental investigation, or investigations of possible violations of our terms or policies, or otherwise to prevent harm. We also retain information from accounts disabled for terms violations for at least a year to prevent repeat abuse or other term violations.

If a Facebook user deletes his account, will any of his data show up in search results and other places on Facebook? Will Facebook retain any of his data for any purpose? If so, please describe what data and for what purposes.

See Response above.

Question 98. Yes or no: does Facebook employ facial-recognition technology?

If so, does Facebook collect user data using facial-recognition technology?

If so, does Facebook collect data on individuals who are not registered Facebook users using facial-recognition technology?

If yes, does Facebook allow third-parties access to its facial-recognition technology or related information obtained as a result of the technology?

If yes, does Facebook allow government entities access to its facial recognition technology and/or the information obtained as a result of the technology?

To the extent that Facebook uses facial-recognition technology, what policies and procedures does Facebook have to safeguard information and data collected using that technology?

Does Facebook offer individuals, whether registered users or not, any opportunity to not be subject to facial-recognition technology or to have data collected using facial-recognition technology deleted?

Yes or no: Will Facebook commit to not using its facial-recognition technology to assemble data on individuals who have never consented to being part of Facebook?

Facebook uses facial recognition technology to provide people with products and features that enhance online experiences for Facebook users while giving them control over this technology. Facebook's facial recognition technology helps people tag their friends in photos; gives people an easier and faster way to privately share their photos with friends; helps people with visual impairments by generating descriptions of photos that people using screen readers can hear as they browse Facebook; lets people know when a photo or video of them has been uploaded to Facebook, even if they are not tagged; and helps prevent people from impersonating other Facebook users.

Facial recognition technology uses machine-learning algorithms to analyze the pixels in photos and videos in which a user is tagged, and the photo used by the person as his or her profile picture, and generates a unique number called a template. When a photo or video is

uploaded to Facebook, Facebook uses the template to attempt to identify someone by determining whether there are any faces in that content, and analyzing the portion of the image in which the face appears to compare it against certain Facebook users depending on the purpose for which facial recognition is being performed.

Facebook has not shared and does not have plans to share or make available to any third party its facial recognition templates. Moreover, these templates do not provide meaningful information on their own; they can be used to identify a person only in conjunction with Facebook's software. They could not be reverse-engineered to recreate someone's face.

Facebook designed its facial-recognition technology and the applications that use it with privacy considerations in mind and incorporated various safeguards and controls that protect both (1) users' ability to control the collection, use, and disclosure of their personal information, and (2) the security of that personal information.

Facebook gives users control over whether Facebook uses facial recognition to recognize them in photos and videos. That control is exercised through users' privacy settings. If a user chooses to turn facial recognition off, Facebook does not create a template for that person or deletes any template it has previously created. Facebook will then be unable to recognize that person in any photos or videos that are uploaded to the service. Facebook also deletes templates of people who delete their Facebook accounts. Additionally, Facebook does not maintain templates for users who have no photos tagged of themselves and do not have a profile photo that is capable of being used to generate a face signature or template (e.g., where a user has no profile photo, where a user's profile photo does not contain a human face, or where a user's profile photo contains multiple untagged faces).

We inform people about our use of facial-recognition technology through the Data Policy, Help Center, posts on Facebook, and direct user notifications. Facebook users are told that they can opt out of facial recognition at any time—in which case Facebook will delete their template and will no longer use facial recognition to identify them.

In creating facial recognition templates, Facebook uses only data that people have voluntarily provided to Facebook: the photos and videos that people have voluntarily uploaded to Facebook (including public profile pictures) and the tags people have applied to those photos and videos. Facebook does not use facial recognition to identify someone to a stranger.

Question 99. Yes or no: does Facebook collect users' audio or visual information for any reason whatsoever, or otherwise activate, monitor, or capture data from a microphone or camera from a user's phone without the user's contemporaneous knowledge and express, contemporaneous consent? If so, please list each and every instance under which Facebook does so.

No, Facebook does not engage in these practices or capture data from a microphone or camera without consent. Of course, we do allow people to take videos on their devices and share those on our platform.

Question 100. Will Facebook commit to not using its platform to gather such audio or visual information surreptitiously?

See Response to Question 99.

Question 101. During the April 11, 2018 House Energy and Commerce Hearing, you stated, “there may be specific things about how you use Facebook, even if you’re not logged in, that we keep track of, to make sure that people aren’t abusing the systems.” You further stated that “in general, we collect data on people who have not signed up for Facebook for security purposes.”

What categories of data does Facebook collect about registered users’ activity on websites and mobile applications other than Facebook?

What categories of data does Facebook collect about individuals who are not registered Facebook users and their activity on websites and mobile applications other than Facebook?

To the extent Facebook collects such data, does Facebook sell or provide this data to third parties?

To the extent Facebook collects such data, has Facebook gained consent from those individuals to collect and use their personal data?

To the extent Facebook gathers such data, what opportunity does Facebook provide to individuals not using Facebook to know, correct, or delete any information Facebook has gathered and retained about them?

Most of your answers to the questions you received on April 10, 2018, and likely most of the answers to these questions for the record, will depend on information that Facebook alone possesses.

Why is/are Facebook’s content-suggesting algorithm(s) secret?

Why are Facebook’s editorial decisions secret?

See Response to Question 93.

When the individual is a Facebook user, we are also able to use this information to personalize their experiences on Facebook, whether or not they are logged out, but we will not target ads to users relying on this information unless the user allows this in their privacy settings. We do not sell or share this information with third-parties.

Question 102. Numerous Americans receive all or a significant portion of their news from Facebook, which, in turn, suggests that news to them based on an algorithm that determines appropriate content based on criteria known only to Facebook.

To what extent will Facebook make public the criteria on which this algorithm relies?

To what extent will Facebook make public any changes that it makes to this or similar algorithms?

Facebook is a distribution platform that reflects the conversations already taking place in society. We want Facebook to be a place where people can discover more news, information, and perspectives, and we are working to build products that help.

As to the questions regarding ranking and algorithmic changes, see Response to Question 66.

Question 103. Facebook conducts numerous social experiments on its users, examining everything from the effects of Facebook on voter turnout to the effects of Facebook on the mood of its users.

Will Facebook commit to not experimenting on its users without express, informed consent in advance?

Will Facebook commit to making the results of any such experiments known publicly?

Will Facebook commit to not experimenting on human subjects at all?

Facebook does research in a variety of fields, from systems infrastructure to user experience to artificial intelligence to social science. We do this work to understand what we should build and how we should build it, with the goal of improving the products and services we make available each day. We're committed to doing research to make Facebook better, but we want to do it in the most responsible way.

In October 2014, we announced a new framework that covers both internal work and research that might be published:

- **Guidelines:** we've given researchers clearer guidelines. If proposed work is focused on studying particular groups or populations (such as people of a certain age) or if it relates to content that may be considered deeply personal (such as emotions) it will go through an enhanced review process before research can begin. The guidelines also require further review if the work involves a collaboration with someone in the academic community.
- **Review:** we've created a panel including our most senior subject-area researchers, along with people from our engineering, research, legal, privacy and policy teams, that will review projects falling within these guidelines. This is in addition to our existing privacy cross-functional review for products and research.
- **Training:** we've incorporated education on our research practices into Facebook's six-week training program, called bootcamp, that new engineers go through, as well as training for others doing research. We'll also include a section on research in the annual privacy and security training that is required of everyone at Facebook.
- **Research website:** our published academic research is now available at a single location (<https://research.facebook.com/>) and will be updated regularly.

We believe in research because it helps us build a better Facebook. Like most companies today, our products are built based on extensive research, experimentation and testing.

It's important to engage with the academic community and publish in peer-reviewed journals, to share technology inventions and because online services such as Facebook can help us understand more about how the world works. We want to do this research in a way that honors the trust users put in us by using Facebook every day. We will continue to learn and improve as we work toward this goal.

Question 104. What, if any, procedures does Facebook employ to verify the identities of individuals who purchase or employ data from Facebook?

Facebook does not sell people's information to anyone, and we never will. We also impose strict restrictions on how our partners can use and disclose the data we provide.

Our Data Policy makes clear the circumstances in which we work with third-party partners who help us provide and improve our Products or who use Facebook Business Tools to grow their businesses, which makes it possible to operate our companies and provide free services to people around the world.

Question 105. Research and reporting by NYU Professor of Marketing Scott Galloway suggests that, combined, Facebook and Google (parent company now known as Alphabet) are together worth approximately \$1.3 trillion. He concludes that this figure exceeds the world's top five advertising agencies (WPP, Omnicom, Publicis, IPG, and Dentsu) with five major media companies (Disney, Time Warner, 21st Century Fox, CBS, and Viacom) and still need to add five major communications companies (AT&T, Verizon, Comcast, Charter, and Dish) approach 90% of Facebook and Google's combined worth.

What business or product lines does Facebook consider itself to be in?

On what basis does Facebook make that determination?

Who does Facebook consider its major competitors in each of these business or product lines?

Of those business or product lines, what market share does Facebook believe that it has?

What other entities provide all of the services that Facebook does in one place or platform, if any?

What other entities provide any of the services that Facebook does?

What is the relevant product market for Facebook (the platform)?

What are the relevant product markets for each of Facebook's products?

What is the relevant geographic market for Facebook (the platform)?

What is the relevant geographic market for each of Facebook’s products?

Given these relevant geographic and product markets, what is Facebook’s market share in each distinct market in which it operates?

What procedures, tools, programs, or calculations does Facebook use to ascertain its market position relevant to its five largest competitors overall (if five exist)?

What procedures, tools, programs, or calculations does Facebook use to ascertain its market position relevant to its five largest competitors in each product market (if five exist)?

In Silicon Valley and around the world, new social apps are emerging all the time. The average American uses eight different apps to communicate with their friends and stay in touch with people. There is a lot of choice, innovation, and activity in this space, with new competitors arising all the time. Facebook’s top priority and core service is to build useful and engaging products that enable people to connect, discover and share through mobile devices and personal computers. Given its broad product offerings, Facebook faces numerous competitors, competing to attract, engage, and retain users, to attract and retain marketers, and to attract and retain developers who build compelling mobile and web applications. For instance, if a user wants to share a photo or video, they can choose between Facebook, DailyMotion, Snapchat, YouTube, Flickr, Twitter, Vimeo, Google Photos, and Pinterest, among many other services. Similarly, if a user is looking to message someone, just to name a few, there’s Apple’s iMessage, Telegram, Skype, Line, Viber, WeChat, Snapchat, and LinkedIn—as well as the traditional text messaging services their mobile phone carrier provides. Equally, companies also have more options than ever when it comes to advertising—from billboards, print and broadcast, to newer platforms like Facebook, Spotify, Twitter, Google, YouTube, Amazon or Snapchat. Facebook represents a small part (in fact, just 6%) of this \$650 billion global advertising ecosystem and much of that has been achieved by helping small businesses—many of whom could never have previously afforded newspaper or TV ads—to cost-effectively reach a wider audience.

Question 106. As you indicated in your testimony, Facebook’s business model relies on advertising to individuals, typically through tailored advertisements. This means that Facebook has monetized access to the information that those individuals have published on Facebook.

To Facebook’s best approximation, what is the total value of all user information that Facebook has acquired or to which Facebook has access?

Facebook generates substantially all of its revenue from selling advertising placements to third parties. Our total revenue and the percentage of which comes from third-party ads is below. This information is from our SEC filings.

2017: 40,653,000,000 (98% from third party ads)

2016: 27,638,000,000 (97% from third party ads)

2015: 17,928,000,000 (95% from third party ads)

2014: 12,466,000,000 (92% from third party ads)

2013: 7,872,000,000 (89% from third party ads)

2012: 5,089,000,000 (84% from third party ads)

2011: 3,711,000,000 (85% from third party ads)

2010: 1,974,000,000 (95% from third party ads)

2009: 777,000,000

2008: 272,000,000

How does Facebook categorize individual pieces of information for purposes of monetizing that information? (For example, Facebook acknowledges that if it is approached by a company selling ski equipment, it will target ads to individuals who have expressed an interest in skiing. We want to know in what ways Facebook organizes this information.)

As explained in our Data Policy, we collect three basic categories of data about people:

- (1) data about things people do and share (and who they connect with) on our services;
- (2) data about the devices people use to access our services; and
- (3) data we receive from partners, including the websites and apps that use our business tools.

Our Data Policy provides more detail about each of the three categories. 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 to edit or delete these interests.

We use data from each of the categories described above to obtain these interests and to personalize every aspect of our services, which is the core value we offer and the thing that makes Facebook services unique from other online experiences. This includes selecting and ranking relevant content, including ads, posts, Page recommendations, to cite but a few examples.

For example, we use the data people provide about their age and gender to help advertisers show ads based on those demographics but also to customize the pronouns on our site and deliver relevant experiences to those users.

We use data about things people do on Facebook, such as the Pages they like, to associate “interests” with their accounts, so we can rank posts relating to those interests higher in NewsFeed, for example, or enable advertisers to reach audiences—i.e., groups of people—that share those interests. For example, if a person has liked Pages about baseball, we might associate them with interests called “baseball” or “sports.”

We use data from devices (such as location data) to help advertisers reach people in particular areas. For example, if people have shared their device locations with Facebook or checked into a specific restaurant, we can show them organic posts from friends who have been in that location or we can show them ads from an advertiser that wants to promote its services in their area or from the restaurant.

We also help advertisers reach people who have given the advertiser their contact information or who have used the advertiser's website or app. For example, advertisers can send us a hashed list of email addresses of people they would like to reach on Facebook. 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).

Again, for people who are new to Facebook, we may have minimal data that we can use to personalize their experience, including their NewsFeed, their recommendations and the content (organic and sponsored) that they see. For people who have used our services for longer, we likely have more data, but the amount of data will depend on the nature of that use and how they have used our controls.

As noted above, in addition to general controls—such as Activity Log—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 as 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. People's use of these controls will, of course, affect the data we use to show them ads.

What types of advertisements does Facebook categorically prohibit?

Section 4 of our Advertising Policies list the types of ads that we categorically prohibit. These include ads that violate Community Standards, ads for illegal products and services, ads with adult content, ads that are misleading or false, ads that include profanity, and many more.

What external controls restrict how Facebook monetizes, sells, rents, or otherwise commercializes an individual's information? Please include (separately) any laws that Facebook views as applicable, any injunctions presently binding Facebook, any regulations directing how Facebook may monetize information, and any publicly available, independent audits of how Facebook monetizes information.

Facebook complies with all applicable laws. In addition, we adhere to the commitments set forth in our Data Policy, which describes how we collect and use data.

What internal controls restrict how Facebook monetizes, sells, rents, or otherwise commercializes an individual's information? Please include (separately) any internal

policies, statements of ethics or principles, directives, guidelines, or prohibitions that Facebook routinely applies in determining whether to use an individual’s personal information for commercial gain.

See Response to previous Question.

Question 107. When an individual chooses to “lock down” or otherwise publicly conceal his Facebook profile, does Facebook:

Continue to use that individual’s private information for commercial gain? (This includes aggregating data as well as targeting advertisements at that individual.)

Continue to retain that individual’s private information for its own archives or records?

When people post on Facebook—whether in a status update or by adding information to their profiles—the ability to input the information is generally accompanied by an audience selector. This audience selector allows the person to choose who will see that piece of information on Facebook—whether they want to make the information public, share it with friends, or keep it for “Only Me.” The tool remembers the audience a user shared with the last time they posted something and uses the same audience when the user shares again unless they change it. This tool appears in multiple places, such as privacy shortcuts and privacy settings. When a person makes a change to the audience selector tool in one place, the change updates the tool everywhere it appears. The audience selector also appears alongside things a user has already shared, so it’s clear who can see each post. After a person shares a post, they have the option to change who it is shared with.

The audience with which someone chooses to share their information is independent of whether we use that information to personalize the ads and other content we show them. Specifically, our Data Policy explains that we may use any information that people share on Facebook “to deliver our Products, including to personalize features and content (including your News Feed, Instagram Feed, Instagram Stories and ads).” However, people can use our Ad Preferences tool to see the list of interests that we use to personalize their advertising. This means that, for example, a person who is interested in cars can continue to share that interest with their friends but tell us not to assign them an interest in ads for ad targeting purposes.

Likewise, the audience of a post does not determine whether a post is retained. Someone can choose to share a post with “Only Me” (meaning that they don’t want anyone to see it but want to retain it in their Facebook account). They may also choose to delete the information entirely. When people choose to delete something they have shared on Facebook, we remove it from the site. In most cases, this information is permanently deleted from our servers; however, some things can only be deleted when a user permanently deletes their account.

Question 108. What are Facebook’s total advertising revenues for each of the calendar years 2001 to 2018?

Our total revenue and the percentage of which comes from third-party ads is below. This information is from our SEC filings.

2017: 40,653,000,000 (98% from third party ads)

2016: 27,638,000,000 (97% from third party ads)

2015: 17,928,000,000 (95% from third party ads)

2014: 12,466,000,000 (92% from third party ads)

2013: 7,872,000,000 (89% from third party ads)

2012: 5,089,000,000 (84% from third party ads)

2011: 3,711,000,000 (85% from third party ads)

2010: 1,974,000,000 (95% from third party ads)

2009: 777,000,000

2008: 272,000,000

What are Facebook's online advertising revenues for each of the calendar years 2001 to 2018?

What are Facebook's five largest competitors for online advertising in each year from 2001 to 2018?

What were each of those competitors' advertising revenues through each of those years?

How many of Facebook's executive staff previously worked at each of those entities?

We expect that our competitors make their numbers available in their SEC filings. And, like many industries across the private sector, many people may work in multiple technology companies throughout the course of their careers.

Question 109. Regardless of place of incorporation, does Facebook consider itself an American company?

Yes, we're an American-based company where ninety percent of our community are outside the US.

Question 110. When Facebook makes policy decisions, are American citizens the company's top priority? If not, what is the company's top priority when it comes to policy decisions?

We are proud to be a US-based company that serves billions of people around the world. While the majority of our employees are located here in the United States, more than 80% of the people who use Facebook are outside this country. We consider the needs of all of our users when making policy decisions. Of course, with headquarters in the US and Ireland, we have particularly strong relationships with policy makers in those regions. We regularly engage with

policy makers around the world, however, and work to take account of regional policy concerns as we build our products and policies for a global user base.

Question 111. Facebook, WhatsApp, and Instagram have all reportedly been blocked or partially blocked from the People's Republic of China (PRC) since 2009.

Facebook, WhatsApp and Instagram are available in Hong Kong and Macau. Facebook and Instagram are blocked in Mainland China. However, these can be accessed by people in Mainland China who employ VPNs. WhatsApp is typically available in Mainland China although we notice availability is often restricted around important events.

Please describe the extent to which these services may be accessed from within the territory of the PRC, including Hong Kong and Macau, and describing in detail any geographical limits or limits on the available content.

Facebook, WhatsApp, and Instagram are available in Hong Kong and Macau. Facebook and Instagram are blocked in Mainland China. However, these can be accessed by people in Mainland China who employ VPNs. WhatsApp is typically available in Mainland China although we notice availability is often restricted around important events.

On what basis does Facebook evaluate whether to honor a foreign government's request to block specific content?

When something on Facebook or Instagram is reported to us as violating local law, but doesn't go against our Community Standards, we may restrict the content's availability only in the country where it is alleged to be illegal after careful legal review. We receive reports from governments and courts, as well from non-government entities such as members of the Facebook community and NGOs.

How does Facebook determine whether to honor a foreign government's request to block specific content or users?

See Response to previous question.

Listed by country, what percentage of requests to block specific content (or users) from foreign governments does Facebook honor in whole or part?

This information is available here: <https://transparency.facebook.com/content-restrictions>.

How does Facebook determine whether to honor the U.S. government's request to block specific content or users?

Our Transparency Report contains data on restrictions we place on content that does not violate community standards but that is alleged to violate local law. We do not have any such reports for the United States.

What percentage of requests to block specific content (or users) from the U.S. government does Facebook honor in whole or part?

See Response to previous question.

Question 112. Yes or no: Has Facebook made any alterations, modifications, or changes to the encryption security of WhatsApp in response to or as a result of the PRC government or any of its agencies or in order to comply with PRC law?

No.

If so, what changes has Facebook made to the encryption security?

Does Facebook program in “back doors” or other mechanisms to decrypt or otherwise decode encrypted information at a government’s request?

No.

If so, under what circumstances does Facebook decrypt such data?

If so, on what platforms does Facebook have such protocols?

Does Facebook make WhatsApp or Facebook information available to the PRC government on a searchable basis?

No.

Question 113. Since 2014, the PRC government has held a World Internet Conference. Charles Smith, the co-founder of the non-profit censorship monitoring website GreatFire, described foreign guests of the Conference as “complicit actors in the Chinese censorship regime [that] are lending legitimacy to Lu Wei, the Cyberspace Administration of China and their heavy-handed approach to Internet governance. They are, in effect, helping to put all Chinese who stand for their constitutional right to free speech behind bars.”

How many Facebook employees have attended the PRC’s World Internet Conference?

There have been four World Internet Conferences. Several Facebook employees have attended one or more of these four conferences.

Have any Facebook employees ever participated on any panels or advisory committees that are held or have been established by the World Internet Conference?

If so, please list the employees and the panels or high-level advisory committees they have participated on.

One Facebook representative, Vaughan Smith, has participated in World Internet Conference panels and keynotes alongside representatives of other leading US technology companies, for example Tim Cook and Sundar Pichai. No employees participated in advisory

committees. Mr. Smith has provided keynotes on AI, innovation and how Facebook is building the knowledge economy.

Has Facebook assisted other countries in designing regimes to monitor or censor Facebook content? If so, which countries, and under what circumstances? Please describe each.

When something on Facebook or Instagram is reported to us as violating local law, but doesn't go against our Community Standards, we may restrict the content's availability only in the country where it is alleged to be illegal after careful legal review. We receive reports from governments and courts, as well from non-government entities such as members of the Facebook community and NGOs. This information is available here:

<https://transparency.facebook.com/content-restrictions>.

Government criticism does not violate our community standards, and we do not evaluate or categorize accounts based on whether they engage in government criticism.

See also Response to Question 111(c).

Has Facebook ever provided any financial support to the World Internet Conference? If yes, please provide and itemize all financial support that has been provided to the World Internet Conference.

Facebook has not paid to participate in the World Internet Conference. In 2016 we paid \$10,000 to rent exhibit space at the event to showcase Oculus VR which is manufactured in China.

Question 114. Has Facebook ever temporarily shut down or limited access to Facebook, WhatsApp, or Instagram within a country or a specific geographic area, at the request of a foreign government or agency, including but not limited to, the PRC, the Islamic Republic of Iran, Syria, the Russian Federation, and Turkey?

If so, please describe each instance Facebook has complied with a foreign government's request to censor content or users, the requesting government, the provided justification for the government request, and a description of the content requested to be removed.

Please describe what if any policies Facebook has in place governing Facebook's responses to government censorship requests.

We do not block access to Facebook products and services in areas where they are otherwise generally available on the basis of specific government requests. We may independently limit access to certain functionality—such as peer-to-peer payments or facial recognition—in some jurisdictions based on legal and regulatory requirements.

In some instances, we may receive requests from governments or other parties to remove content that does not violate our Community Standards but is alleged to contravene local law. When we receive such requests, we conduct a careful review to confirm whether the report is legally valid and is consistent with international norms, as well as assess the impact of our response on the availability of other speech. When we comply with a request, we restrict the

content only within the relevant jurisdiction. We publish details of content restrictions made pursuant to local law, as well as details of our process for handling these requests, in our Transparency Report (<https://transparency.facebook.com/content-restrictions>).

Questions from Ranking Member Nelson

Question 1. While the primary focus of the April 10 hearing was on Cambridge Analytica and Facebook’s privacy and data security policies, concerns were heard about many other issues from Members on both sides of the aisle. Within this context, please detail specific steps that Facebook is taking to address: (1) “fake news”, (2) foreign government interference in American elections, (3) illegal sex trafficking, and (4) copyright infringement of digital content.

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 with 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.

Foreign Interference: 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 prevent bad actors from using misinformation to undermine the democratic process. 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 the 10 most important changes we have made:

- **Ads transparency.** Advertising should be transparent: users should be able to see all the ads an advertiser is currently running on Facebook, Instagram, and Messenger. We are taking steps to help users assess the content they see on Facebook. For example, for ads with political content, we’ve created an archive that will hold ads with political content for seven years—including for information about ad impressions and spend, as well as demographic data such as age, gender, and location. People in Canada and Ireland can already see all the ads that a Page is running on Facebook—and we’re launching this globally in June. Further, advertisers will now need to confirm their ID and location before being able to run any ads with political content in the US. All ads with political content will also clearly state who paid for them. 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.

- **Verification and labeling.** We are working hard to regain the trust of our community. Success would consist of minimizing or eliminating abuse of our platform and keeping our community safe. We have a number of specific goals that we will use to measure our progress in these efforts. First, we are increasing the number of people working on safety and security at Facebook, to 20,000. We have significantly expanded the number of people who work specifically on election integrity, including people who investigate this specific kind of abuse by foreign actors. Those specialists find and remove more of these actors. Second, we work to improve threat intelligence sharing across our industry, including, we hope, by having other companies join us in formalizing these efforts. This is a fight against sophisticated actors, and our entire industry needs to work together to respond quickly and effectively. Third, we are bringing greater transparency to election ads on Facebook by requiring more disclosure from people who want to run election ads about who is paying for the ads and by making it possible to see all of the ads that an advertiser is running, regardless of the targeting. We believe that these efforts will help to educate our community and to arm users, media, civil society, and the government with information that will make it easier to identify more sophisticated abuse to us and to law enforcement.
- **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.** We have gotten increasingly better at finding and disabling fake accounts. We're now at the point that we block millions of fake accounts each day at the point of creation before they do 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.** (see above).
- **Significant investments in security.** We're doubling the number of people working on safety and security from 10,000 last year to over 20,000 this year. We expect these investments to impact our profitability. But the safety of people using Facebook needs to come before profit.
- **Industry collaboration.** In April, we joined 34 global tech and security companies in signing a TechAccord pact to help improve security for everyone.

- **Intelligence sharing with government.** 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.** In recent months, we've started to deploy 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 plan to continue 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. We are taking steps to enhance trust in the authenticity of activity on our platform, including increasing ads transparency, implementing a more robust ads review process, imposing tighter content restrictions, and exploring how to add additional authenticity safeguards. We also have improved information sharing about these issues among our industry partners.

Copyright: Facebook takes intellectual property rights seriously and believes they are important to promoting expression, creativity, and innovation in our community. Facebook's Terms of Service do not allow people to post content that violates someone else's intellectual property rights, including copyright and trademark. We publish information about the intellectual property reports we receive in our bi-annual Transparency Report, which can be accessed at <https://transparency.facebook.com/>.

Sex trafficking: Facebook is committed to making our platform a safe place, especially for individuals who may be vulnerable. We have a long history of working successfully with governments to address a wide variety of threats to our platform, including child exploitation.

When we learn of a situation involving physical abuse, child exploitation, or an imminent threat of harm to a person, we immediately report the situation to first responders or the National Center for Missing and Exploited Children (NCMEC). Further, as part of official investigations, government officials sometimes request data about people who use Facebook. We have processes in place to handle these government requests, and we disclose account records in accordance with our terms of service and applicable law. We also have a global team that strives to respond within minutes to emergency requests from law enforcement.

Our relationship with NCMEC also extends to an effort that we launched in 2015 to send AMBER Alerts to the Facebook community to help find missing children. When police determine that a case qualifies for an AMBER Alert, the alert is issued by the NCMEC and distributed through the Facebook system with any available information, including a photograph

of the missing child, a license plate number, and the names and descriptions of the child and suspected abductor. Law enforcement determines the range of the target area for each alert. We know the chances of finding a missing child increase when more people are on the lookout, especially in the critical first hours. Our goal is to help get these alerts out quickly to the people who are in the best position to help, and a number of missing children have been found through AMBER Alerts on Facebook.

Further, we work tirelessly to identify and report child exploitation images (CEI) to appropriate authorities. We identify CEI through a combination of automated and manual review. On the automated review side, we use image hashing to identify known CEI. On the manual review side, we provide in-depth training to content reviewers on how to identify possible CEI. Confirmed CEI is reported to the NCMEC, which then forwards this information to appropriate authorities. When we report content to the NCMEC, we preserve account information in accordance with applicable law, which can help further law enforcement investigations. We also reach out to law enforcement authorities in serious cases to ensure that our reports are received and acted upon.

Question 2. Some commentators worry that the Internet is dominated by a few large platforms with little competition or accountability. Facebook is obviously considered to be one of those key, dominant platforms.

A. Please comment on how American laws should hold large Internet platforms accountable when things go wrong?

B. What is Facebook's legal and ethical responsibility as an Internet platform with billions of global users?

Our mission is to give people the power to build community and bring the world closer together—a mission that is inherently global and enhanced by a global scope. As the internet becomes more important in people's lives, the real question is about the right set of regulations that should apply to all internet services, regardless of size. Across the board, we have a responsibility to not just build tools, but to make sure that they're used in ways that are positive for our users. It will take some time to work through all the changes we need to make across the company, but Facebook is committed to getting this right.

Question 3. If large Internet platforms compromise consumer privacy and/or facilitate the theft of original content, what should be the federal government's response? What should be the obligations of the platforms?

We take intellectual property rights seriously at Facebook and work closely with the motion picture industries and other rights holders worldwide to help them protect their copyrights and other IP. Our measures target potential piracy across our products, including Facebook Live, and continue to be enhanced and expanded. These include a global notice-and-takedown program, a comprehensive repeat infringer policy, integration with the content recognition service Audible Magic, and our proprietary video- and audio-matching technology called Rights Manager. More information about these measures can be found in our Intellectual Property Help Center, Transparency Report, and Rights Manager website.

Question 4. In general, as reflected in the General Data Protection Regulation (GDPR), the European Union (EU) is considered to require stronger data and privacy protections than the United States. According to press reports, Facebook will be moving 1.5 billion users outside of the scope of the EU's GDPR. Please explicitly lay out how Facebook's compliance with the GDPR will affect all Facebook users, including American users. That is, to what extent will the GDPR's requirements and protections extend to Americans and users outside Europe?

The press reports referred to in this question pertain to the legal entity with which Facebook users contract when they use the service, which changed in some jurisdictions as a part of the most recent updates to our Terms of Service and Data Policy. This change did not impact people who live in the United States, who contract with Facebook, Inc. under both our new and old policies.

The substantive protections in our user agreements offered by Facebook Ireland 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 (DPO) 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.

In any case, the controls and settings that Facebook is enabling as part of GDPR are available to people around the world, including settings for controlling our use of face recognition on Facebook and for controlling our ability to use data we collect off Facebook Company Products to target ads. We recently began providing 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 consent to our updated terms. We provide the same tools for access, rectification, erasure, data portability and others to people in the US and rest of world that we provide in Europe, and many of those tools (like our Download Your Information tool, ad preferences tool, and Activity Log) have been available globally for many years.

Questions from Senator Cantwell

SUMMARY

I understand that last week you announced your support for legislation that would regulate political ads on internet platforms. By your own report, Facebook has removed 70 Facebook accounts, 138 Facebook Pages, and 65 Instagram accounts run by the Russian government-connected troll farm and election interference group known as the Internet Research Agency.

I want to explore the distinction between paid political ads and the troll and bot activity deployed by Russia that was designed to meddle with and influence US elections.

Question 1.

What tools do we have to address this going forward? If we pass the Honest Ads Act, won't we still have a problem with bots and trolls that aren't running traditional paid "political ads"?

We have always believed that Facebook is a place for authentic dialogue and that the best way to ensure authenticity is to require people to use the names they are known by. Fake accounts undermine this objective and are closely related to the creation and spread of inauthentic communication such as spam and disinformation. We also prohibit the use of automated means to access our platform. We rely on both automated and manual review in our efforts to effectively detect and deactivate fake accounts, including bots, and we are now taking steps to strengthen both. For example, we continually update our technical systems to identify, checkpoint, and remove inauthentic accounts. We block millions of attempts to register fake accounts every day. These systems examine thousands of detailed account attributes and prioritize signals that are more difficult for bad actors to disguise.

Question 2.

Do we need a new definition of paid advertising or political expenditures that reaches bots and troll activity that are backed by foreign national interests?

We're committed to addressing this, and we have a number of efforts underway. Facebook has generally dealt with bots and troll activity via its Authenticity policy. Already, we build and update technical systems every day to better identify and remove inauthentic accounts, which also helps reduce the distribution of material that can be spread by accounts that violate our policies. Each day, we block millions of fake accounts at registration. Our systems examine thousands of account attributes and focus on detecting behaviors that are very difficult for bad actors to fake, including their connections to others on our platform. By constantly improving our techniques, we also aim to reduce the incentives for bad actors who rely on distribution to make their efforts worthwhile.

For example, the Internet Research Agency, based in St. Petersburg, is a "troll farm" and generally thought to be aligned with the Russian government. Facebook has determined that Internet Research Agency users violated Facebook's authenticity policy and has been working to

remove them from the platform. This has resulted in the removal of numerous Facebook and Instagram accounts, as well as the content connected with those accounts. Facebook has found that many trolls are motivated by financial incentives and is taking steps to disrupt those incentives to discourage the behavior. While working to limit the impact of bots and trolls, Facebook is striving to strike the right balance between enabling free expression and ensuring that its platform is safe. Facebook's policies are aimed at encouraging expression and respectful dialogue.

Question 3.

Would you commit to working on a way to address the bots and troll problem in a way that does not compromise free speech?

Yes, see Response to Question 2.

Question 4.

In your testimony you talked about your use of artificial intelligence to combat hate speech, bots, and trolls. What do you feel is the correct regulatory or other approach Congress should take to address artificial intelligence or other emerging technologies?

Artificial Intelligence (AI) is a very promising technology that has many applications. Fairness, transparency and accountability should guide its development. Presently, AI systems make decisions in ways that people don't really understand. Thus, society needs to invest further in developing AI systems which are more transparent. Facebook has AI teams working on developing the philosophical, as well as technical, foundations for this work. We discussed our AI ethics work during the keynote of our recent developer's conference (at minute 47): <https://www.facebook.com/FacebookforDevelopers/videos/10155609688618553/>.

Question 5.

How does Facebook plan to address the leveraging of its social engineering tools developed to optimize advertising revenue by state sponsored actors and geopolitical forces that seek to influence democratic elections and political outcomes?

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 prevent bad actors from using misinformation to undermine the democratic process.

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 the 10 most important changes we have made:

1. Ads 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 content, we've created an archive that will hold ads with political content for seven years—including information about ad impressions and spend, as well as demographic data such

as age, gender and location. People in Canada and Ireland can already see all the ads that a Page is running on Facebook—and we’re launching this globally in June.

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

3. 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.

4. 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.

5. 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 with 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.

A key focus is working to disrupt the economics of fake news. For example, preventing the creation of fake accounts that spread it, banning sites that engage in this behavior from using our ad products, and demoting articles found to be false by fact checkers in News Feed—causing it to lose 80% of its traffic. We now work with independent fact checkers in the US, France, Germany, Ireland, the Netherlands, Italy, Mexico, Colombia, India, Indonesia and the Philippines with plans to scale to more countries in the coming months.

6. Significant investments in security. We’re doubling the number of people working on safety and security from 10,000 last year to over 20,000 this year. We expect these investments to impact our profitability. But the safety of people using Facebook needs to come before profit.

7. Industry collaboration. Recently, we joined 34 global tech and security companies in signing a TechAccord pact to help improve security for everyone.

8. 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.

9. Tracking 40+ elections. In recent months, we've started to deploy 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 plan to continue 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.

10. 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.

Question 6.

How should Congress address the leveraging of social engineering tools developed to optimize advertising revenue on technology platforms, by state sponsored actors and geopolitical forces that seek to influence democratic elections and political outcomes?

From its earliest days, Facebook has always been focused on security. These efforts are continuous and involve regular contact with law enforcement authorities in the US and around the world. Elections are particularly sensitive events for Facebook's security operations, and as the role of Facebook's service plays in promoting political dialogue and debate has grown, so has the attention of its security team. To address these concerns, Facebook is taking steps to enhance trust in the authenticity of activity on its platform, including increasing ads transparency, implementing a more robust ads review process, imposing tighter content restrictions, and exploring how to add additional authenticity safeguards. We welcome a dialog with government about how to address these societal issues.

Questions from Senator Blumenthal

Facebook's Download Your Information Tool

During the hearing, I asked not only whether Facebook users should be able to access their information, but specifically whether it would provide its users “all of the information that you collect as a result of purchases from data brokers, as well as tracking them?” You affirmatively stated that Facebook has a “Download Your Information (DYI) tool that allows people to see and to take out all of the information that they have put into Facebook or that Facebook knows about them.”

However, in a March 7, 2018 correspondence provided to the U.K. Parliament regarding Paul-Olivier Dehaye's legal request for personal data, Facebook's Privacy Operations Team acknowledged that the DYI tool does not provide records stored in its 'Hive' database. This answer appears to confirm that the Facebook 'Pixel' web tracking system and other records are stored and combined with profile information, but not provided to users. Since then, *WIRED* magazine and academic researchers have noted the omission from the DYI tool of other pieces of data that Facebook is known to collect.

Question 1. What specific pieces of data does Facebook collect that are not provided through the DYI tool? Please provide exact labels and descriptions of the types of data and its source, rather than broad categories or intent, including but not limited to web tracking data, location history, ad interactions and advertiser targeting data, third party applications, and derived inferences.

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, will 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, delete this information from their account, and turn off Facebook's ability to store it associated with their account going forward.

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.

Facebook's Web Tracking

While users can more readily understand the types of data that Facebook collects directly from them, Facebook's data collection practices regarding non-users and from other sources are opaque. For example, Facebook collects data from its social plugins, Pixel, and other similar properties ("web tracking data") that provide a deep understanding about an individual's web browsing habits.

Question 2. Would an employee with appropriate technical permissions to the Hive database be able to generate a list of websites viewed by a Facebook user, where such websites contained a Facebook tracking property?

We have strict policy controls and technical restrictions so employees only access the data they need to do their jobs—for example to fix bugs, manage customer support issues or respond to valid legal requests. Employees who abuse these controls will be fired. Further information is available in our Cookies Policy, available at <http://facebook.com/ads/about>.

Question 3. Is web tracking data used for inferring an individual's interests or other characteristics? Are those inferences used in advertising?

Yes, but only for Facebook users. We do not use web browsing data to show ads to non-users or otherwise store profiles about non-users. Our goal is to show people content (including advertising) that is relevant to their interests. We use information people have provided on Facebook—such as things they've liked or posts they've engaged with—to help determine what people will be interested in. Like most online advertising companies, we also inform our judgments about what ads to show based on apps and websites that people use off of Facebook. People can turn off our use of web browser data and other data from third-party partners to show them ads through a control in Ads Preferences. They can also customize their advertising experience by removing interests that they do not want to inform the Facebook ads they see. In addition, a person's browser or device may offer settings that allow users to choose whether browser cookies are set and to delete them.

Question 4. Does Facebook provide users and non-users with the ability to disable the collection (not merely the use) of web tracking? Does Facebook allow users to delete this data without requiring the deletion of their accounts?

When people visit apps or websites that feature our technologies—like the Facebook Like or Comment button—our servers automatically log (i) standard browser or app records of the fact that a particular device or user visited the website or app (this connection to Facebook's servers occurs automatically when a person visits a website or app that contains our technologies, such as a Like button, and is an inherent function of Internet design); and (ii) any additional information the publisher of the app or website chooses to share with Facebook about the person's activities on that site (such as the fact that a purchase was made on the site). This is a standard feature of the Internet, and most websites and apps share this same information with multiple different third-parties whenever people visit their website or app. For example, the Senate Commerce Committee's website shares information with Google and its affiliate DoubleClick and with the analytics company Webtrends. This means that, when a person visits the Committee's website, it sends browser information about their visit to each one of those third

parties. More information about how this works is available at <https://newsroom.fb.com/news/2018/04/data-off-facebook/>.

When the person visiting a website featuring Facebook's tools is not a registered Facebook user, Facebook does not have information identifying that individual, and it does not create profiles for this individual.

We use the browser and app logs that apps and websites send to us—described above—in the following ways for non-Facebook users. First, these logs are critical to protecting the security of Facebook and to detecting or preventing fake account access. For example, if a browser has visited hundreds of sites in the last five minutes, that's a sign the device might be a bot, which would be an important signal of a potentially inauthentic account if that browser then attempted to register for an account. Second, we aggregate those logs to provide summaries and insights to websites and apps about how many people visit or use their product or use specific features like our Like button—but without providing any information about a specific person. We do not create profiles for non-Facebook users, nor do we use browser and app logs for non-Facebook users to show targeted ads from our advertisers to them or otherwise seek to personalize the content they see. However, we may take the opportunity to show a general ad that is unrelated to the attributes of the person or an ad encouraging the non-user to sign up for Facebook.

When the individual is a Facebook user, we are also able to use this information to personalize their experiences on Facebook, whether or not they are logged out, but we will not target ads to users relying on this information unless the user allows this in their privacy settings. We do not sell or share this information with third-parties.

We recently announced plans to build on this by introducing Clear History, a new feature that will enable users to see the websites and apps that send us information when they use them, delete this information from their accounts, and turn off our ability to store it associated with their accounts going forward.

Question 5. One academic study from 2015 raised concerns about the privacy risks of web tracking data collected from health-related web pages, including an example of Facebook collecting information from the inclusion of a Facebook Like button on the CDC's page about HIV. Does Facebook impose any limitation on itself regarding the collection and use (including inferences) of web tracking data collected from health-related pages or any other themes of websites?

Websites and apps choose whether they use Facebook services to make their content and ads more engaging and relevant and whether they share browser data or other information with Facebook or other companies when people visit their sites. These services include:

- Social plugins, such as our Like and Share buttons, which make other sites more social and help people share content on Facebook;
- Facebook Login, which lets people use their Facebook account to log into another website or app;

- Facebook Analytics, which helps websites and apps better understand how people use their services; and
- Facebook ads and measurement tools, which enable websites and apps to show ads from Facebook advertisers, to run their own ads on Facebook or elsewhere, and to understand the effectiveness of their ads.

Many companies offer these types of services and, like Facebook, they also get information from the apps and sites that use them. Twitter, Pinterest, and LinkedIn all have similar Like and Share buttons to help people share things on their services. Google has a popular analytics service. And Amazon, Google, and Twitter all offer login features. These companies—and many others—also offer advertising services. In fact, most websites and apps send the same information to multiple companies each time users visit them.

For example, when a user visits a website, their browser (for example Chrome, Safari or Firefox) sends a request to the site’s server. The browser shares a user’s IP address, so the website knows where on the internet to send the site content. The website also gets information about the browser and operating system (for example Android or Windows) they’re using because not all browsers and devices support the same features. It also gets cookies, which are identifiers that websites use to know if a user has visited before.

A website typically sends two things back to a user’s browser: first, content from that site; and second, instructions for the browser to send the user’s request to the other companies providing content or services on the site. So, when a website uses one of our services, our users’ browsers send the same kinds of information to Facebook as the website receives. We also get information about which website or app our users are using, which is necessary to know when to provide our tools.

Our policies include a range of restrictions on the use of these tools for health-related advertising. For example, we do not allow ads that discriminate based on disability, medical or genetic condition. Ads also may not contain content that directly or indirectly asserts or implies a person’s disability, medical condition (including physical or mental health), or certain other traits. And ads generally may not request health information, including physical health, mental health, medical treatments, medical conditions, or disabilities. And we prohibit anyone from using our pixel to send us data that includes health, financial information, or other categories of sensitive information.

In addition, we also enable ad targeting options—called “interests” and “behaviors”—that are based on people’s activities on Facebook, and when, where, and how they connect to the Internet (such as the kind of device they use and their mobile carrier). These options do not reflect people’s personal characteristics, but we still take precautions to limit the potential for advertisers to misuse them. For example, we do not create interest or behavior segments that suggest the people in the segment are members of sensitive groups such as people who have certain medical conditions.

Question 6. What changes, if any, is Facebook making to limit the amount of data that Facebook itself collects about users and non-users?

As explained in our Data Policy, we collect three basic categories of data about people: (1) data about things people do and share (and who they connect with) on our services, (2) data about the devices people use to access our services, and (3) data we receive from partners, including the websites and apps that use our business tools. Our Data Policy provides more detail about each of the three categories.

We use this information for a variety of purposes, including to provide, personalize, and improve our products, provide measurement, analytics, and other business services, promote safety and security, to communicate with people who use our services, and to research and innovate to promote the social good. We provide more information in our Data Policy about these uses as well.

Our policies limit our retention of the data that we receive in several ways. Specifically, we store data until it is no longer necessary to provide our services and Facebook products, or until a person's account is deleted—whichever comes first. This is a case-by-case determination that depends on things like the nature of the data, why it is collected and processed, and relevant legal or operational retention needs. For example, when a user searches for something on Facebook, they can access and delete that query from within their search history at any time, but the log of that search is deleted after 6 months. If they submit a copy of their government-issued ID for account verification purposes, we delete that copy 30 days after submission. If a user posts something on their Facebook profile, then that information would be retained until they delete it or until they delete their account.

We also have other policies that are designed to limit our retention of other types of information about people. For example, if a user visits a site with the “Like” button or another social plugin, we receive cookie information that we use to help show them a personalized experience on that site as well as Facebook, to help maintain and improve our service, and to protect both the user and Facebook from malicious activity. We delete or anonymize it within 90 days.

In general, when a user deletes their account, we delete things they have posted, such as their photos and status updates, and they won't be able to recover that information later. (Information that others have shared about them isn't part of their account and won't be deleted.)

There are some limited exceptions to these policies: For instance, information can be accessed and preserved for an extended period when it is the subject of a legal request or obligation, governmental investigation, or investigations of possible violations of our terms or policies, or otherwise to prevent harm. We also retain information from accounts disabled for terms violations for at least a year to prevent repeat abuse or other term violations.

We collect very little data about non-users (unless they choose to communicate directly with us) and do not create profiles or track browsing history for people who are not registered users of Facebook, for example.

Particularly in the past few months, we've realized that we need to take a broader view of our responsibility to our community. Part of that effort is continuing our ongoing efforts to identify ways that we can improve our privacy practices. This includes restricting the way that developers can get information from Facebook and announcing plans to build Clear History, a new feature that will enable users to see the websites and apps that send us information when

they use them, delete this information from their accounts, and turn off our ability to store it associated with their accounts going forward.

Onavo Protect

When Facebook bought a VPN service, Onavo Protect, the purchase was portrayed as a way for your company to build more efficient mobile products. Since 2016, you have encouraged users to install the Onavo application as a way to “keep you and your data safe,” although it does not brand itself as a Facebook product. Onavo is a particularly sensitive product since it provides your company access to all of the Internet traffic being generated by the device. Wall Street Journal and other publications have reported that Facebook has used the data captured from the Onavo for market analytics on competitive services.

Question 7. Does Facebook use traffic information collected from Onavo to monitor the adoption or popularity of non-Facebook applications?

When people first install the iOS version of the Onavo Protect app, we explain that Onavo uses a VPN that “helps keep you and your data safe by understanding when you visit potentially malicious or harmful websites and giving you a warning.” In addition, the first screen that a person sees when installing the app explains, under a heading that reads “Data Analysis”:

“When you use our VPN, we collect the info that is sent to, and received from, your mobile device. This includes information about: your device and its location, apps installed on your device and how you use those apps, the websites you visit, and the amount of data use.

This helps us improve and operate the Onavo service by analyzing your use of websites, apps and data. Because we’re a part of Facebook, we also use this info to improve Facebook products and services, gain insights into the products and services people value, and build better experiences.”

People must tap a button marked “Accept & Continue” after seeing this information in a full-screen interstitial before they can use the app.

The Android version of the Onavo Protect app offers data management features (e.g., the ability to block apps from using background data) that do not require users to enable the app’s VPN.

For both versions of the app, we communicate repeatedly and up front—in the App Store description, in Onavo’s Privacy Policy, and in-line at the time the user first opens the app after downloading it—that Onavo is part of Facebook and what that means for how Onavo Protect handles data in other ways.

More broadly, websites and apps have used market research services for years. We use Onavo, App Annie, comScore, and publicly available tools to help us understand the market and improve all our services. When people download Onavo to manage their data usage and help secure their connection, we are clear about the information we collect and how it is used. Like other VPNs, when the Onavo VPN is enabled, Onavo Protect helps create a secure connection, including when people are on public Wi-Fi. As part of this process, Onavo receives their mobile

data traffic. This helps us improve and operate the Onavo service. Because we're part of Facebook, we also use this information to improve Facebook products and services. We let people know about this activity, and other ways that Onavo uses, analyzes, and shares data (for example, the apps installed on users' devices) in the App Store descriptions, and when they first open the app after downloading it.

Facebook does not use Onavo data for Facebook product uses, nor does it append any Onavo data or data about individuals' app usage to Facebook accounts.

Question 8. Has Facebook ever used the Onavo data in decisions to purchase another company or develop a product to compete against another company?

See Response to Question 7.

Question 9. Does Facebook associate Onavo traffic information with profile data from its social networking sites, including for analytic purposes?

No. See Response to Question 7.

Facebook and Academic Research

Facebook's users place a significant amount of trust in the company to keep its data safe and protect the integrity of the platform. While Facebook has now developed a well-regarded ethical review processes and it is commendable that the company has supported academic research, any process is fallible and at least one of its experiments on "emotional contagion" was highly criticized by the academic community. One of the researchers behind the Cambridge Analytica application, Dr. Aleksandr Kogan, had frequently collaborated with Facebook on social science research based on its data, including a paper where Facebook provided data on every friendship formed in 2011 in every country in the world at the national aggregate level. Facebook users almost certainly are unaware that their data is used for scientific research by outside researchers nor do they have a credible understanding of the accountability of these relationships.

Question 10. Has Facebook ever provided any third party researcher with direct access to non-anonymized user data?

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.

Question 11. Do users have the ability to opt out of such experiments?

No, 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.

Question 12. Has a researcher ever been found to have misused access to the non-anonymized user data? Please describe any such incidents.

We are investigating all apps that, like Aleksandr Kogan's, had access to large amounts of information before we changed our platform in 2014 to reduce data access. The investigation process is in full swing, and it has two phases. First, a comprehensive review to identify every app that had access to this amount of Facebook data. And second, where we have concerns, we will conduct interviews, make requests for information (RFI)—which ask a series of detailed questions about the app and the data it has access to—and perform audits that may include on-site inspections. We have large teams of internal and external experts working hard to investigate these apps as quickly as possible. To date thousands of apps have been investigated and around 200 (from a handful of developers: Kogan, AIQ, Cube You, the Cambridge Psychometrics Center, and myPersonality) have been suspended—pending a thorough investigation into whether they did in fact misuse any data. Where we find evidence that these or other apps did misuse data, we will ban them and notify people whose data was shared with these apps. Additionally, we have suspended an additional 14 apps, which were installed by around one thousand people. They were all created after 2014, after we made changes to more tightly restrict our platform APIs to prevent abuse. However, these apps appear to be linked to AIQ, which was affiliated with Cambridge Analytica. So, we have suspended them while we investigate further. Any app that refuses to take part in or fails our audit will be banned.

Question 13. Does Facebook believe it would have a responsibility to report such incidents to the FTC under the consent decree? If such incidents have occurred, has Facebook reported them to the FTC?

The July 27, 2012 Consent Order memorializes the agreement between Facebook and the FTC and does not require ongoing reporting.

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 ongoing review by an independent assessor (Sections IV and V). Facebook accurately represented the operation of its developer Platform and the circumstances under which people could share data (including friends data) with developers, honored the restrictions of all privacy settings that covered developer access to data, and implemented a comprehensive privacy program build on industry-leading controls and principles, which has undergone ongoing review by an independent assessor approved by the FTC.

Cambridge Analytica Timeline Questions

There have been conflicting reports regarding the timeline of Facebook's response to the "thisisyourdigitallife" application developed for Cambridge Analytica. Please provide specific information about Facebook's response to the matter.

Question 14. With respect to the harvesting of user data from the "thisisyourdigitallife" application, for each the following (a) Cambridge Analytica, (b) Christopher Wylie, and (c) Dr. Kogan, on what date did Facebook:

- 1. First contact that party about the data collected from the application?**

2. Seek certification that the partys copy of the data was destroyed?

3. Receive the certification from party?

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. (SCL)/Cambridge Analytica. If this occurred, 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 executed and provided to Facebook 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 it 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 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.

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.

Question 15. Was Facebook aware at that time that Cambridge Analytica had developed other platform applications to collect user data? What applications did it delete due to associations with Cambridge Analytica and when were they removed from the platform?

Our investigation of Cambridge Analytica's advertising activities is ongoing, and we have banned Cambridge Analytica from purchasing ads on our platform. Cambridge Analytica generally utilized custom audiences, some of which were created from contact lists and other identifiers that it generated and uploaded to our system to identify the people it wanted to deliver ads to on Facebook, and in some instances, refined those audiences with additional targeting attributes.

Facebook's "People You May Know" Feature

Facebook's "People You May Know" feature has drawn attention for disclosures that reveal sensitive relationships, such as psychiatrists who have reported that their clients were recommended to each other.

Question 16. What pieces of data does Facebook use for the PYMK feature? Has it ever used data collected from data brokers for this purpose?

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

Question 17. Has PYMK ever used location to make recommendations and does it currently? If so, is this based on device reported geolocation or IP address?

PYMK uses country-level location to help users find friends.

Question 18. Does Facebook provide users with the ability to opt out of data collected from them or data about them being used by PYMK?

See Response to Question 16.

Question 19. Has the PYMK feature ever bypassed the privacy controls in order to perform its analytics for recommendations? For example, if a user's friends list is set to private, will Facebook still use this data to make recommendations to others?

See Response to Question 16.

Other Cambridge Analytics

Over a month ago, Mr. Zuckerberg stated that one of Facebook's next responsibilities was to "make sure that there aren't any other Cambridge Analytics out there." One would expect that review process would include identifying past cases where Facebook identified or took action against third-party developers over their data collection practices.

Question 20. When the company Klout automatically created accounts and assigned social popularity scores for the children of Facebook users, did Facebook send a deletion letter or exercise its right to audit?

In 2011, Facebook contacted Klout regarding potential violations of Facebook policies. Facebook determined that these issues had been resolved by Dec. 2011.

We use a variety of tools to enforce Facebook policies against violating parties, including developers. We review tens of thousands of apps per year and regularly disapprove noncompliant apps as part of our proactive review process. We also use tools like cease and desist letters, account suspensions, letter agreements, and civil litigation. For example, since 2006, Facebook has sent over 1,150 cease-and-desist letters to over 1,600 targets. In 2017, we took action against about 370,000 apps, ranging from imposing certain restrictions to removal of the app from the platform. Moreover, we have required parties who have procured our data without authorization to delete that data. We have invested significant resources in these efforts. Facebook is presently investigating apps that had access to large amounts of information before we changed our platform policies in 2014 to significantly reduce the data apps could access. To date around 200 apps (from a handful of developers: Kogan, AIQ, Cube You, the Cambridge Psychometrics Center, and myPersonality) have been suspended—pending a thorough investigation into whether they did in fact misuse any data.

Additionally, we have suspended an additional 14 apps, which were installed by around one thousand people. They were all created after 2014, after we made changes to more tightly restrict our platform APIs to prevent abuse. However, these apps appear to be linked to AIQ, which was affiliated with Cambridge Analytica. So, we have suspended them while we investigate further. Any app that refuses to take part in or fails our audit will be banned.

Question 21. How many times was Facebook made aware of privacy breaches by applications?

Facebook's policies regarding third-party usage of its platform technologies have prohibited—and continue to prohibit—those third-party app developers from selling or licensing user data obtained from Facebook and from sharing any user data obtained from Facebook with any ad network, data broker or other advertising or monetization-related service. We will investigate all apps that had access to large amounts of information before we changed our platform in 2014 to reduce data access, and we will conduct a full audit of any app with suspicious activity.

Question 22. How many times did Facebook send a deletion letter to an application developer for strictly privacy violations?

We use a variety of tools to enforce Facebook policies against violating parties, including developers. We review tens of thousands of apps per year and regularly disapprove noncompliant apps as part of our proactive review process. We also use tools like cease and desist letters, account suspensions, letter agreements, and civil litigation. For example, since 2006, Facebook has sent over 1,150 cease-and-desist letters to over 1,600 targets. In 2017, we took action against about 370,000 apps, ranging from imposing certain restrictions to removal of the app from the platform. Moreover, we have required parties who have procured our data without authorization to delete that data. We have invested significant resources in these efforts. Facebook is presently investigating apps that had access to large amounts of information before we changed our platform policies in 2014 to significantly reduce the data apps could access. To date around 200 apps (from a handful of developers: Kogan, AIQ, Cube You, the Cambridge Psychometrics Center, and myPersonality) have been suspended—pending a thorough investigation into whether they did in fact misuse any data.

Question 23. How many times did Facebook perform an audit on an application for strictly privacy violations?

See Response to Question 22.

Question 24. How many times did Facebook initiate litigation for strictly privacy violations?

See Response to Question 22.

Question 25. How many times did Facebook impose a moratorium or ban on an application developer for strictly privacy violations?

See Response to Question 22.

Question 26. Does Facebook plan to provide public disclosure of incidents where it finds that user data was improperly obtained or transferred by third-party application developers?

We are in the process of investigating every app that had access to a large amount of information before we changed our Platform in 2014. The investigation process is in full swing, and it has two phases. First, a comprehensive review to identify every app that had access to this amount of Facebook data and to focus on apps that present reason for deeper investigation. And second, where we have concerns, we will conduct interviews, make requests for information (RFI)—which ask a series of detailed questions about the app and the data it has access to—and perform audits using expert firms that may include on-site inspections. We have large teams of internal and external experts working hard to investigate these apps as quickly as possible. To date thousands of apps have been investigated and around 200 apps have been suspended—pending a thorough investigation into whether they did in fact misuse any data. Where we find evidence that these or other apps did misuse data, we will ban them and let people know.

These apps relate to a handful of developers: Kogan, AIQ, Cube You, the Cambridge Psychometrics Center, and myPersonality, with many of the suspended apps being affiliated with the same entity. Many of these suspensions include apps that appear to be “test” apps that were

never released to the public, and therefore would not have acquired significant user data, although our investigation into these apps is ongoing.

Additionally, we have suspended an additional 14 apps, which were installed by around one thousand people. They were all created after 2014, after we made changes to more tightly restrict our platform APIs to prevent abuse. However, these apps appear to be linked to AIQ, which was affiliated with Cambridge Analytica. So, we have suspended them while we investigate further. Any app that refuses to take part in or fails our audit will be banned.

We will commit to briefing your staff on future developments.

Facebook Privacy Settings

This month, Facebook began to roll out changes to comply with new European data protection rules. These updates include a new consent process that affects how Facebook uses sensitive data and whether facial recognition is enabled, among other factors.

Question 27. Has Facebook engaged in user testing or other analysis that assessed how platform changes and interface design influence the adoption of certain privacy settings?

We routinely test new products and consent flows before rolling them out broadly to ensure that there are no bugs or unintended behaviors that would lead to an unintended or negative user experience. 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. We will continue to monitor how these and other privacy settings perform with users. It’s important to us that people have the information they need to make the privacy choices that are right for them.

Question 28. Has Facebook ever tested platform changes and interface design to determine whether it would lead to users allowing more permissive privacy settings?

At Facebook, we make decisions about privacy through a cross-functional, cross-disciplinary effort 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 is responsible for reviewing product launches, major changes, and privacy-related bug fixes to products and features to ensure that privacy policies and procedures are consistently applied and

that key privacy decisions are implemented for the product. This approach has several key benefits.

First, it is designed to consider privacy early in the product development process. This allows us to consider the benefits that a feature is intended to have for people who use our services, how data will be used to deliver those benefits, and how we can build features from the ground up that include privacy protections to enable those benefits while protecting people's information and putting them in control.

Second, while complying with our obligations is critically important, taking a cross-disciplinary approach to privacy encourages us to think about data protection as more than just a compliance exercise. Instead, we evaluate how to design privacy into the features that we build and consider this from the perspective of things like how we design interfaces that make data use intuitive, taking a consistent approach to privacy across our services, and building protections in how our software is engineered. Accordingly, while we scale our privacy review process depending on the complexity of a particular data use, reviews typically involve experts who evaluate proposed data practices from the perspective of multiple disciplines.

As part of our consent agreement with the Federal Trade Commission, we submit a report to the FTC every two years. That report is based on assessments conducted by an independent third party on a bi-annual basis, which require us to submit evidence to demonstrate the effectiveness of the program.

EU Data Protection Regulations

In Europe, under new data protection regulations, Facebook will be required to provide users with more clear opportunities to provide consent and afford more protections to that data. While Facebook has stated that it will offer some of those protections for users outside of Europe, it has not committed to providing all of these protection. I am interested in what rules Congress should put into place for such data.

Question 29. Would Facebook support a requirement that users be provided with clear and plain information about the use of their data?

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.

Question 30. Would Facebook support a requirement that users be allowed to download and take their data to competitive services?

Facebook already allows users to download a copy of their information from Facebook. This functionality, which we've offered for many years, includes numerous categories of data, including About Me, Account Status History, Apps, Chat, Follower, Following, Friends, Messages, Networks, Notes, and more. We recently launched improvements to our "Download

Your Information” tool, including to give people choices about whether they want to download only certain types of information and about the format in which they want to receive the download, to make it easier for people to use their information once they’ve retrieved it.

Question 31. Would Facebook support a requirement that users are assured that their data is actually deleted when they request its deletion or close their account?

In general, when a user deletes their account, we delete things they have posted, such as their photos and status updates, and they won’t be able to recover that information later. (Information that others have shared about them isn’t part of their account and won’t be deleted.)

There are some limited exceptions to these policies: For instance, information can be accessed and preserved for an extended period when it is the subject of a legal request or obligation, governmental investigation, or investigations of possible violations of our terms or policies, or otherwise to prevent harm. We also retain information from accounts disabled for terms violations for at least a year to prevent repeat abuse or other term violations.

Question 32. Would Facebook support a requirement of mandatory and timely disclosure of breaches?

Facebook is generally open to the idea of breach notification requirements, particularly legislation that would centralize reporting and ensure a consistent approach across the United States. For example, in Europe, the GDPR requires notification to a lead supervisory authority, rather than individual member states, in cases of a data breach. In the United States, however, there is no centralized notification scheme, and instead, reporting obligations vary widely across all 50 states. This complexity makes it harder to respond appropriately and swiftly to protect people in the event of a data breach. We believe this is an important issue and an area that is ripe for thoughtful regulation.

Question 33. Would Facebook support a requirement for a baseline technical and organizational measures to ensure adequate data security?

Facebook is generally not opposed to regulation but wants to ensure it is the right regulation. The issues facing the industry are complex, multi-faceted, and affect an important part of peoples’ lives. As such, Facebook is absolutely committed to working with regulators, like Congress, to craft the right regulations. Facebook would be happy to review any proposed legislation and provide comments.

Russian Interference

As early as June 2015, the New York Times Magazine had documented the Internet Research Agency’s interest in interfering with American politics, and even named specific Facebook accounts associated in the disinformation effort. The way that Facebook is designed, outsiders have very little insight into these efforts. And yet, the Russian media outlet RBC had identified accounts that were paying to spread content several months before Facebook took notice. New York Times also claims that as early as November 2016, Facebook’s Chief Security Officer Alex Stamos had uncovered evidence that Russian operatives used the platform to weaponized information obtained from the hacking of the DNC and the Clinton campaign.

In a CNN interview, Mr. Zuckerberg for the first time disclosed that Facebook had found “a lot of different accounts coming from Macedonia” to spread false news during the Alabama special election. That election, another one decided by only small margin, was months ago. Mr. Zuckerberg acknowledged that Facebook expects there will be attempts to interfere in the midterm elections with newer tactics, a belief shared by the intelligence community.

Question 34. Will you commit to providing Congress with information about disinformation and propaganda campaigns on a timely basis prior to the midterm elections?

We recently outlined steps we are taking on election integrity here:
<https://newsroom.fb.com/news/2018/03/hard-questions-election-security/>.

Further, pursuant to the new transparency measures Facebook is launching, all advertisers who want to run ads with political content targeted at the US will have to confirm their identity and location by providing either a US driver’s license or passport, last four digits of their social security number, and a residential mailing address. Ads that include political content and appear on Facebook or Instagram will include a “Paid for by” disclaimer provided by the advertisers that shows the name of the funding source for the ad.

Question 35. The New York Times reports details of Russian interference were removed from the April 2017 report “Information Operations and Facebook” by management due to political and business reasons. Will Facebook provide Congress with the original draft of the report?

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 prevent bad actors from using misinformation to undermine the democratic process.

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 the 10 most important changes we have made:

- **Ads 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 content, we’ve created an archive that will hold ads with political content for seven years—including information about ad impressions and spend, as well as demographic data such as age, gender, and location. People in Canada and Ireland can already see all the ads that a Page is running on Facebook—and we’re launching this globally in June.
- **Verification and labeling.** Every advertiser will now need confirm their ID and location before being able to run any ads with political content in the US. All ads with political 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 with 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.** We're doubling the number of people working on safety and security from 10,000 last year to over 20,000 this year. 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 34 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.** In recent months, we've started to deploy 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 plan to continue 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.

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

Hate Speech

Over the past months, human rights organizations and other civil society groups have raised attention to concerns over Facebook’s insufficient response to hate speech in countries where there is a credible threat of violence. In addition to Myanmar, the New York Times recently published an article on how mob violence against Muslims in Sri Lanka was spurred by a baseless rumor that a Muslim restaurant owner was secretly feeding sterilization pills to women from the Sinhalese-Buddhist community.

Mr. Zuckerberg and other members of Facebook management have expressed a renewed commitment to providing resources to address these threats, including taking action to address those who generate hate speech. As Mr. Zuckerberg noted, AI will not be able to resolve such complex matters in the near or medium term, necessitating teams that deal with local languages and context. While Facebook currently has approximately 1,200 German content reviewers to comply with regulations, it only has plans to hire “dozens” of Burmese content reviewers. Hiring staff with reviewers, market specialists and analysts with the appropriate expertise can be difficult, but these reports of violence demonstrate the human cost of insufficient community resources to handle content and complaints.

Question 36. What “specific product changes” will you be making to address hate speech in such countries? Will the new product changes enable content that violates Facebook’s Community Standards to be removed within 24 hours?

We’ve been too slow to deal with the hate and violence in places like Myanmar and Sri Lanka. The challenges we face in a country that has fast come online are very different than those in other parts of the world, and we are investing in people, technology, and programs to help address them as effectively as possible.

We are increasing the number of Burmese and Sinhalese-language content reviewers as we continue to grow and invest in Myanmar and Sri Lanka. Our goal is always to have the right number of people with the right native language capabilities to ensure incoming reports are reviewed quickly and effectively. That said, there is more to tackling this problem than reported

content. A lot of abuse may go unreported, which is why we are supplementing our hiring with investments in technology and programs.

We are building new tools so that we can more quickly and effectively detect abusive, hateful, or false content. We have, for example, designated several hate figures and organizations for repeatedly violating our hate speech policies, which has led to the removal of accounts and content that support, praise, or represent these individuals or organizations. We are also investing in artificial intelligence that will help us improve our understanding of dangerous content.

We are further strengthening our civil society partner network so that we have a better understanding of local context and challenges. We are focusing on digital literacy education with local partners in Myanmar and Sri Lanka. For example, we launched a local language version of our Community Standards (<https://www.facebook.com/safety/resources/myanmar>) to educate new users on how to use Facebook responsibly in 2015 and we have been promoting these actively in Myanmar, reaching over 8 million people through promotional posts on our platform alone. We've also rolled out several education programs and workshops with local partners to update them on our policies and tools so that they can use this information in outreach to communities around the country. One example of our education initiatives is our work with the team that developed the Panzagar initiative (<https://www.facebook.com/supportflowerspeech>) to develop the Panzagar counterspeech Facebook stickers to empower people in Myanmar to share positive messages online. We also recently released locally illustrated false news tips, which were promoted on Facebook and in consumer print publications. We have a dedicated Safety Page for Myanmar (<https://www.facebook.com/safety/resources/myanmarand>) and have delivered hard copies of our local language Community Standards and safety and security tips to civil society groups in Myanmar who have distributed them around the country for trainings. Similarly, in Sri Lanka, we ran a promotion in English, Sinhalese, and Tamil at the top of News Feeds in April 2017 to educate people on our Community Standards, in particular hate speech. The content has been viewed almost 100M times by almost 4M people.

Question 37. Does Facebook believe that it has hired or will hire within the year a sufficient number of content reviewers and established local emergency points of contact for all regions where its platform could inadvertently facilitate communal violence?

We are investing in people, technology, and programs to help address the very serious challenges we have seen in places like Myanmar and Sri Lanka.

Our content review teams around the world—which grew by 3,000 people last year—work 24 hours a day and in over 50 languages.

Over the last two years, we have added dozens more Burmese language reviewers to handle reports from users across our services, and we plan to more than double the number of content reviewers focused on user reports. We also have increased the number of people across the company working on Myanmar-related issues and we have a special product team working to better understand the local challenges and build the right tools to help keep people in the country safe. We will continue to hire more staff dedicated to Myanmar, including Burmese speakers and policy experts.

In Sri Lanka, we are increasing the number of Sinhalese language experts sevenfold.

From a programmatic perspective, we will continue to work with experts to develop safety resources and counter-speech campaigns in these regions and conduct regular training for civil society and community groups on using our tools.

Facebook is committed to continuing to provide a platform where people can raise awareness about human rights abuses around the globe, and we have a track record of partnering with experts and local organizations on these issues. For example, we have been part of the Global Network Initiative (GNI) since 2013. That organization brings together industry, civil society, academics, and socially-responsible investors to address freedom-of-expression and privacy issues online. An independent assessor conducted a human-rights-impact assessment of Facebook to confirm that we comply with GNI's principles.

Question 38. What product changes, operational decisions, and resource allocations has Facebook made in order to avoid future risks such as those made abundantly clear in Myanmar and Sri Lanka?

We are working to enable freedom of expression around the globe and ensure that our platform is safe. Our Community Standards account for situations in which people may be raising awareness of and/or condemning violence; however, they prohibit hate speech and celebrating graphic violence. Drawing that line can be complex, which is why we work with experts and external groups, including local civil society organizations in places like Myanmar and Sri Lanka, to ensure that we are taking local context and challenges into account. Our content review team, which includes native language speakers, carefully reviews reports that we receive from the public, media, civil society, and governments. We remove content that violates our policies, regardless of who posted the content (including the government). We have also been working with local communities and NGOs for years in these regions to educate people about hate speech, news literacy, and our policies. For example, we have introduced an illustrated, Myanmar language specific copy of our community standards and a customized safety Page, which we work with our local partners to promote, and we recently ran a series of public service ads in Myanmar that we developed with the News Literacy Project to help inform people about these important issues.

Question 39. What emergency processes for escalation do you have in place for situations where there is content inciting people to violence, such as what happened in Sri Lanka?

We have clear rules against hate speech and content that incites violence, and we remove such content as soon as we're made aware of it. In response to the situation in Sri Lanka, we're building up teams that deal with reported content, working with civil society and government to learn more about local context and changing language, and exploring the use of technology to help. We want to provide direct reporting channels to civil society partners so that they can alert us to offline activity that might prompt an increase in violating content on Facebook. We work with local civil society organizations to understand what types of reporting channels would best serve their specific communities and are engaging with organizations in Sri Lanka to understand what more we can do. We are committed to having the right policies, products, people, and partnerships in place to help keep our community in Sri Lanka safe.

Question 40. In the context of Sri Lanka and Myanmar, rumors present a credible threat of violence and have resulted in violence. Are rumors such as those in Sri Lanka interpreted

as violations under your existing “credible threat” policy? How do your systems or reporting mechanisms account for such country or context specific threats? Given how quickly such content can lead to violence, do you apply different processes or response time targets to prioritize content categorized as hate speech?

We require everyone on Facebook to comply with our Community Standards, and we carefully review reports of threatening language to identify serious threats of harm to public and personal safety. We recognize our services have an important role to play in countries that are fast coming online. That’s why we’re investing in people, technology, and programs to address the challenges we face in these countries. We’ve added more local language reviewers, established dedicated product teams, rolled out better reporting tools and appeals, and are removing fake accounts, hate groups and individuals. We remove credible threats of physical harm to individuals and specific threats of theft, vandalism, or other financial harm. We also prohibit the use of Facebook to facilitate or organize criminal activity that causes physical harm to people, businesses or animals, or financial damage to people or businesses, and we work with law enforcement when we believe there is a genuine risk of physical harm or direct threats to public safety. As part of our work in places like Sri Lanka and Myanmar, we are strengthening our relationships with civil society organizations to ensure we are taking local context, challenges, and tensions into account.

Question 41. The anti-Muslim monk, U Wirathu, was reportedly banned by Facebook in January 2018 after having been frequently reported for hate content. Despite several bans, he was able to recreate a presence on the platform on several occasions and there are to this day accounts which carry his name. What mechanisms do you have in place to remove users who repeatedly breach Facebook’s Community Standards and what actions are you taking to guarantee their permanent removal?

Our Community Standards (<https://www.facebook.com/communitystandards>) prohibit hate speech that targets people based on their race, ethnic identity, or religion. We remove violating content when it is reported to us. We also have designated several hate figures and hate organizations in Myanmar. These include Wirathu, Thuseitta, Ma Ba Tha, and Parmaukkha. This means these individuals or organizations are not allowed a presence on Facebook, and we will remove accounts and content that support, praise or represent these individuals or organizations.

In addition to removing content that violates our Community Standards or Page Terms, we disable the accounts of repeat infringers in appropriate circumstances.

Over the last several months, we have proactively searched for and removed content on the platform that praises, supports, or represents Wirathu.

Human Rights - Iran

Iranian women’s rights and pro-democracy advocates have reported that copyright infringement and content reporting mechanisms have been instrumentalized by pro-government actors to take down their Instagram pages and Facebook Groups over the past several years. While community reporting mechanisms are necessary, and often legally required, for operating a platform as large as Facebook, the threat posed by abusive reporting also demonstrates the need for human reviewers. Likewise, the trolling, hacking, and impersonation that frequently target Iranian dissidents also necessitate teams that are

empowered to deal with the Persian language and the Iranian context. However, many activists have struggled to establish relationships or receive help from Facebook to have such issues addressed.

Question 42. What measures has Facebook taken to address the abusive use of copyright reporting mechanisms being used to take down Iranian content?

We recognize that individuals and entities may purposefully report content en masse in an attempt to stifle speech. That is why we believe content must be reviewed with the appropriate context.

We are proud that our platform has been used to inspire people to stand up for their beliefs and values, even in the face of intimidating opposition, and we regularly provide tools and programmatic resources to activists and journalists. We also make materials available to ensure activists and journalists are able to use Facebook safely.

Based on the foundation established in the Universal Declaration of Human Rights and the UN Guiding Principles on Business and Human Rights, Facebook joined the ICT-sector specific Global Network Initiative in 2013. As part of our commitments as a GNI member, we routinely conduct human rights impact assessments of our product and policy decisions and engage with external stakeholders to inform this work. We are also independently assessed against our compliance with the GNI Principles every two years.

Question 43. What measures, such as verification of accounts, has Facebook taken to address the impersonation of Iranian activists, cultural dissidents, and other public figures?

Claiming to be another person violates our Community Standards, and we want to make it harder for anyone to be impersonated on our platform. Users can also report accounts that are impersonating them. We've developed several techniques to help detect and block this type of abuse. At the time someone receives a friend request, our systems are designed to check whether the recipient already has a friend with the same name, along with a variety of other factors that help us determine if an interaction is legitimate. Further, we recently announced new features that use face recognition technology that may help detect when someone is using another user's image as their profile photo—which helps stop impersonation. This is an area we're continually working to improve so that we can provide a safe and secure experience on Facebook.

Questions from Senator Schatz

You said at the hearing that Facebook users own and control their data. But I am not persuaded that the company has done an adequate job explaining, for example, what specific information the company collects about individuals, how that information is being used and kept safe, and how they can easily delete or modify it. If you and your company are committed to putting privacy first, I urge that you answer these questions in a precise, accurate, but straightforward way. I understand your legal team will be reviewing this, but I hope you resist complexity and answer these questions in a way that any American could understand.

Question 1. Please list and describe all of the types and categories of data that Facebook collects and how Facebook uses this data. This includes, but is not limited to, data collected:

- **on the Facebook platform (e.g., posts, messages, and search history);**
- **off the Facebook platform (quantify how ubiquitous Facebook's plugins are on the web, for instance);**
- **on products offered by Facebook family companies;**
- **on specific devices (e.g., smartphone microphone and camera, other apps, data from the operating system);**
- **via third-party companies and app developers;**
- **from data brokers; and**
- **from publishers.**

For each, describe whether users own the data, and what options users have to modify or delete the data.

We believe that it's important to communicate with people about the information that we collect and how people can control it. That 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 as people are using the Facebook service itself.

We've heard loud and clear that privacy settings and other important tools are too hard to find and that we must do more to keep people informed. So, we're taking 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're also updating our Data Policy to better spell out what data we collect and how we use it in Facebook, Instagram, Messenger, and other products.

In response to your specific questions, depending on which Services a person uses, we collect different kinds of information from or about them. This is described in our Data Policy:

- **Things Users and others do and provide.** Information and content users provide. We collect the content, communications and other information users provide when they use our Products, including when they sign up for an account, create or share content, and message or communicate with others. This can include information in or about the content they provide (like metadata), such as the location of a photo or the date a file was created. It can also include what users see through features we provide, such as our camera, so we can do things like suggest masks and filters that they might like, or give users tips on using camera formats. Our systems automatically process content and communications users provide to analyze context and what's in them for the purposes described below. Learn more about how people can control who can see the things they share.
 - Data with special protections: Users can choose to provide information in their Facebook profile fields or Life Events about their religious views, political views, who they are “interested in,” or their health. This and other information (such as racial or ethnic origin, philosophical beliefs, or trade union membership) could be subject to special protections under the laws of their country.
- **Networks and connections.** We collect information about the people, Pages, accounts, hashtags, and groups users are connected to and how they interact with them across our Products, such as people a user communicates with the most or groups users are part of. We also collect contact information if they choose to upload, sync, or import it from a device (such as an address book or call log or SMS log history), which we use for things like helping them and others find people they may know and for the other purposes listed below.
- **People's usage.** We collect information about how people use our Products, such as the types of content they view or engage with; the features they use; the actions they take; the people or accounts they interact with; and the time, frequency, and duration of their activities. For example, we log when they're using and have last used our Products, and what posts, videos, and other content they view on our Products. We also collect information about how they use features like our camera.
- **Information about transactions made on our Products.** If people use our Products for purchases or other financial transactions (such as when users make a purchase in a game or make a donation), we collect information about the purchase or transaction. This includes payment information, such as their credit or debit card number and

other card information; other account and authentication information; and billing, shipping, and contact details.

- **Things others do and information they provide about users.** We also receive and analyze content, communications, and information that other people provide when they use our Products. This can include information about them, such as when others share or comment on a photo of a user, send a message to them, or upload, sync or import their contact information.
- **Device Information.** As described below, we collect information from and about the computers, phones, connected TVs and other web-connected devices they use that integrate with our Products, and we combine this information across different devices they use. For example, we use information collected about their use of our Products on their phone to better personalize the content (including ads) or features they see when they use our Products on another device, such as their laptop or tablet, or to measure whether they took an action in response to an ad we showed they on their phone on a different device.

Information we obtain from these devices includes:

- Device attributes: information such as the operating system, hardware and software versions, battery level, signal strength, available storage space, browser type, app and file names and types, and plugins.
- Device operations: information about operations and behaviors performed on the device, such as whether a window is foregrounded or backgrounded, or mouse movements (which can help distinguish humans from bots).
- Identifiers: unique identifiers, device IDs, and other identifiers, such as from games, apps or accounts people use, and Family Device IDs (or other identifiers unique to Facebook Company Products associated with the same device or account).
- Device signals: Bluetooth signals, and information about nearby Wi-Fi access points, beacons, and cell towers.
- Data from device settings: information users allow us to receive through device settings people turn on, such as access to their GPS location, camera, or photos.
- Network and connections: information such as the name of users' mobile operator or ISP, language, time zone, mobile phone number, IP address, connection speed and, in some cases, information about other devices that are nearby or on users' network, so we can do things like help people stream a video.
- Cookie data: data from cookies stored on a user's device, including cookie IDs and settings. Learn more about how we use cookies in the Facebook

Cookies Policy (<https://www.facebook.com/policies/cookies/>) and Instagram Cookies Policy (<https://www.instagram.com/legal/cookies/>).

- **Information from partners.** Advertisers, app developers, and publishers can send us information through Facebook Business Tools they use, including our social plug-ins (such as the Like button), Facebook Login, our APIs and SDKs, or the Facebook pixel. These partners provide information about users' activities off Facebook—including information about a user's device, websites users visit, purchases users make, the ads they see, and how they use their services—whether or not they have a Facebook account or are logged into Facebook. For example, a game developer could use our API to tell us what games users play, or a business could tell us about a purchase a user made in its store. We also receive information about a user's online and offline actions and purchases from third-party data providers who have the rights to provide us with their information. Partners receive user data when users visit or use their services or through third parties they work with. We require each of these partners to have lawful rights to collect, use and share user data before providing any data to us.

People own what they share on Facebook, and they can manage things like who sees their posts and the information they choose to include on their profile.

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 to edit or delete these interests. They can choose not to see ads from a particular advertiser or not to see ads based on their use of third-party websites and apps. They also can choose not to see ads off Facebook that are based on the interests we derive from their activities on Facebook.

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 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, will 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/>.

And we recently announced plans to build Clear History. This feature will enable users to see the websites and apps that send us information when they use them, delete this information from their accounts, and turn off our ability to store it associated with their accounts going forward. Apps and websites that use features such as the Like button or Facebook Analytics send us information to make their content and ads better. We also use this information to make users’ experiences on Facebook better. If a user clears their history or use the new setting, we’ll remove identifying information so a history of the websites and apps they’ve used won’t be associated with their account. We’ll still provide apps and websites with aggregated analytics—for example, we can build reports when we’re sent this information so we can tell developers if their

apps are more popular with men or women in a certain age group. We can do this without storing the information in a way that's associated with a user's account, and as always, we don't tell advertisers who a user is.

Question 2. What data does Facebook collect about non-users? For example, when a user first joins Facebook, what data has Facebook already typically collected about them? Assume that the new user is an average American and active web user with many friends who are already on Facebook. List the attributes that Facebook would typically know about the new user and where that information comes from. If Facebook collects information about non-users, what is the purpose?

Facebook does not create profiles or track website visits for people without a Facebook account.

When people visit apps or websites that feature our technologies—like the Facebook Like or Comment button—our servers automatically log (i) standard browser or app records of the fact that a particular device or user visited the website or app (this connection to Facebook's servers occurs automatically when a person visits a website or app that contains our technologies, such as a Like button, and is an inherent function of Internet design); and (ii) any additional information the publisher of the app or website chooses to share with Facebook about the person's activities on that site (such as the fact that a purchase was made on the site). This is a standard feature of the Internet, and most websites and apps share this same information with multiple different third-parties whenever people visit their website or app. For example, the Senate Commerce Committee's website shares information with Google and its affiliate DoubleClick and with the analytics company Webtrends. This means that, when a person visits the Committee's website, it sends browser information about their visit to each one of those third parties. More information about how this works is available at <https://newsroom.fb.com/news/2018/04/data-off-facebook/>.

When the person visiting a website featuring Facebook's tools is not a registered Facebook user, Facebook does not have information identifying that individual, and it does not create profiles for this individual.

We use the browser and app logs that apps and websites send to us—described above—in the following ways for non-Facebook users. First, these logs are critical to protecting the security of Facebook and to detecting or preventing fake account access. For example, if a browser has visited hundreds of sites in the last five minutes, that's a sign the device might be a bot, which would be an important signal of a potentially inauthentic account if that browser then attempted to register for an account. Second, we aggregate those logs to provide summaries and insights to websites and apps about how many people visit or use their product or use specific features like our Like button—but without providing any information about a specific person. We do not create profiles for non-Facebook users, nor do we use browser and app logs for non-Facebook users to show targeted ads from our advertisers to them or otherwise seek to personalize the content they see. However, we may take the opportunity to show a general ad that is unrelated to the attributes of the person or an ad encouraging the non-user to sign up for Facebook.

Question 3. Last year, how many Facebook users clicked on their privacy settings at least once? What was the average time a user spent adjusting their privacy controls? How often does an average user go into their privacy settings (per year, for instance)? In 2017, how many times did Facebook modify the user experience of its privacy settings to better suit its users? What other analytics of this kind does Facebook measure?

Privacy is at the core of everything we do, and our approach to privacy starts with our commitment to transparency and control. Our threefold approach to transparency includes, first, whenever possible, providing information on the data we collect and use and how people can control it in context and in our products. Second, we provide information about how we collect and use data in our user agreements and related educational materials. And third, we enable people to learn more about the specific data we have about them through interactive tools such as Download Your Information, which lets people download a file containing data that they may want to take to another service, and Access Your Information, a tool we are launching that will let people more easily access and manage their data on Facebook.

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. People can control the audience for their posts and the apps that can receive their data. They can see and delete the history of their activities on Facebook, and, if they no longer want to use Facebook, they can delete their account and the data associated with it. Of course, we recognize 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. 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 are launching a new settings menu that features core privacy settings in a single place. We are always working to help people understand and control how their data shapes their experience on Facebook.

Question 4. At the hearing, you said that you don't believe that enough users read Facebook's terms-of-service policy. Facebook has some of tech's smartest UX and behavioral experts, which is evident by a platform that millions of people use for hours each week. How is Facebook applying its UX and behavioral expertise to track and improve user engagement in this area? What does Facebook know about its users' understanding of its terms-of-service? For example, how long do users take to read Facebook's policies, on average? What does this number indicate about whether users have actually read the material?

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 as people are using the Facebook service itself.

As to your specific question, there is no single number that measures how much time people spend understanding how Facebook services work, in large part because 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, 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, over the last 18 months, we’ve run a global 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.

Question 5. Recently you said Facebook would “make all controls and settings the same everywhere, not just in Europe.” Please describe these controls and settings and what they do? Would the modification of these controls and settings apply in the U.S. only to new users or to all users? Would Facebook commit to default those settings and controls to minimize, to the greatest extent, the collection and use of users’ data? What changes will U.S. users see in their settings and controls after this change is implemented? And what features and protections (including but not limited to controls and settings) will European Facebook users have that will differ from U.S. users after the company implements GDPR?

The GDPR requires companies to obtain explicit consent to process certain kinds of data (“special categories of data” like biometric data). We are seeking explicit consent 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 target ads. We recently began providing 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 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.

The controls and settings that Facebook is enabling as part of GDPR are already available to other users around the world, including in the US. We also provide identical levels of transparency in our user agreements and in product notices to people in the US that we are providing under GDPR.

In the US, where these settings are already in place, people will have a mechanism to maintain their current choice or to change it. In each of these cases, we want people to make the choice—not Facebook—so nobody’s settings will change as part of this roll out unless they choose to change an existing setting.

And we also provide the same tools for access, rectification, erasure, data portability and others to users in in the US and rest of world that we provide in Europe, and many of those tools (like our Download Your Information tool, Ads Preferences tool, and Activity Log) have been available globally for many years.

Questions from Senator Markey

1. Mr. Zuckerberg, your company has stated that it has “no plans” to include advertisements on Messenger Kids. Will you pledge that Facebook will never incorporate advertising into Messenger Kids or any future products for children 12 and under?

We have no plans to include advertising in Messenger Kids. Moreover, there are no in-app purchases, and we do not use the data in Messenger Kids to advertise to kids or their parents. In developing the app, we assembled a committee of advisors, including experts in child development, online safety, and media and children’s health, and we continue to work with them on an ongoing basis. In addition, we conducted roundtables with parents from around the country to ensure we were addressing their concerns and built the controls they need and want in the app. We are committed to approaching all efforts related to children 12 and under thoughtfully, and with the guidance and input of experts and parents.

2. In your response to my letter on the topic of Messenger Kids, you stated that your company will not “automatically” create a Facebook account for Messenger Kids users when those children turn 13. Will you commit to not share children’s information for targeted advertisements, once young users turn 13?

As we stated in our response to your earlier letter, we will not automatically create a Facebook account for Messenger Kids users, or automatically transition a Messenger Kids account into a Facebook account once a child turns 13. Contained within that commitment and our commitment not to use data collected within Messenger Kids to market to kids or their parents is a commitment that we will not automatically enable third parties to send targeted ads to children who have used Messenger Kids when the child turns 13.

Questions from Senator Udall

DATA PROTECTION ON FACEBOOK

QUESTION 1: The General Data Protection Regulation or “GDPR”, which will go into effect on May 25th of this year. Will Facebook provide the same privacy protections for consent, retention, data portability, and transparency to American consumers that it will provide to EU consumers?

The controls and settings that Facebook is enabling as part of GDPR are available to people around the world, including settings for controlling our use of face recognition on Facebook and for controlling our ability to use data we collect off Facebook Company Products to target ads. We recently began providing 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 consent to our updated terms. We provide the same tools for access, rectification, erasure, data portability and others to people in the US and rest of world that we provide in Europe, and many of those tools (like our Download Your Information tool, ad preferences tool, and Activity Log) have been available globally for many years. We also provide identical levels of transparency in our user agreements and in product notices to people in the United States that we are providing under GDPR.

QUESTION 2: What kind of privacy review is required to make a change to Facebook that impacts user privacy? When did that level of review become mandatory?

At Facebook, we make decisions about privacy through a cross-functional, cross-disciplinary effort overseen by the Chief Privacy Officer 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 is responsible for reviewing product launches, major changes, and privacy-related bug fixes to products and features to ensure that privacy policies and procedures are consistently applied and that key privacy decisions are implemented for the product. This approach has several key benefits:

- First, it is designed to consider privacy early in the product development process. This allows us to consider the benefits that a feature is intended to have for people who use our services, how data will be used to deliver those benefits, and how we can build features from the ground up that include privacy protections to enable those benefits while protecting people’s information and putting them in control.
- Second, while complying with our obligations is critically important, taking a cross-disciplinary approach to privacy encourages us to think about data protection as more than just a compliance exercise. Instead, we evaluate how to design privacy into the features that we build and consider this from the perspective of things like how we design interfaces that make data use intuitive, taking a consistent approach to privacy across our services, and building protections in how our software is engineered. Accordingly, while we scale our privacy review process depending on the complexity of a particular data use, reviews typically involve experts who evaluate proposed data practices from the perspective of multiple disciplines.

As part of our consent agreement with the Federal Trade Commission, we submit a report to the FTC every two years. That report is based on assessments conducted by an independent third party on a bi-annual basis, which require us to submit evidence to demonstrate the effectiveness of the program.

QUESTION 3: Before that level of review was required, what checks were in place to ensure new features would not adversely impact users' privacy? What level of seniority was required of employees to approve a launch of such a privacy-impacting feature? For example, have you ever allowed an intern make changes that impacts customers' privacy?

See Response to Question 2.

QUESTION 4: Has Facebook ever launched a feature that had to be turned off because of the privacy concerns? If yes, how many times has that happened, and how many users were impacted? Did you notify the users who were impacted?

See Response to Question 2.

RUSSIA/CAMBRIDGE ANALYTICA:

QUESTION 5: Between 2010 and 2015, 3rd party applications were able to keep data indefinitely. Can you say how many applications downloaded app users' data, their friends' data, or their personal messages in this period of time?

We are in the process of investigating every app that had access to a large amount of information before we changed our Platform in 2014. The investigation process is in full swing, and it has two phases. First, a comprehensive review to identify every app that had access to this amount of Facebook data and to focus on apps that present reason for deeper investigation. And second, where we have concerns, we will conduct interviews, make requests for information (RFI)—which ask a series of detailed questions about the app and the data it has access to—and perform audits using expert firms that may include on-site inspections. We have large teams of internal and external experts working hard to investigate these apps as quickly as possible. To date thousands of apps have been investigated and around 200 apps have been suspended—pending a thorough investigation into whether they did in fact misuse any data. Where we find evidence that these or other apps did misuse data, we will ban them and let people know.

These apps relate to a handful of developers: Kogan, AIQ, Cube You, the Cambridge Psychometrics Center, and myPersonality, with many of the suspended apps being affiliated with the same entity. Many of these suspensions include apps that appear to be “test” apps that were never released to the public, and therefore would not have acquired significant user data, although our investigation into these apps is ongoing.

Additionally, we have suspended an additional 14 apps, which were installed by around one thousand people. They were all created after 2014, after we made changes to more tightly restrict our platform APIs to prevent abuse. However, these apps appear to be linked to AIQ, which was affiliated with Cambridge Analytica. So, we have suspended them while we investigate further. Any app that refuses to take part in or fails our audit will be banned.

We will commit to briefing your staff on future developments.

QUESTION 6: Given the recent reports about Cambridge Analytica and the years of poor security around your data, what measures will be put into place to ensure that advertisers are not targeting ads using ill-gotten data?

We are not aware of any evidence to suggest that Kogan shared data obtained through his app with Russia or other foreign governments, but our investigation is ongoing. We are in the process of investigating every app that had access to a large amount of information before we changed our Platform in 2014.

In April 2014, we significantly restricted the types of data generally available to app developers and required apps seeking additional categories of data to undergo proactive review by our internal teams. We rejected more than half of the apps seeking these permissions, including the second version of Kogan's app.

We review apps to ensure that the requested permissions clearly improve the user experience and that the data obtained is tied to an experience within the app. We conduct a variety of manual and automated checks of applications on the platform for Policy compliance, as well as random sampling. 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.

Recently, we announced a number of additional steps we're taking to address concerns raised by Kogan's app.

- **Review our platform.** We will investigate all apps that had access to large amounts of data before the platform changes we announced in 2014, and we will audit any app where we identify suspicious activity. If we identify misuses of data, we'll take immediate action, including banning the app from our platform and pursuing legal action if appropriate.
- **Tell people about data misuse.** We will tell people about apps that have misused their data. This includes building a way for people to know if their data might have been accessed via the app. Moving forward, if we remove an app for misusing data, we will tell everyone who used it.
- **Turn off access for unused apps.** If someone has not used an app within the last three months, we will turn off the app's access to their data.
- **Restrict Facebook Login data.** We are changing Login, so that the only data that an app can request without app review will include name, profile photo, and email address. Requesting any other data will require approval from Facebook. We will 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've permitted those apps to use. But we're making it easier for

people to see what apps they use and the information they have shared with those apps.

- **Reward people who find vulnerabilities.** We launched the Data Abuse Bounty program so that people can report to us any misuses of data by app developers.
- **Update our policies.** We have updated our terms and Data Policy to explain how we use data and how data is shared with app developers.

QUESTION 7: Will your team re-architect the Facebook platform software architecture to ensure that 3rd party applications do not have the ability to store and share data?

In April 2014, we announced that we would more tightly restrict our platform 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. A small number of developers asked for and were granted short-term extensions beyond the one-year transition period, the longest of which lasted several months. These extensions ended several years ago. A transition period of this kind is standard when platforms implement significant changes to their technology base and was necessary here to avoid disrupting the experience of millions of people. New apps that launched after April 30, 2014 were required to use our more restrictive platform APIs. We required apps seeking additional categories of data to undergo proactive review by our internal teams. We rejected more than half of the apps seeking these permissions, including the second version of Kogan's app.

We review apps to ensure that the requested permissions clearly improve the user experience and that the data obtained is tied to an experience within the app. We conduct a variety of manual and automated checks of applications on the platform for Policy compliance, as well as random sampling. 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.

Recently, we announced a number of additional steps we're taking to address concerns raised by Kogan's app.

- **Review our platform.** We will investigate all apps that had access to large amounts of data before the platform changes we announced in 2014, and we will audit any app where we identify suspicious activity. If we identify misuses of data, we'll take immediate action, including banning the app from our platform and pursuing legal action if appropriate.
- **Tell people about data misuse.** We will tell people about apps that have misused their data. This includes building a way for people to know if their data might have been accessed via the app. Moving forward, if we remove an app for misusing data, we will tell everyone who used it.

- **Turn off access for unused apps.** If someone has not used an app within the last three months, we will turn off the app's access to their data.
- **Restrict Facebook Login data.** We are changing Login, so that the only data that an app can request without app review will include name, profile photo, and email address. Requesting any other data will require approval from Facebook. We will 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've permitted those apps to use. But we're making it easier for people to see what apps they use and the information they have shared with those apps.
- **Reward people who find vulnerabilities.** We launched the Data Abuse Bounty program so that people can report to us any misuses of data by app developers.
- **Update our policies.** We have updated our terms and Data Policy to explain how we use data and how data is shared with app developers.

We are investing so much in security that our costs will increase significantly. But we want to be clear about what our priority is: protecting our community is more important than maximizing our profits.

As our CEO Mark Zuckerberg has said, when you are building something unprecedented like Facebook, there are going to be mistakes. What people should hold us accountable for is learning from the mistakes and continually doing better—and, at the end of the day, making sure that we're building things that people like and that make their lives better.

QUESTION 8: How will you prevent another developer like Kogan from creating a viral app for the expressed purpose of gathering data and downloading, storing, and sharing that data?

See Response to Question 7.

QUESTION 9: How do you know that there are no other copies of the data that Kogan acquired from Facebook?

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, was accounted for and destroyed. Based on recent allegations, we have reopened our investigation into the veracity of these certifications and have hired a forensic auditor to conduct a forensic audit of Cambridge Analytica's systems. We are currently paused on the audit at the request of the UK Information Commissioner's Office request, which is conducting a regulatory investigation into Cambridge Analytica (based in the UK), and we hope to move forward with that audit soon.

We have suspended SCL/Cambridge Analytica from purchasing advertising on Facebook.

QUESTION 10: A March 2018 online article in *Quartz* reported that Facebook employees and Cambridge Analytica employees were both working in the Trump Campaign San Antonio headquarters.⁴ How will you ensure that your advertising salespeople are not engaging with entities previously identified for violating your terms of service?

No one from Facebook was assigned full-time to the Trump campaign, or full-time to the Clinton campaign. We offered identical support to both the Trump and Clinton campaigns, and had teams assigned to both. Everyone had access to the same tools, which are the same tools that every campaign is offered. We continuously work to ensure that we comply with all applicable laws and policies. While our investigation is ongoing, our review indicates that Facebook employees did not identify any issues involving the improper use of Facebook data in the course of their interactions with Cambridge Analytica during the 2016 US Presidential campaign.

QUESTION 11: In a recent press conference,⁵ you state that you are fully confident you are making progress against foreign actor manipulating the Facebook platform. Will you provide Congress and the American people auditable periodic reports about the progress you and your team are making on fighting disinformation on your platform?

We have worked to notify people about this issue, broadly, through our white paper in April 2017, Information Operations on Facebook, and our disclosures about the IRA last fall. We have also been publishing updates on these issues in our Newsroom.

THIRD PARTY APPLICATIONS:

QUESTION 12: How many times has Facebook enforced your terms of services against 3rd party application for misuse of data?

We use a variety of tools to enforce Facebook policies against violating parties, including developers. We review tens of thousands of apps per year and regularly disapprove noncompliant apps as part of our proactive review process. We also use tools like cease and desist letters, account suspensions, letter agreements, and civil litigation. For example, since 2006, Facebook has sent over 1,150 cease-and-desist letters to over 1,600 targets. In 2017, we took action against about 370,000 apps, ranging from imposing certain restrictions to removal of the app from the platform. Moreover, we have required parties who have procured our data without authorization to delete that data. We have invested significant resources in these efforts. Facebook is presently investigating apps that had access to large amounts of information before we changed our platform policies in 2014 to significantly reduce the data apps could access. To date around 200 apps (from a handful of developers: Kogan, AIQ, Cube You, the Cambridge Psychometrics

⁴ Kozlowska, Hanna. 20 March 2018. Facebook and Cambridge Analytica worked side by side at a Trump campaign office in San Antonio. <https://qz.com/1233579/facebook-and-cambridge-analytica-worked-side-by-side-at-a-trump-campaign-office-in-san-antonio/>.

⁵ Facebook. 4 April 2018. "Hard Questions: Q&A with Mark Zuckerberg on Protecting People's Information". <https://newsroom.fb.com/news/2018/04/hard-questions-protecting-peoples-information/>.

Center, and myPersonality) have been suspended—pending a thorough investigation into whether they did in fact misuse any data.

QUESTION 13: It's clear that, over the course of the Facebook platform program, enforcement of the Platform Policy has been reactive rather than proactive. Of all the 3rd party applications, how many such applications have been reviewed in the past 8 years? How many 3rd party applications have been removed from the platform due to violations of the terms of service?

See Response Question 12.

QUESTION 14: According to your Platform Policy, if you exceed 5 million monthly active users or 100M API calls per day, developers may be subject to additional terms. What are the additional terms? How many 3rd party applications are currently subject to additional terms?

In circumstances where developers make a high volume of API calls, Facebook may impose additional terms, which are generally negotiated and vary depending on which APIs are at issue.

In addition, Facebook has a set of APIs that enable certain partners, primarily operating systems and device manufacturers, to provide people with Facebook-like experiences (e.g., Facebook apps, news feed notifications, address book syncs) in their products. We developed these APIs, which are commonly known as “device-integrated APIs,” in the early days of mobile when the demand for Facebook outpaced our ability to build versions of our product that worked on every phone or operating system. Several dozen companies still used them at the start of the year, including Amazon, Apple, Blackberry, HTC, Microsoft, Huawei, Lenovo and Samsung, among others. On April 23, 2018, we announced that we would wind down these APIs. So far over 30 of these partnerships have been ended, including with Huawei.

These device-integrated APIs are different from the platform APIs that were used by Alexandr Kogan, which were the focus of the hearing and went to the heart of the Cambridge Analytica matter. Third party developers using our platform APIs built new, social experiences incorporating information that Facebook users brought with them; by contrast, the very point of our device-integrated APIs was to enable other companies to create Facebook functionality, primarily for devices and operating systems. The experiences that partners built using our device-integrated APIs were reviewed and approved by Facebook, and partners could not integrate the user's Facebook features without the user's permission.

QUESTION 15: For the Platform Policy for Messenger, how do you ensure that malicious actors are not using bots using the Messenger API to spread disinformation to users at a mass scale?

Businesses large and small are using bots for Messenger to connect with their customers in a way that is convenient, functional, and enables them to connect with customers at scale. We give people control of their experience. We offer a set of tools that allow a person to block or mute a bot or business at any time and people can also report bots where the Facebook Community Operations team will review and take action if appropriate. Finally, a few months ago we announced that bot developers are now required to have business verification for

apps/bots that need access to specialized APIs as a result of our ongoing efforts to ensure integrity across our platforms.

FACEBOOK - SUITE OF APPLICATION - ONAVO VPN:

QUESTION 16: Do know whether customers who download the virtual private network, or VPN, of Facebook’s subsidiary Onavo’s understand that any activity occurring on their mobile device is being collected and stored by Facebook? Doesn’t this practice violate the privacy consumers expect of a VPN?

When people first install the iOS version of the Onavo Protect app, we explain that Onavo uses a VPN that “helps keep you and your data safe by understanding when you visit potentially malicious or harmful websites and giving you a warning.” In addition, the first screen that a person sees when installing the app explains, under a heading that reads “Data Analysis”:

“When you use our VPN, we collect the info that is sent to, and received from, your mobile device. This includes information about: your device and its location, apps installed on your device and how you use those apps, the websites you visit, and the amount of data use.

This helps us improve and operate the Onavo service by analyzing your use of websites, apps and data. Because we’re a part of Facebook, we also use this info to improve Facebook products and services, gain insights into the products and services people value, and build better experiences.”

People must tap a button marked “Accept & Continue” after seeing this information in a full-screen interstitial before they can use the app.

The Android version of the Onavo Protect app offers data management features (e.g., the ability to block apps from using background data) that do not require users to enable the app’s VPN.

For both versions of the app, we communicate repeatedly and up front—in the App Store description, in Onavo’s Privacy Policy, and in-line at the time the user first opens the app after downloading it—that Onavo is part of Facebook and what that means for how Onavo Protect handles data in other ways.

More broadly, websites and apps have used market research services for years. We use Onavo, App Annie, comScore, and publicly available tools to help us understand the market and improve all our services. When people download Onavo to manage their data usage and help secure their connection, we are clear about the information we collect and how it is used. Like other VPNs, when the Onavo VPN is enabled, Onavo Protect helps create a secure connection, including when people are on public Wi-Fi. As part of this process, Onavo receives their mobile data traffic. This helps us improve and operate the Onavo service. Because we’re part of Facebook, we also use this information to improve Facebook products and services. We let people know about this activity, and other ways that Onavo uses, analyzes, and shares data (for example, the apps installed on users’ devices) in the App Store descriptions, and when they first open the app after downloading it.

Facebook does not use Onavo data for Facebook product uses, nor does it append any Onavo data or data about individuals' app usage to Facebook accounts.

QUESTION 17: According to this Wall Street Journal article, Facebook uses data collected from the Onavo suite of applications to monitor potentially competitive application⁶. Since the acquisition in 2013, how specifically has Facebook used information from Onavo to inform acquisitions as well as product development?

See Response to Question 16.

TERMS OF SERVICE:

QUESTION 18: Has Facebook ever disclosed to its users which “third parties partners” have access to user information? If no, will you publish this list so that users know which outside parties have access to their information?

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 platform. Users can view and edit the categories of information that apps they have used have access to through the App Dashboard.

USER TRACKING:

QUESTION 19: Does Facebook can “track a user’s Internet browsing activity, even after that user has logged off of the Facebook platform”? If yes, how Facebook discloses that kind of tracking to its users? And can users opt-out of this kind of tracking?

When people visit apps or websites that feature our technologies—like the Facebook Like or Comment button—our servers automatically log (i) standard browser or app records of the fact that a particular device or user visited the website or app (this connection to Facebook’s servers occurs automatically when a person visits a website or app that contains our technologies, such as a Like button, and is an inherent function of Internet design); and (ii) any additional information the publisher of the app or website chooses to share with Facebook about the person’s activities on that site (such as the fact that a purchase was made on the site). This is a standard feature of the Internet, and most websites and apps share this same information with multiple different third-parties whenever people visit their website or app. For example, the Senate Commerce Committee’s website shares information with Google and its affiliate DoubleClick and with the analytics company Webtrends. This means that, when a person visits the Committee’s website, it sends browser information about their visit to each one of those third

⁶ Seetharaman, Deepa and Morris, Betsy. “Facebook’s Onavo Gives Social-Media Firm Inside Peek at Rivals’ Users.” 13 August 2017. *Wall Street Journal*.

parties. More information about how this works is available at <https://newsroom.fb.com/news/2018/04/data-off-facebook/>.

When the person visiting a website featuring Facebook’s tools is not a registered Facebook user, Facebook does not have information identifying that individual, and it does not create profiles for this individual.

We use the browser and app logs that apps and websites send to us—described above—in the following ways for non-Facebook users. First, these logs are critical to protecting the security of Facebook and to detecting or preventing fake account access. For example, if a browser has visited hundreds of sites in the last five minutes, that’s a sign the device might be a bot, which would be an important signal of a potentially inauthentic account if that browser then attempted to register for an account. Second, we aggregate those logs to provide summaries and insights to websites and apps about how many people visit or use their product, or use specific features like our Like button—but without providing any information about a specific person. We do not create profiles for non-Facebook users, nor do we use browser and app logs for non-Facebook users to show targeted ads from our advertisers to them or otherwise seek to personalize the content they see. However, we may take the opportunity to show a general ad that is unrelated to the attributes of the person or an ad encouraging the non-user to sign up for Facebook.

When the individual is a Facebook user, we are also able to use this information to personalize their experiences on Facebook, whether or not they are logged out, but we will not target ads to users relying on this information unless the user allows this in their privacy settings. We do not sell or share this information with third-parties.

QUESTION 20: How many Facebook “Like” buttons there are on non-Facebook web pages?

Facebook does not publish tracking software. When people visit apps or websites that feature our technologies—like the Facebook Like or Comment button—our servers automatically log (i) standard browser or app records of the fact that a particular device or user visited the website or app (this connection to Facebook’s servers occurs automatically when a person visits a website or app that contains our technologies, such as a Like button, and is an inherent function of Internet design); and (ii) any additional information the publisher of the app or website chooses to share with Facebook about the person’s activities on that site (such as the fact that a purchase was made on the site).

This is a standard feature of the Internet, and most websites and apps share this same information with multiple different third-parties whenever people visit their website or app. For example, the Senate Commerce Committee’s website shares information with Google and its affiliate DoubleClick and with the analytics company Webtrends. This means that, when a person visits the Committee’s website, it sends browser information about their visit to each one of those third parties. More information about how this works is available at <https://newsroom.fb.com/news/2018/04/data-off-facebook/>.

During the week prior to April 16, 2018, on sites that use Facebook services: the Like button appeared on 8.4M websites, the Share button on 931K websites covering 275M webpages, and there were 2.2M Facebook pixels installed on websites.

QUESTION 21: How many Facebook “Share” buttons there are on non-Facebook web pages?

See Response to Question 20.

QUESTION 22: How many non-Facebook websites have Facebook pixel code?

See Response to Question 20.

QUESTION 23: While users can download their user generated data using the “Download Your Information” tool, how can users download data that Facebook has inferred about them?

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 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, will 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, delete this information from their account, and turn off Facebook’s ability to store it associated with their account going forward.

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.

QUESTION 24: How many websites have Facebook-tracking software on them? What percentage of all internet sites have Facebook-tracking software?

See Response to Question 20.

QUESTION 25: According to a Gizmodo report⁷, Facebook collects data on people using Shadow Profiles. Do you collect data on people who are not Facebook users? Please describe the process for non-Facebook users can employ to delete any data collected about them by the company.

⁷ Hill, Kasmir. 07 November 2017. How Facebook Figures Out Everyone You’ve Met. *Gizmodo*. <https://gizmodo.com/how-facebook-figures-out-everyone-youve-ever-met-1819822691?IR=T>.

Yes. If a person doesn't have a Facebook account but believes Facebook may have information about them, they can contact us to request a copy of your information. A contact form is available at <https://www.facebook.com/help/contact/180237885820953>. However, Facebook does not create profiles about or track web or app browser behavior of non-users.

When people visit apps or websites that feature our technologies—like the Facebook Like or Comment button—our servers automatically log (i) standard browser or app records of the fact that a particular device or user visited the website or app (this connection to Facebook's servers occurs automatically when a person visits a website or app that contains our technologies, such as a Like button, and is an inherent function of Internet design); and (ii) any additional information the publisher of the app or website chooses to share with Facebook about the person's activities on that site (such as the fact that a purchase was made on the site). This is a standard feature of the Internet, and most websites and apps share this same information with multiple different third-parties whenever people visit their website or app. For example, the Senate Commerce Committee's website shares information with Google and its affiliate DoubleClick and with the analytics company Webtrends. This means that, when a person visits the Committee's website, it sends browser information about their visit to each one of those third parties. More information about how this works is available at <https://newsroom.fb.com/news/2018/04/data-off-facebook/>.

When the person visiting a website featuring Facebook's tools is not a registered Facebook user, Facebook does not have information identifying that individual, and it does not create profiles for this individual.

We use the browser and app logs that apps and websites send to us—described above—in the following ways for non-Facebook users. First, these logs are critical to protecting the security of Facebook and to detecting or preventing fake account access. For example, if a browser has visited hundreds of sites in the last five minutes, that's a sign the device might be a bot, which would be an important signal of a potentially inauthentic account if that browser then attempted to register for an account. Second, we aggregate those logs to provide summaries and insights to websites and apps about how many people visit or use their product, or use specific features like our Like button—but without providing any information about a specific person. We do not create profiles for non-Facebook users, nor do we use browser and app logs for non-Facebook users to show targeted ads from our advertisers to them or otherwise seek to personalize the content they see. However, we may take the opportunity to show a general ad that is unrelated to the attributes of the person or an ad encouraging the non-user to sign up for Facebook.

Questions from Senator Peters

- 1. A major challenge artificial intelligence (AI) and machine learning developers need to address is the ability to ensure prolonged safety, security, and fairness of the systems. This is especially true of systems designed to work in complex environments that may be difficult to replicate in training and testing, or systems that are designed for significant learning after deployment. One approach to address this challenge is to implement standards or principles guiding the development of AI systems. However, you referenced AI more than 30 times in your testimony on Capitol Hill, and many of those references were in different contexts. This seems to imply Facebook has assumed a broad or vague definition of AI. I fear that a vague definition will make it difficult to implement clear, unambiguous standards or principles to guide the fair, safe, and secure application of AI and algorithms.**
 - a. What how does Facebook define AI?**
 - b. How is Facebook currently working to build trust in its usage of AI? Specifically, has your company developed a set of principles to guide your development and use of AI systems? If so, what are they? Please also provide details on how these principles are being implemented.**
 - c. How will these principles improve the transparency of decision-making AI systems?**
 - d. How will these principles prevent a system designed to learn after deployment from developing unacceptable behavior over time?**

We are focused on both the technical and the ethical aspects of artificial intelligence. We believe these two should go hand-in-hand together in order to fulfill our commitment to being fair, transparent and accountable in our development and use of AI. Facebook has AI teams working on developing the philosophical, as well as technical, foundations for this work. Facebook is also one of the co-founders and members of the Partnership on AI (PAI), a collaborative and multi-stakeholder organization established to study and formulate best practices on AI technologies, to advance the public's understanding of AI, and to serve as an open platform for discussion and engagement about AI and its influences on people and society. The thematic pillars that structure the work we're doing in the scope of the PAI—safety, fairness, transparency and accountability—are the principles that we believe industry should follow and promote when building and deploying AI systems. The PAI's Fair, Transparent and Accountable AI Working Group is also working alongside industry, academia, and civil society to develop best practices around the development and fielding of fair, explainable, and accountable AI systems.

We believe that over the long term, building AI tools is the scalable way to identify and root out most content that violates our policies. We are making substantial investments in building and improving these tools. We already use artificial intelligence to help us identify threats of real world harm from terrorists and others. For example, the use of AI and other automation to stop the spread of terrorist content is showing promise. Today, 99 percent of the ISIS and Al Qaeda related terror content we remove from Facebook is content we detect before

anyone in our community has flagged it to us, and in some cases, before it goes live on the site. We do this primarily through the use of automated systems like photo and video matching and text-based machine learning. We also use AI to help find child exploitation images, hate speech, discriminatory ads, and other prohibited content.

2. Mr. Zuckerberg, you said recently that Facebook is more like a government than a traditional company. Facebook is a community of over 2 billion people from every country in the world. You have also said you hope to grow the number of Facebook employees working on security of the user community to 20,000 by the end of the year. A city like Flint, Michigan has a population of 100,000 and roughly 100 uniformed police officers. Your company is aiming to have one cop on the beat for every 100,000 of its 2 billion users.

a. Is this going to be adequate to prevent another misuse of consumer data like we saw with Cambridge Analytica?

We are doubling the size of our security and content review teams (from 10,000 to 20,000) over the course of this year. We currently have approximately 15,000 people working on these teams.

b. How are you making the efforts of these employees transparent and accountable to your users?

We are taking significant steps to increase our transparency. For example, we have published the internal guidelines we use to enforce our Community Standards here: <https://newsroom.fb.com/news/2018/04/comprehensive-community-standards/>. We decided to publish these internal guidelines for two reasons. First, the guidelines will help people understand where we draw the line on nuanced 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—over time.

We also recently publicized data around enforcement of our Community Standards in a Community Standards Enforcement Report (<https://transparency.facebook.com/community-standards-enforcement>). The report details our enforcement efforts between October 2017 to March 2018, and it covers six areas: graphic violence, adult nudity and sexual activity, terrorist propaganda, hate speech, spam, and fake accounts. The numbers show you:

- How much content people saw that violates our standards;
- How much content we removed; and
- How much content we detected proactively using our technology—before people who use Facebook reported it.

The data we published is the same information we use to measure our progress internally. We believe this increased transparency will lead to increased accountability and responsibility over time.

- 3. Facebook has made some changes in light of the 2016 US Presidential election and the fact that your platform allowed for the proliferation of fake news. You've since developed tools that try to tamp down on this activity—pulling down fake accounts and destroying bots.**
- a. You have described the content on your platform during elections held since 2016, both foreign and domestic, as “cleaner”—but what metrics are you using to evaluate the real effectiveness of the changes you have made?**
 - b. Once you have a true understanding of the impact these tools have—how can you communicate the changes to users so they can be confident that what they are viewing is real and not there for the purpose of manipulating them?**
 - c. Consumers are skeptical of the content on your platform, how can you gain back their trust?**

We are working hard to regain the trust of our community.

Success would consist of minimizing or eliminating abuse of our platform and keeping our community safe. We have a number of specific goals that we will use to measure our progress in these efforts. First, we are increasing the number of people working on safety and security at Facebook, to 20,000. We have significantly expanded the number of people who work specifically on election integrity, including people who investigate this specific kind of abuse by foreign actors. Those specialists find and remove more of these actors. Second, we work to improve threat intelligence sharing across our industry, including, we hope, by having other companies join us in formalizing these efforts. This is a fight against sophisticated actors, and our entire industry needs to work together to respond quickly and effectively. Third, we are bringing greater transparency to election ads on Facebook by requiring more disclosure from people who want to run election ads about who is paying for the ads and by making it possible to see all of the ads that an advertiser is running, regardless of the targeting. We believe that these efforts will help to educate our community and to arm users, media, civil society, and the government with information that will make it easier to identify more sophisticated abuse to us and to law enforcement.

We have gotten increasingly better at finding and disabling fake accounts. We're now at the point that we block millions of fake accounts each day at the point of creation before they do any harm.

We are taking steps to help users assess the content they see on Facebook. For example, for ads with political content, we've created an archive that will hold ads with political content for seven years—including for information about ad impressions and spend, as well as demographic data such as age, gender and location. People in Canada and Ireland can already see all the ads that a Page is running on Facebook—and we're launching this globally in June. Further, advertisers will now need to confirm their ID and location before being able to run any ads with political content in the US. All ads with political content will also clearly state who paid for them. 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.

4. How did Facebook, prior to the 2016 US Presidential election, identify and evaluate fake or troll accounts, and how have your processes changed since then?

- a. What steps are taken once Facebook has identified fake or troll accounts and, specifically, how much of your response is consumer-facing? Will a user ever truly know the extent to which they were influenced by a fake account?**

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.

- b. Is it true that Facebook does not authenticate the administrators of group and organization pages in the same manner it authenticates individual accounts? Will you take a different approach going forward?**

We have 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 users additional context about Pages to effectively assess their content. For example, a user can see whether a Page has changed its name.

5. Current sector-specific privacy laws and state privacy laws, as well as currently proposed federal legislation that address data privacy and security, often narrowly define personal information to include identifiers like a person's name, social security number, and bank information. But definitions of personal information currently do not cover information like social media "likes" and certain choices and activities online that bad actors have at worst used to manipulate voters and at best used to deliver targeted advertisements.

- a. What do you think Cambridge Analytica has taught us about what should be considered personal information?**
- b. Should definitions of personal information be updated to include an individual's activities like search activity and social media "likes"?**

Facebook is generally not opposed to regulation but wants to ensure it is the right regulation. The issues facing the industry are complex, multi-faceted, and affect an important part of peoples' lives. As such, Facebook is absolutely committed to working with regulators,

like Congress, to craft the right regulations. Facebook would be happy to review any proposed legislation and provide comments.

6. Who do you consider to be Facebook’s customers (i.e., what stakeholders directly provide Facebook with revenue)? To the extent that the customers are not the end users of the platform, how will Facebook reconcile the privacy expectations and interests of both sets of stakeholders?

In the words of Facebook CEO and Founder Mark Zuckerberg, “Facebook is an idealistic and optimistic company. For most of our existence, we focused on all the good that connecting people can bring. As Facebook has grown, people everywhere have gotten a powerful new tool to stay connected to the people they love, make their voices heard, and build communities and businesses.” Our product is social media—the ability to connect users with the people that matter to them, wherever they are in the world. It’s the same with a free search engine, website or newspaper. The core product is reading the news or finding information—and the ads exist to fund that experience. Our priority is protecting our community, and that is more important than maximizing our profits.

7. Does Facebook intend to provide its users with a comprehensive listing of all apps and services that have accessed their Facebook data? In such a listing, would Facebook include information about which data points were accessed, when they were accessed, and how they were accessed?

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 platform. Users can view and edit the categories of information that apps they have used have access to through the App Dashboard.

8. What mechanisms does Facebook have in place to monitor third parties who have access to user data once the data is delivered? If a user deletes their data on Facebook, how does Facebook ensure that third parties with access to their data have also deleted it?

We use a variety of tools to enforce Facebook policies against violating parties, including developers. We review tens of thousands of apps per year. With the exception of Account Information (name, email, gender, birthday, current city, and profile picture URL), apps may maintain user data obtained from us only for as long as necessary for their business purpose and must delete the information if they stop using the Facebook Platform. Further, developers are required to keep the data maintained on their systems up to date.

9. What mechanisms—beyond self-reporting—are currently in place, or will be in place in the future, to enable independent academic and journalistic validation of Facebook’s

current and future claims that the platform has removed bad actors who have abused or compromised user data and privacy?

App Review. We have large teams of internal and external experts working hard to investigate these apps as quickly as possible. To date, thousands of apps have been investigated and around 200 (from a handful of developers) have been suspended—pending a thorough investigation into whether they did in fact misuse any data.

The App Review process introduced in 2014 requires developers who create an app that asks for more than certain basic user information from installers 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 it is approved following such review can the app ask for users' permission to get their data. Facebook has rejected more than half of the apps submitted for App Review between April 2014 and April 2018.

New Developer Requirements. We are in the process of investigating every app that had access to a large amount of information before we changed our Platform in 2014. If we find suspicious activity, we will take immediate steps to investigate (including a full forensic audit) or take enforcement actions against the app. If we determine that there has been improper use of data, we will ban those developers and notify everyone affected. Facebook is launching the Data Abuse Bounty to reward people who report any misuse of data by app developers. The Data Abuse Bounty, inspired by the existing bug bounty program that we use to uncover and address security issues, will help us identify violations of our policies.

Further, Facebook's Platform Policy makes clear to app developers the relevant requirements regarding users' privacy that apply to apps operating on the Platform, including the requirements to give users choice and control, and to respect user privacy. Application developers explicitly agree to Facebook's Statement of Rights and Responsibilities and Platform Policy when they set up their Facebook accounts. The Platform Policy imposes a variety of obligations on app developers regarding the features, functionality, data collection and usage, and content for apps on the Platform, as well as Facebook's right to take enforcement action if an application violates the Platform Policy.

Clear History. We have also worked with regulators, legislators, and privacy experts on updates that make data settings and tools easier to find. For example, we recently announced plans to build Clear History. This feature will enable users to see the websites and apps that send us information when they use them, delete this information from their accounts, and turn off our ability to store it associated with their accounts going forward. When developing tools such as Clear History, we will work 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 heard specific demands for controls like these at a session we held at our headquarters two weeks ago. We're looking forward to doing more.

Measuring Misinformation Through Academic Commission. In April, Facebook also announced a new initiative to help provide independent research about the role of social media in

elections, as well as democracy more generally. In the coming weeks, the commission will lead a request for proposals to measure the volume and effects of misinformation on Facebook. They will then manage a peer review process to select which scholars will receive funding for their research, and access to privacy-protected data sets from Facebook. This will help keep us accountable and track our progress over time.

Elections. We know that outside experts, researchers, and academics can also help by analyzing political advertising on Facebook. It's why we're working closely with our newly-formed Election Commission and other stakeholders to launch an API for the archive of ads with political content. We also recognize that news coverage of elections and important issues is distinct from advocacy or electoral ads, even if those news stories receive paid distribution on Facebook. We're working closely with news partners and are committed to updating the archive to help differentiate between news and non-news content.

Questions from Senator Duckworth

Question 1: According to the New York Times and other media outlets, fair housing advocates recently filed a lawsuit in federal court arguing that “Facebook continues to discriminate against certain groups, including women, disabled veterans and single mothers, in the way that it allows advertisers to target the audience for their ads.” Despite repeated announcements by Facebook suggesting that your company will remedy this disturbing practice, third-party organizations have tested your platform repeatedly to exclude certain minorities. Unfortunately, many of these tests of your platform were successful and this issue has been known to Facebook for several years.

Please explain in detail why Facebook provided housing advertisers with targeting options to exclude users based on “ethnic affinity” in clear violation of Federal law. Following third-party demonstrations of how a housing advertiser could unlawfully use Facebook to discriminate against certain protected classes of housing customers, please describe in detail the specific actions Facebook took to end the practice and make sure that Facebook’s user tools actually reflect Facebook’s written policies that claim to prohibit using Facebook’s targeting options to discriminate. As Chairman and Chief Executive Officer, please describe how you personally responded to the public reports demonstrating that Facebook’s targeting options had enabled unlawful discrimination in housing. Please provide any company documents, in hard copy or electronic form, addressing the implementation of Facebook advertising targeting options and any associated risk that such an option could result in violations of Federal legal prohibitions against discrimination in housing. If Facebook has no such documents, please provide a detailed justification as to why the company did not, or does not, have a compliance protocol or office dedicated to enforcing Fair Housing laws.

We want our advertising tools to help promote inclusion and diversity of all kinds. Discrimination has no place on Facebook, and we make this clear to advertisers in a number of ways. Everyone on Facebook must agree to our Terms when they sign up to use our service. In so doing, they agree not to engage in discriminatory conduct on Facebook. In addition, our Advertising Policies (available at <https://www.facebook.com/policies/ads/>) include an explicit and detailed anti-discrimination policy that prohibits discriminatory ads or the use of our audience selection tools for discriminatory purposes.

In late 2016, we began building machine learning tools (called “classifiers”) that were intended to automatically identify, at the point of creation, advertisements offering housing, employment or credit opportunities (referred to here generally as “housing, employment and credit ads”). We built these classifiers so that when we identified one of these kinds of ads, we could: (1) prevent the use of our “multicultural affinity” targeting options in connection with the ad, and (2) for the use of any other kind of targeting, require that the advertiser certify compliance with our anti-discrimination policy and applicable anti-discrimination laws.

We trained the classifiers before we launched them, including by using search terms provided by your office in January 2017. After the classifiers launched in approximately February 2017, we anticipated that, through machine learning, they would become better over time at distinguishing ads offering housing, employment, or credit opportunities from other types

of ads. We also expected that we would receive feedback about the performance of the tool that would enable us to detect problems and improve the classifiers over time.

In practice, the classifiers did not improve over time as much as we had anticipated. Rather, they became both over- and under-inclusive, identifying and requiring self-certification for hundreds of thousands of ads each day that may have had nothing to do with housing, employment, or credit offers, while missing ads that may have contained such offers.

There were two principal reasons for this failure. First, a key aspect of our ad-review process involves the random sampling of ads that are live on Facebook for the purpose of reassessing those ads' compliance with our Advertising Policies. When we identify ads that should have been flagged as being in violation of our policies, we use that information to improve our review processes, including our machine learning classifiers. In hindsight, our training set was not sufficiently comprehensive and did not include an evolving set of housing, credit and employment ads that should have been flagged by our classifiers to better train our models. We also failed to fully account for the lack of feedback we would likely receive about the performance of these classifiers through other channels—feedback we typically rely on to alert us to performance issues. For example, advertisers whose ads should have been (but were not) identified through this process would have had no reason to report a problem.

We take these limitations very seriously, and we regret that they prevented us from providing the oversight we had hoped to provide. Since they were brought to our attention in November 2017, we have taken significant steps to remedy them. These steps include the following:

- We have integrated all of the classifiers and targeting prohibitions into the random sampling process we use to gather feedback about the performance of our ad review processes.
- We are adding more than 1,000 people to our global ads review teams over the next year to allow for more human review of the ads placed on our platform.
- We have built teams whose role it is to pressure test our policy-enforcement products to identify potential performance issues.

In addition to addressing the issues with housing, employment and credit classifiers to more accurately identify such ads, as of January 2018, we have implemented the following additional changes with regard to multicultural affinity targeting more generally:

- We disabled the use of multicultural affinity exclusion targeting for all ads; this prohibition is no longer limited to housing, employment and credit ads.
- We now require self-certification of compliance with our anti-discrimination policies and applicable anti-discrimination laws for any use of multicultural affinity targeting, regardless of the type of ad.
- We have undertaken a review of our ad-targeting tools generally, with an eye toward identifying the potential for the tools to be abused.

- As a result of that review, we disabled the use of other exclusion targeting categories that we determined, on their face, may have been misunderstood to identify a group of Facebook users based on race, color, national origin or ancestry.

Question 2: What is Facebook doing to protect Veterans, women and other minorities to ensure that advertisements on your platform do not discriminate against them in possible violation of federal laws? Is Facebook aware of an investigation by the U.S. Department of Housing and Urban Development regarding these issues and is Facebook cooperating with an investigation? When were you alerted that an investigation(s) had begun? Do you believe that violators of Federal laws prohibiting discrimination, such as the protections contained in the Fair Housing Act, should be held accountable?

Discriminatory advertising has no place on Facebook's platform and Facebook removes such content as soon as it becomes aware of it. Facebook's policies prohibit advertisers from discriminating against people on personal attributes such as race, ethnicity, color, national origin, religion, age, sex, sexual orientation, gender identity, family status, disability, and medical or genetic conditions. Facebook educates advertisers on our anti-discrimination policy, and in some cases, requires the advertisers to certify compliance with Facebook's anti-discrimination policy and anti-discrimination laws.

Facebook also uses machine learning to help identify ads that offer housing, employment, or credit opportunities. When an advertiser attempts to show an ad that Facebook identifies as offering a housing, employment, or credit opportunity and includes Facebook's multicultural advertising segments, Facebook will disapprove the ad. Facebook also requires advertisers to certify that they are complying with Facebook's updated anti-discrimination policy and anti-discrimination laws when the advertiser attempts to show a housing, employment, or credit opportunity and uses any other audience segment on Facebook.

Facebook has been actively engaged with the US Department of Housing and Urban Development (HUD) since at least the Fall of 2016. As part of the engagement, Facebook has focused on addressing the concern that advertisers may seek to engage in discriminatory advertising on Facebook's platform. In connection with this engagement, Facebook has made numerous modifications and improvements to its ad policies, practices, and tools.

Question 3: I'm glad to hear that Facebook plans to extend the European Union's General Data Protection Regulations (GDPR) to U.S. users. By what date does Facebook plan on extending those protections to U.S. users? In doing so, is Facebook affirming that all data generated by a user is the property of that user and is subject to protections outlined in the General Data Protection Regulations, including rights to access, rectification, erasure, data portability, among others?

We confirm that we provide 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, and many of those tools (like our Download Your Information tool, Ad Preferences tool, and Activity Log) have been available globally for many years. We have recently begun providing direct notice of these controls and our updated terms of service to people around the world (including in the US), allowing people to choose whether or not to enable or disable these settings or to consent to our updated terms. The controls and settings that Facebook is enabling

as part of GDPR are available to people around the world, including settings for controlling our use of face recognition on Facebook and for controlling our ability to use data we collect off Facebook Company Products to target ads.

Question 4: The European Union’s deadline for full implementation of their General Data Protection Regulations (GDPR) is May 25, 2018. While you have said publically that Facebook plans to extend General Data Protection Regulations (GDPR) across its platform “in spirit,” including to users in the U.S., recent media reporting suggests that Facebook’s commitment to GDPR implementation across its platform is questionable. In your view, what does implementation of GDPR “in spirit” mean? If Facebook were to be found violating GDPR protections for non-European Union users, what recourse do those users have, legal or otherwise, to remedy a complaint?

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 GDPR include settings for controlling our use of face recognition on Facebook and for controlling our ability to use data we collect off Facebook Company Products to target ads. We recently began providing 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 consent to our updated terms. Many of these tools (like our Download Your Information tool, ad preferences tool, and Activity Log) have been available globally for many years.

The substantive protections in our user agreements offered by Facebook Ireland 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.

Facebook is subject to ongoing oversight by the Federal Trade Commission with respect to its privacy commitments to people and its implementation of privacy settings, under a Consent Order with the FTC. Facebook is subject to the authority of the Irish Data Protection Commissioner, its lead regulator, under the GDPR in the European Union.

Question 5: As reported by Politico on April 17, 2018, Facebook has enlisted the help of conservative organizations to push back against GDPR and other potential regulatory efforts in the U.S. Is Facebook coordinating with political organizations to consider or address potential state or federal regulatory actions?

When the GDPR was finalized, we realized it was an opportunity to invest even more heavily in privacy. We not only wanted to comply with the law, but also go beyond our obligations to build new and improved privacy experiences for everyone on Facebook. To that end, as we often do, we sought feedback from people with a variety of perspectives on privacy,

including people who use our services, regulators and government officials, privacy and policy experts, and designers. We are applying the same protections, controls, and transparency to people in the US and around the world that we are providing to people in Europe under GDPR.

[Removed by Committee due to inadvertent submission of unrelated material]

[Removed by Committee due to inadvertent submission of unrelated material]

Questions from Senator Cortez Masto

Children's Data

Question 1. Does Instagram have an age limit requirement similar to the 13 years old Facebook requires?

Yes, Instagram requires everyone to be at least 13 years old before they can create an account (and in some jurisdictions, this age limit may be higher).

Question 2. How vulnerable or widely utilized have children's (18 or younger) data been in both Facebook and your other platforms?

We take the privacy, safety, and security of all those who use our platform very seriously, and when it comes to minors (13 to 18 years old), we provide special protections and resources.

We also provide special protections for teens on Facebook and Messenger. We provide education before allowing teens to post publicly. We don't show search results based on specific profile data (high school, birthday/age, and hometown, or current city) of teens to unconnected adults when the adults search on Facebook. Unconnected adults can't message minors who are 13-17. And, we prohibit search engines off Facebook from indexing minors' profiles. And, we have age limits for advertisements. For example, ads for dating sites, financial services, and other products or services are gated to users under 18.

We provide special resources to help ensure that they enjoy a safe and secure experience. For example, we recently announced the launch of our Youth Portal, which is available in 60 languages at <https://www.facebook.com/safety/youth>. This portal is a central place for teens that includes:

- **Education.** Information on how to get the most out of products like Pages, Groups, Events, and Profile, while staying safe. Plus, information on the types of data Facebook collects and how we use it.
- **Peer Voices.** First-person accounts from teens around the world about how they are using technology in new and creative ways.
- **Ways to control user experience.** Tips on things like security, reporting content, and deciding who can see what teens share.
- **Advice.** Guidelines for how to safely get the most out of the internet.

Instagram also will be providing information to teens to show them where they can learn about all of the tools on Instagram to manage their privacy and stay safe online, including how to use the new Access and Download tools to understand what they have shared online and learn how to delete things they no longer want to share. We are also making this information available in formats specifically designed for young users, including video tutorials for our privacy and safety tools, and teen-friendly FAQs about the Instagram Terms of Use, Data Policy, safety features, and Community Guidelines.

Instagram has also launched new content on Instagram Together, including videos and FAQs about privacy controls; information on how to use safety features, including comment controls, blocking accounts, reporting abuse, spam, or troubling messages; information on responsible social media use; and FAQs about safety on Instagram. We will be reaching out to users under 18 on Instagram to encourage them to learn more on Instagram Together, available at <https://www.instagram-together.com/>.

Further, we have content restrictions and reporting features for everyone, including minors. We have Community Standards that prohibit hate speech, bullying, intimidation, and other kinds of harmful behavior. We encourage people to report posts and rely on our team of content reviewers around the world to review reported content. Our reviewers are trained to look for violations and enforce our policies consistently and as objectively as possible. When reviewed by our team, we hide certain graphic content from users under 18 (and include a warning for adults). We are also working to improve our ability to get our community help in real time, especially in instances where someone is expressing thoughts of suicide or self-harm, by expanding our use of proactive detection, working with safety experts and first-responders, and dedicating more reviewers from our Community Operations team.

Question 3. How many children (18 or younger) had their data taken during the Cambridge Analytica breach?

The Children's Online Privacy Protection Act (COPPA) requires parental consent and notification in specific instances involving the collection and use of data about children under the age of 13. Facebook does not allow children under the age of 13 on its service or collect data about children under 13 that would trigger parental consent or notification.

Question 4. Are you notifying parents about their children's exposed data?

See Response to Question 3.

Discriminatory Advertising

Question 1. Please provide a detailed description, including screenshots if applicable, of the nondiscrimination compliance certification that Facebook currently requires advertisers to complete.

Please refer to our letter to you dated May 16, 2018.

Question 2. Please provide a complete list of the characteristics, categories, descriptors, and/or interests that Facebook allows advertisers to select in order to target certain users for inclusion in an advertisement's audience.

Please refer to our letter to you dated May 16, 2018. Please note, however, that in limited cases and for the purpose of running ads that are not related to housing, employment or credit, we are re-enabling the ability of advertisers to exclude people from their audiences based on family status but are reviewing this as a targeting option.

Question 3. Please provide a complete list of the characteristics, categories, descriptors, and/or interests that Facebook allows advertisers to select in order to exclude certain users from an advertisement's audience.

See Response to Question 2.

Question 4. Are there any characteristics, categories, descriptors, and/or interests that Facebook had previously permitted advertisers to select, but that Facebook no longer allows to be selected as targeting or exclusion criteria? If so, please provide a complete list of those characteristics, categories, descriptors, and/or interests.

See Response to Question 2.

Question 5. Are there certain characteristics, categories, descriptors, and/or interests that Facebook has never allowed advertisers to select for the purpose of targeting or excluding users from an advertisement's audience? If so, please provide a complete list of those characteristics, categories, descriptors, and/or interests.

See Response to Question 2.

Question 6. Please describe the process that Facebook uses to determine whether a characteristic, category, descriptor, or interest will be available for selection as a targeting or exclusion criteria. If Facebook has a written policy governing this determination, please provide a copy.

We have Community Standards that prohibit hate speech, bullying, intimidation, and other kinds of harmful behavior. We hold advertisers to even stricter advertising policies to protect users from things like discriminatory ads. We don't want advertising to be used for hate or discrimination, and our policies reflect that. For example, our Advertising Policies make it clear that advertisers may not discriminate against people based on personal attributes such as race, ethnicity, color, national origin, religion, age, sex, sexual orientation, gender identity, family status, disability, and medical or genetic condition. The Policies also prohibit asserting or implying that a person belongs to one of these groups.

We educate advertisers on our anti-discrimination policy, and when we detect that an advertiser is attempting to run a housing, employment or credit ad, we require the advertiser to certify compliance with our anti-discrimination policy and anti-discrimination laws. We are committed to getting better at enforcing our advertising policies. We review many ads proactively using automated and manual tools, and reactively when people hide, block, or mark ads as offensive. We are taking aggressive steps to strengthen both our automated and our manual review. We are also expanding our global ads review teams and investing more in machine learning to better understand when to flag and take down ads, such as ads that use our multicultural affinity segments in connection with offers of housing, employment or credit opportunities.

Question 7. Regardless of whether the characteristics are described as demographic, behavioral, or interest-based criteria, does Facebook allow employment, housing, credit

advertisements to be targeted to users on the basis of protected characteristics, including race, national origin, religion, sex, gender, disability, age, and familial status?

See Response to Question 2.

Question 8. Regardless of whether the characteristics are described as demographic, behavioral, or interest-based criteria, does Facebook allow advertisers for employment and housing to exclude users on the basis of protected characteristics, including race, national origin, religion, sex, gender, disability, age, and familial status?

See Response to Question 2.

Question 9. Has Facebook reviewed characteristics/categories available for advertising to select or exclude when targeting that can be used as “proxies” for protected characteristics? If so, what is Facebook’s policy regarding the continued availability of that characteristic as a targeting or exclusion criteria and has Facebook ever removed categories that were being used as “proxies” for protected categories? How does Facebook go about determining which such categories could potentially be used as “proxies” for discrimination?

See Response to Question 6.

Question 10. Does Facebook allow employment, housing, and credit advertisements to be targeted to users on the basis of categories that may be reasonable proxies for protected characteristics?

See Response to Question 2.

Question 11. Does Facebook allow employment, housing, and credit advertisements to be targeted to users on the basis of their sexual orientation or gender identity?

Please refer to our letter to you dated May 16, 2018.

Question 12. In Facebook’s December 20, 2017 press release, Rob Goldman, VP of Ads, wrote that Facebook “proactively look[s] for bad ads, and investigate[s] concerns when they are raised.” Please describe Facebook’s process for monitoring ads for possible violations of Title VII, the Fair Housing Act, the Americans with Disabilities Act, and Title II of the Genetic Information Nondiscrimination Act.

Please refer to our letter to you dated May 16, 2018.

Question 13. Does Facebook “proactively look” for ads that may be discriminatory on the basis of each protected characteristic before they are posted to the platform?

Please refer to our letter to you dated May 16, 2018.

Question 14. Does Facebook have defined, written policies for determining whether an employment, housing, or credit ad is discriminatory on the basis of each protected

characteristic, and a procedure for deleting such ads? If so, please provide copies of such policies.

Please refer to our letter to you dated May 16, 2018.

Question 15. Has Facebook ever proactively deleted an employment, housing, or credit ad on the grounds that it discriminated on the basis of a protected characteristic? If so, how many such ads has Facebook deleted, broken down by each protected characteristic?

Please refer to our letter to you dated May 16, 2018.

Question 16. Has Facebook ever deleted an employment, housing, or credit ad on the grounds that it discriminated on the basis of a protected characteristic in response to a user complaint? If so, how many such ads has Facebook deleted, broken down by each protected characteristic?

Please refer to our letter to you dated May 16, 2018.

Question 17. Has Facebook ever barred a businesses or ad companies from using its services because of discriminatory ads? How many? Please detail the process Facebook has for addressing discriminatory advertisers, once identified.

Please refer to our letter to you dated May 16, 2018.

Question 18. Many state and local nondiscrimination laws go further than federal statutes prohibiting discrimination against protected classes. Does Facebook require advertisers to certify that they will comply with state and local nondiscrimination laws?

Please refer to our letter to you dated May 16, 2018.

Question 19. Does Facebook “proactively look” at employment, housing, and credit ads to evaluate their compliance with state and local nondiscrimination laws?

Please refer to our letter to you dated May 16, 2018.

Question 20. Does Facebook respond to user complaints about employment, housing, and credit ads that may violate state and local nondiscrimination laws? If so, how?

Please refer to our letter to you dated May 16, 2018.

Question 21. Please provide a timeline and any relevant documentation of interactions with the U.S. Department of Housing and Urban Development on Facebook’s advertisement policies.

Please refer to our letter to you dated May 16, 2018.

Question 22. Please provide a detailed description of any other U.S. federal agencies that have contacted Facebook regarding the issue of discriminatory advertising on the Facebook platform.

We regularly work cooperatively with regulators that may have questions about our platform and are happy to answer questions.

Question 23. Please describe when this contact took place and a detailed description of the agency's inquiry and interaction with Facebook, as well as Facebook's response.

See Response to Question 22.

Question 24. Will Facebook commit to having an outside entity conducting a Civil Rights Audit of its platform and advertising practices? If so, will Facebook commit to meaningfully consulting civil rights organizations on the perimeters of the Civil Rights Audit? Will Facebook commit to making the results of such audit accessible to the public?

Relman, Dane & Colfax, a respected civil rights law firm, will carry out a comprehensive civil rights assessment of Facebook's services and internal operations. Laura Murphy, a national civil liberties and civil rights leader, will help guide this process—getting feedback directly from civil rights groups, like The Leadership Conference on Civil and Human Rights, and help advise Facebook on the best path forward.

Discrimination and Diversity in Tech Community

Over the past few months, our country has been reckoning with some hard truths about the way that women and minorities are treated in the workplace. And I think this is a moment for all types of organizations, including tech giants like the one represented here, to take a clear-eyed accounting of their culture and practices, to take responsibility for what hasn't worked, and to renew their commitments to make meaningful improvements. The Equal Employment Opportunity Commission's 2016 report on "Diversity in High Tech" found that that women, African Americans, and Hispanics are all represented at significantly lower levels in high tech than in private industry as a whole. And while recent internal studies by you at Facebook, and Google, have showed some progress in the hiring of women, there has not been equal improvement in the representation of people of color and other underrepresented groups.

Question 1. What does diversity mean to you, and how do you want it reflected in your operations?

With a global community of over two billion people on Facebook, greater diversity and inclusivity are critical to achieving our mission. Studies have shown that cognitive diversity on teams that are working on hard problems produces better results. Diversity helps us build better products, make better decisions and better serve our community. In order to achieve that, we have developed programming to attract and retain more people from traditionally underrepresented groups which include women, people of color, veterans and people with disabilities.

We are not where we would like to be, but we are encouraged that representation for people from underrepresented groups at Facebook has increased. We've grown Black and

Hispanic representation by 1 percent each (2 percent combined) between our first report in 2014 and our most recent report in 2017:

- Black Representation: from 2 percent to 3 percent
- Hispanic Representation: from 4 percent to 5 percent
- Black Non-Tech: from 2 percent to 6 percent
- Hispanic Non-Tech: from 6 percent to 8 percent
- Black Leadership: from 2 percent to 3 percent
- Hispanic Leadership: from 3 percent to 4 percent
- Black and Hispanic Tech have stayed at 1 percent and 3 percent

As of August 2017, the number of women globally increased from 33 percent to 35 percent:

- Women in Tech: from 17 percent to 19 percent
- Women in Non-Tech: from 47 percent to 55 percent
- Women in Leadership: from 23 percent to 28 percent
- Women made up 27 percent of all new graduate hires in engineering and 21 percent of all new technical hires at Facebook.

We seek to promote diversity in a variety of ways, and we want to highlight three programs in particular. First, we have adopted our Diverse Slate Approach (DSA) to interviewing job candidates. The more people that hirers interview who don't look or think like them, the more likely they are to hire someone from a diverse background. To hardwire this behavior at Facebook, we introduced our DSA in 2015 and have since rolled it out globally. DSA sets the expectation that hiring managers will consider candidates from underrepresented backgrounds when interviewing for an open position.

Second, we are working to reduce unconscious bias. Our publicly available Managing Unconscious Bias class encourages our people to challenge and correct bias as soon as they see it—in others, and in themselves. We've also doubled down by adding two additional internal programs: Managing Inclusion, which trains managers to understand the issues that affect marginalized communities, and Be The Ally, which gives everyone the common language, tools, and space to practice supporting others.

Third, we have created Facebook University. We want to increase access and opportunity for students with an interest in software engineering, business, and analytics. Facebook University (FBU) gives underrepresented students extra training and mentorship earlier in their

college education. We started FBU in 2013 with 30 students and expect to have 280 in 2018. More than 500 students have graduated from this program, with many returning to Facebook for internships and full-time jobs.

Finally, we have many partnerships to move the numbers nationally such as Black Girls Code, All Star Code, Hack the Hood, The Hidden Genius Project, Level Playing Field Institute, Yes We Code, Streetcode Academy, Dev Color, Dev Bootcamp and Techbridge. And, we now recruit at 300 Universities—including historically black colleges and universities (HBCUs) like Spelman, Morehouse, Howard, NCA&T, and Morgan State (EIR) and the HBCU Faculty Summit.

We're committed to building a more diverse, inclusive Facebook. Much like our approach to launching new products on our platform, we are willing to experiment and listen to feedback.

Question 2. How are your entities working to address issues of discrimination, or lack of diversity, in your own workforce?

See Response to Question 1.

Question 3. Do you believe those efforts are sufficient and what do you believe is needed throughout the tech sector to address the mistreatment of some, and the need to expand ladders of opportunities for everyone?

See Response to Question 1.

Like most companies, Facebook files numerous patents on its emerging technology and I'd like to raise concerns about some of the patents that your company has recently filed.

One is titled "[Socioeconomic group classification based on user features](#)" which is technology that would allow Facebook to group users into upper, middle, and working classes based on user action. It was recently discovered that Facebook has allowed advertisers to discriminate on the base of age.

Question 1. How can we be confident that your company will crack down on discriminatory behavior as it is developing technology to group users into class?

Discriminatory advertising has no place on Facebook's platform and Facebook removes such content as soon as it becomes aware of it. Facebook's policies prohibit advertisers from discriminating against people on personal attributes such as race, ethnicity, color, national origin, religion, age, sex, sexual orientation, gender identity, family status, disability, and medical or genetic conditions. Facebook educates advertisers on our anti-discrimination policy, and in some cases, requires the advertisers to certify compliance with Facebook's anti-discrimination policy and anti-discrimination laws.

Facebook also uses machine learning to help identify ads that offer housing, employment, or credit opportunities. When an advertiser attempts to show an ad that Facebook identifies as

offering a housing, employment, or credit opportunity and includes Facebook's multicultural advertising segments, Facebook will disapprove the ad. Facebook also requires advertisers to certify that they are complying with Facebook's updated anti-discrimination policy and anti-discrimination laws when the advertiser attempts to show a housing, employment, or credit opportunity and uses any other audience segment on Facebook.

Question 2. What other uses could this patent possibly have?

See Response to Question 1.

Equal Pay Day

Mr. Zuckerberg, the date you appeared before was Equal Pay Day in America, which symbolizes the number of extra days a typical woman who works full-time, year-round must work into this year to be paid what a typical white man got paid. Women are still only paid 80 cents on the dollar compared to men. It's estimated that women employed full time in the in the U.S. will lose nearly \$900 billion to the wage gap this year. I'm passionate about getting underrepresented folks into the job opportunities that our tech revolution provides, and equal pay goes along with creating those ladders of opportunities.

Question 1. Is this an issue you are aware of and active on within your operations?

At Facebook, women and men receive equal pay for equal work and have done so for many years. This is an absolute minimum standard for a diverse business such as ours and we continually review our hiring and compensation practices to ensure this remains the case. Compensation at Facebook is made up of base salary, cash bonus or commission, and equity in the company. We work hard to avoid unconscious bias affecting how much people get paid. Managers don't make decisions about compensation increases—instead, we use a formulaic approach that determines pay based on performance and level.

Opportunities for advancement and leadership within the company are also crucial. For our women employees, we run a series of development workshops and training programs designed to provide a strong network of support, along with the tools they need to be the best leaders they can be across different levels in the company. We hold ourselves accountable because this matters to us. In 2017, the number of women employees globally rose from 33 percent to 35 percent and the number of women in technical roles increased from 17 percent to 19 percent. Between 2014 when we first publicly reported our representation data and 2017, the number of women in leadership roles has increased from 23 percent to 28 percent.

We are committed to increasing the representation of women at all levels. We know we're not where we need to be, and we're committed to making real progress.

With a global community of over two billion people on Facebook, greater diversity and inclusivity are critical to achieving our mission. Studies have shown that cognitive diversity on teams that are working on hard problems produces better results. Diversity helps us build better products, make better decisions, and better serve our community.

Question 2. Can you provide us confirmation, including figures, that your pay for women matches their male counterparts?

See Response to Question 1.

Question 3. And that you appropriately compensate all of your employees based on their job title and value to the company?

See Response to Question 1.

Facebook's self-regulation of Campaign and Issue Ads & the Honest Ads Act

You recently announced that political ads run on Facebook are now going to be subject to heightened transparency requirements, such as including disclaimers stating who paid for the ad, and making it easier for viewers to see the ads that a page is running. I think this is a good first step but there are several questions I have regarding its implementation and how you will enforce this new policy:

Question 1. What if you have an organization, let's call them "XYZ," who wants to post an issue or political ad, but they have never filed reports with the FEC, they are not registered with the IRS as a nonprofit, and they don't appear to have a website?

We now require more thorough documentation from advertisers who want to run ads with political content. Any person who wants to run one of these ads must upload an identification document and provide the last four digits of their Social Security number. They also must prove residency in the US by providing a residential mailing address. Once they provide the address, we mail a letter with a code that the person must provide to us in order to become authorized to run ads with political content.

Question 2. You have said that advertisers running political ads and issue ads will have to be "authorized," and that Facebook will confirm their identity and location before running ads. What does it mean to "confirm their identity?"

See Response to Question 1.

Question 3. Walk me through how this ad would be treated under Facebook's new policies.

See Response to Question 1.

Question 4. So, this ad will say "paid for by XYZ." But there is no public record of XYZ, besides the fact that they have a Facebook page. Would you let a mysterious group like this run an ad on Facebook without any further information about who they are?

See Response to Question 1.

Question 5. Will you require any further verification from this group?

See Response to Question 1.

Question 6. Will these transparency measures you are discussing tell you who paid the Facebook page to run the ad? In other words, will Facebook disclose the sources of funding for these political ads?

Once verified as described above in response to Question 1, these advertisers will have to include a disclosure in these ads, which reads: “Paid for by.” When users click on the disclosure, they will be able to see details about the advertiser. These ads will also all appear in a searchable archive, available at www.facebook.com/politicalcontentads, which includes information about how much the advertiser spent on the ad, how many people saw it, and general demographic information about the people who saw it.

Question 7. What if a foreign government gave money to a Facebook page with a U.S. address to run political ads? Would you tell that to viewers?

These are real challenges and reflect problems largely outside our control, but we will continue to work to improve our enforcement of ads that violate our policies.

Question 8. What if a foreign government gave money to a Facebook page through a series of shell companies or LLCs?

See Response to Question 7.

Question 9. How will Facebook know who the real donors to this group are?

See Response to Question 7.

Question 10. How is Facebook defining a “political ad” and an “issue ad” subject to these heightened transparency measures?

Our Political Advertising Policy (https://www.facebook.com/policies/ads/restricted_content/political) applies to any ad that:

- Is made by, on behalf of or about a current or former candidate for public office, a political party, a political action committee or advocates for the outcome of an election to public office;
- Relates to any election, referendum or ballot initiative, including “get out the vote” or election information campaigns;
- Relates to any national legislative issue of public importance in any place where the ad is being run; or
- Is regulated as political advertising.

We further define “national legislative issue of public importance” as including twenty issues. Ads that take a position on one or more of these issues are covered by the policy. To develop this initial list (which we expect to evolve over time), we worked with the non-partisan Comparative Agendas Project and many other stakeholders from across the political spectrum.

We determine whether an ad is subject to our Political Advertising policy based on the content of the ad.

Question 11. Is the “political ad/issue ad” determination based on the content of a particular ad, or the identity of the advertiser running the ad, or some other criteria?

See Response to Question 10.

Facebook sells several types of ads, including sponsored ads that appear directly in a user’s newsfeed, and smaller ads that appear on the right column. Studies show that a large volume of political ads from the 2016 election ran in the right column rather than in a user’s newsfeed.

Question 12. Will all types of ads sold by Facebook, including smaller ads, be subject to these heightened transparency measures?

Yes, all ads with political content will be subject to this policy.

Question 13. You mentioned that the disclaimers Facebook is going to implement will say which Facebook page paid for the ad. Will it tell you exactly what organization or individual is behind that page?

We require the advertiser to disclose who paid for an ad with political content—regardless of whether that is an individual or an organization.

Question 14. Rob Goldman, the Vice President of Ads at your company, indicated that you are working with the “third parties” to develop these parameters. Who are these “third parties?”

See Response to Question 11.

Question 15. Will these ad transparency measures also apply to state and local elections?

Our Political Advertising policy applies to all advertisers running ads with political content. The products we have launched (authorization, disclaimer, and archive) are available to all advertisers running ads with political content to users in the US.

Question 16. Will these same measures apply to other platforms owned by Facebook, like Instagram?

Yes, the measures will apply to ads with political content shown on Instagram.

New Employees-- Content Review

In your testimony, you note that Facebook plans to hire an additional 5,000 workers for its security and content review teams, for a total of 20,000 workers by the end of this year. But Facebook first announced the plan for a 20,000 person security team in late October of last year, in response to concerns about Russian interference in the election.

Question 1. Given the additional revelations about the role of Cambridge Analytica and other third party apps in compromising the privacy and personal information of at least 87 million users, do you still believe 20,000 is the appropriate level of staffing for Facebook's security team?

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.

We are doubling the size of our security and content review teams (from 10,000 to 20,000) over the course of this year. We currently have approximately 15,000 people working on these teams.

Of that 15,000, more than 7,500 people review content around the world.

- Our content review team is global and reviews reports in over 50 languages.
- Reports are reviewed 24 hours a day, 7 days a week and the vast majority of reports are reviewed within 24 hours.
- Our goal is always to have the right number of skilled people with the right language capabilities to ensure incoming reports are reviewed quickly and efficiently.
- We hire people with native language and other specialist skills according to the needs we see from incoming reports.
- The team also includes specialists in areas like child safety, hate speech and counter-terrorism, software engineers to develop review systems, quality control managers, policy specialists, legal specialists, and general reviewers.

We are also using machine learning to better detect and action on content and people that should not be using our platform.

For example, we incorporated learnings from interference in previous elections to better detect and stop false accounts from spreading misinformation in more recent elections.

We recently shared how we are using machine learning to prevent bad actors like terrorists or scammers from using our platform (<https://www.facebook.com/notes/facebook-security/introducing-new-machine-learning-techniques-to-help-stop-scams/10155213964780766/>).

We employ a mix of full-time employees, contractors and vendor partners to assist with content review and help us scale globally.

We partner with reputable vendors who are required to comply with specific obligations, including provisions for resiliency, support, transparency, and user privacy.

Question 2. Will these new security and content review workers be direct employees of Facebook, or do you plan to outsource this work to third party entities?

See Response to Question 1.

Question 3. If the security review work is outsourced, how will Facebook vet those contractors, subcontractors, and employees and where will those employees be located?

See Response to Question 1.

Question 4. And how can Facebook assure its users that there will be transparency and accountability for any future breaches of privacy if the company is outsourcing its security work?

See Response to Question 1.

Future Facebook Technology

One of your recent patent is titled “[Dynamic eye tracking calibration](#)” and another is called “[Techniques for emotion detection and content delivery](#)”. The patent for the eye tracking technology says that “the (eye) calibration process is performed automatically in the background while the user uses a device.” The second patent would use a device’s camera to monitor your emotions and “display content based upon a received emotion type.”

Question 1. How does Facebook plan to use this technology?

Like many companies, we apply for a wide variety of patents to protect our intellectual property. Right now we’re not building technology to identify people with eye-tracking cameras. However, we’re always exploring how new technologies and methods can improve our services, and eye-based identification is one way that we could potentially reduce consumer friction and add security for people when they log into Oculus or access Oculus content.

If we implement this technology in the future, we will absolutely do so with people’s privacy in mind, just as we do with movement information (which we anonymize in our systems).

As we continue to develop new virtual reality products and services, we’re committed to being transparent and open about the information that we collect and how we use it, as well as any ways that changes over time.

Question 2. Will users be fully aware that their eyes and emotions are being tracked?

See Response to Question 1.

Question 3. Is Facebook confident it has the proper data security in place to have this intimate level of data on users?

See Response to Question 1.

Question 4. Facebook has reportedly been developing an in-home digital assistant similar to products like Alexa, will this also be tracking this granular level of data?

See Response to Question 1.

Question 5. The second patent says that content will be delivered on a person's perceived emotion type. Couldn't this be potentially dangerous in amplifying hateful messages?

See Response to Question 1.

Question 6. If a person focuses on an image of say, a propaganda image of immigrants, will this technology deliver more of this content?

See Response to Question 1.

China's Facebook Access

In July 2009, the Chinese government blocked Facebook in China. The precise reason for that action remains obscure, but it fits into an overall pattern. The Chinese government is unwilling to allow a social media platform—foreign or domestic—to operate in China unless it agrees to abide by Chinese law. First, a social media platform must agree to censor content and conversations in line with directives from China's information authorities. And second, businesses that collect data from Chinese individuals can only store that data in China where, presumably, it would be easier for the Chinese government to access, via legal means or otherwise. You've made no secret of your desire to see Facebook available once again in China.

Question 1. Could you please reveal to the committee whether you are willing to agree to either of these requirements?

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.

Question 2. And will Facebook pledge to guarantee its future Chinese users the same level of privacy protection it gives its users in the U.S. and the European Union?

Everyone in the world deserves good privacy protection.

Consent Agreement

Question 1. The FTC consent agreement with Facebook requires an independent, biennial audit of Facebook’s privacy controls—when exactly have those audits been conducted, and what were the results?

To date, three independent privacy assessments prepared by PwC have been completed and submitted to the FTC: a 180-Day Assessment (dated April 16, 2013), a biennial privacy assessment covering the period between February 12, 2013 and February 11, 2015 (dated April 13, 2015), and a biennial privacy assessment covering the period between February 12, 2015 and February 11, 2017 (dated April 12, 2017). In each of these assessments, PwC determined that Facebook’s privacy controls were operating with sufficient effectiveness to protect the privacy information covered under the FTC Consent Order.

Question 2. Did Facebook inform any of its auditors of the Cambridge Analytica data leak? Did any of Facebook’s auditors know about the Cambridge Analytic data leak?

Facebook routinely undertakes internal and external reviews, including undergoing biennial assessments under Facebook’s consent agreement with the Federal Trade Commission, which focus on the functioning of privacy controls that are part of Facebook’s privacy program. As a part of the assessments, our independent assessors (PwC) have onsite access to our personnel and records, and we provide them with such access to information and personnel as they request in order to perform their work. PwC is also permitted to conduct a number of tests to determine whether the privacy controls in place under our privacy program—including controls relating to developer’s access to information—are working properly. In its capacity as independent assessor, PwC evaluates the sufficiency of our controls through independent testing and requesting information that we provide to conduct that evaluation. Their focus is on evaluating the operation and sufficiency of our controls, rather than specific incidents.

Kogan’s violation of Facebook’s Platform Policies was widely reported at the time Facebook learned about it, including reporting in *The Guardian* on December 11, 2015, which reported that Kogan and his company, GSR, may have passed information Kogan’s app had obtained from Facebook users to SCL Elections Ltd. No data was transferred to Kogan’s app unless it was authorized by the users who installed his app, so there was not a data leak from Facebook’s systems. However, based on public reports and testimony, it appears that Kogan may have improperly transferred data to Cambridge Analytica in violation of our policies.

Question 3. Does Facebook choose which policies and procedures the auditors look at? Please explain in detail how these policies and procedures are chosen? Does the 3rd party auditor have any say on what policies and procedures are examined? Does the FTC have any input on how an audit is structured?

Facebook’s privacy assessments are conducted pursuant to the July 27, 2012 Consent Order. They 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 incorporated GAPP principles in designing its privacy

program and related controls, which are considered industry leading principles for protecting the privacy and security of personal information. Facebook provided the FTC with summaries of these controls and engaged extensively with the FTC regarding the structure of its privacy program. Facebook has submitted copies of each assessment to the FTC.

Question 4. Will Facebook commit to making the entirety of PwC audit submitted to the Federal Trade Commission in 2017 public? If not, please describe in detail why.

The privacy assessments conducted by PwC contain both Facebook's and PwC's sensitive business information that are confidential in order to prevent competitive harm and to ensure the integrity of Facebook's privacy program, including the steps that we take to protect people's information. We have furnished these reports to the FTC and are prepared to review the reports with regulators and lawmakers with appropriate assurances that confidential information or information that could be exploited to circumvent Facebook's privacy protections will not be disclosed publicly.

Question 5. During the negotiations with the FTC in 2011, were you asked by them to remove the capability to expose friends from having their data utilized without their direct permission?

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. The Consent Order memorializes the agreement between Facebook 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 apps they used. Facebook voluntarily changed this feature of Platform in 2014, however.

It is worth noting that in 2011, Facebook offered more control and protection over the availability of friends data to apps than any other digital platform at the time, including mobile app platforms, which generally permitted apps to access user data and their friends' data without consent or any control. By contrast, Facebook notified users of each category of data an app could access—including friends data—before the user consented to the app, and also provided all users with controls that would prevent their friends from sharing their data with apps on Facebook's platform.

Hospital Data Sharing Project

It was reported by CNBC on April 5 that your company was in talks with top hospitals and other medical groups as recently as March 2018 about a proposal to share data you possess with the patients. As of now, the project is reportedly "on hiatus" so that Facebook can do a better job of protecting individuals' data.

Question 1. Please provide us the specific privacy concerns Facebook has with compiling your users' data with medical data possessed by the hospitals?

The medical industry has long understood that there are general health benefits to having a close-knit circle of family and friends. But deeper research into this link is needed to help

medical professionals develop specific treatment and intervention plans that take social connection into account. With this in mind, last year Facebook began discussions with leading medical institutions, including the American College of Cardiology and the Stanford University School of Medicine, to explore whether scientific research using fully-anonymized Facebook data could help the medical community advance our understanding in this area. This work has not progressed past the planning phase, and we have not received, shared, or analyzed anyone's data.

In March, we decided that we should pause these discussions so we can focus on other important work, including doing a better job of protecting people's data and being clearer with them about how that data is used in our products and services.

Question 2. Would you share any internal documents that led Facebook to put this project on hiatus?

See Response to Question 1.

Data Details & FB Messenger Data

Based on the FTC-Facebook consent order, your company collects a great deal of personal information on its users including—the location (e.g., city or state), age, sex, birthday, “Interested in” responses (i.e. whether a user is interested in men or women), Relationship Status, Likes and Interests, Education (e.g., level of education, current enrollment in high school or college, affiliation with a particular college, and choice of major in college), and name of employer of individuals.

Question 1. Do you collect any other specific information you have on individual Facebook users?

In response to your specific questions, depending on which Services a person uses, we collect different kinds of information from or about them. This is described in our Data Policy:

- **Things you and others do and provide.** Information and content you provide. We collect the content, communications and other information you provide when you use our Products, including when you sign up for an account, create or share content, and message or communicate with others. This can include information in or about the content you provide (like metadata), such as the location of a photo or the date a file was created. It can also include what you see through features we provide, such as our camera, so we can do things like suggest masks and filters that you might like, or give you tips on using camera formats. Our systems automatically process content and communications you and others provide to analyze context and what's in them for the purposes described below. Learn more about how you can control who can see the things you share.
 - **Data with special protections:** You can choose to provide information in your Facebook profile fields or Life Events about your religious views, political views, who you are “interested in,” or your health. This and other

information (such as racial or ethnic origin, philosophical beliefs or trade union membership) could be subject to special protections under the laws of your country.

- **Networks and connections.** We collect information about the people, Pages, accounts, hashtags, and groups you are connected to and how you interact with them across our Products, such as people you communicate with the most or groups you are part of. We also collect contact information if you choose to upload, sync or import it from a device (such as an address book or call log or SMS log history), which we use for things like helping you and others find people you may know and for the other purposes listed below.
- **Your usage.** We collect information about how you use our Products, such as the types of content you view or engage with; the features you use; the actions you take; the people or accounts you interact with; and the time, frequency and duration of your activities. For example, we log when you're using and have last used our Products, and what posts, videos and other content you view on our Products. We also collect information about how you use features like our camera.
- **Information about transactions made on our Products.** If you use our Products for purchases or other financial transactions (such as when you make a purchase in a game or make a donation), we collect information about the purchase or transaction. This includes payment information, such as your credit or debit card number and other card information; other account and authentication information; and billing, shipping and contact details.
- **Things others do and information they provide about you.** We also receive and analyze content, communications and information that other people provide when they use our Products. This can include information about you, such as when others share or comment on a photo of you, send a message to you, or upload, sync, or import your contact information.
- **Device Information.** As described below, we collect information from and about the computers, phones, connected TVs and other web-connected devices you use that integrate with our Products, and we combine this information across different devices you use. For example, we use information collected about your use of our Products on your phone to better personalize the content (including ads) or features you see when you use our Products on another device, such as your laptop or tablet, or to measure whether you took an action in response to an ad we showed you on your phone on a different device.

Information we obtain from these devices includes:

- **Device attributes:** information such as the operating system, hardware and software versions, battery level, signal strength, available storage space, browser type, app and file names and types, and plugins.

- **Device operations:** information about operations and behaviors performed on the device, such as whether a window is foregrounded or backgrounded, or mouse movements (which can help distinguish humans from bots).
- **Identifiers:** unique identifiers, device IDs, and other identifiers, such as from games, apps or accounts you use, and Family Device IDs (or other identifiers unique to Facebook Company Products associated with the same device or account).
- **Device signals:** Bluetooth signals, and information about nearby Wi-Fi access points, beacons, and cell towers.
- **Data from device settings:** information you allow us to receive through device settings you turn on, such as access to your GPS location, camera, or photos.
- **Network and connections:** information such as the name of your mobile operator or ISP, language, time zone, mobile phone number, IP address, connection speed and, in some cases, information about other devices that are nearby or on your network, so we can do things like help you stream a video from your phone to your TV.
- **Cookie data:** data from cookies stored on your device, including cookie IDs and settings. Learn more about how we use cookies in the Facebook Cookies Policy (available at <https://www.facebook.com/policies/cookies/>) and Instagram Cookies Policy (available at <https://www.instagram.com/legal/cookies/>).
- **Information from partners.** Advertisers, app developers, and publishers can send us information through Facebook Business Tools they use, including our social plug-ins (such as the Like button), Facebook Login, our APIs and SDKs, or the Facebook pixel. These partners provide information about your activities off Facebook—including information about your device, websites you visit, purchases you make, the ads you see, and how you use their services—whether or not you have a Facebook account or are logged into Facebook. For example, a game developer could use our API to tell us what games you play, or a business could tell us about a purchase you made in its store. We also receive information about your online and offline actions and purchases from third-party data providers who have the rights to provide us with your information. Partners receive your data when you visit or use their services or through third parties they work with. We require each of these partners to have lawful rights to collect, use, and share your data before providing any data to us.

Question 2. Are you tracking and collecting information and data from within your messenger chat tool? If so, what specific data are you collecting?

See Response to Question 1.

Question 3. What about your other platforms, like Instagram, what type of data are you tracking there?

Our Instagram Data Policy describes the data we collect and is available at <https://help.instagram.com/519522125107875>.

Question 4. Are you preserving broad and full conversations?

See Response to Questions 1 and 3.

Question 5. Is that something you would have available to provide law enforcement?

We reach out to law enforcement if we learn of content that we believe reflects a credible threat of imminent harm. We have been able to provide support to authorities around the world that are responding to the threat of terrorism, including in cases where law enforcement has been able to disrupt attacks and prevent harm. Further, as part of official investigations, government officials sometimes request data about people who use Facebook. We have strict processes in place to handle these government requests, and we disclose account records in accordance with our terms of service and applicable law. We publish more information on the standards that govern our release of information to law enforcement in our Law Enforcement Guidelines at <https://www.facebook.com/safety/groups/law/guidelines/>, and release statistics on the frequency with which we receive and comply with law enforcement requests at <https://transparency.facebook.com/>.

Data Protection on Facebook

Question 1. Has Facebook ever launched a feature that had to be turned off because of the privacy implications?

Protecting people's information is at the heart of everything we do, and as our CEO has recently stated, we are serious about doing what it takes to protect our community. We have developed extensive systems and processes that are designed to protect our data and user data, to prevent data loss, to disable undesirable accounts and activities on our platform, and to prevent or detect security breaches. In addition to comprehensive privacy reviews, we put products through rigorous data security testing. We also meet with regulators, legislators, and privacy experts around the world to get input on our data practices and policies.

At Facebook, we make decisions about privacy through a cross-functional, cross-disciplinary effort 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 is responsible for reviewing product launches, major changes, and privacy-related bug fixes to products and features to ensure that privacy policies and procedures are consistently applied and that key privacy decisions are implemented for the product. This approach has several key benefits.

- First, it is designed to consider privacy early in the product development process. This allows us to consider the benefits that a feature is intended to have for people who use our services, how data will be used to deliver those benefits, and how we can build features from the ground up that include privacy protections to enable those benefits while protecting people’s information and putting them in control.
- Second, while complying with our obligations is critically important, taking a cross-disciplinary approach to privacy encourages us to think about data protection as more than just a compliance exercise. Instead, we evaluate how to design privacy into the features that we build, and consider this from the perspective of things like how we design interfaces that make data use intuitive, taking a consistent approach to privacy across our services, and building protections in how our software is engineered. Accordingly, while we scale our privacy review process depending on the complexity of a particular data use, reviews typically involve experts who evaluate proposed data practices from the perspective of multiple disciplines.

As part of our consent agreement with the Federal Trade Commission, we submit a report to the FTC every two years. That report is based on assessments conducted by an independent third party on a biennial basis, which require us to submit evidence to demonstrate the effectiveness of the program.

Question 2. If so, how many times has that happened, and how many users were impacted?

See Response to Question 1.

Question 3. Did you notify the users who were impacted?

See Response to Question 1.

Facebook tracking software:

Question 1. How many websites have Facebook tracking software on them?

Facebook does not publish tracking software. When people visit apps or websites that feature our technologies—like the Facebook Like or Comment button—our servers automatically log (i) standard browser or app records of the fact that a particular device or user visited the website or app (this connection to Facebook’s servers occurs automatically when a person visits a website or app that contains our technologies, such as a Like button, and is an inherent function of Internet design); and (ii) any additional information the publisher of the app or website chooses to share with Facebook about the person’s activities on that site (such as the fact that a purchase was made on the site). This is a standard feature of the Internet, and most websites and apps share this same information with multiple different third-parties whenever people visit their website or app. For example, the Senate Commerce Committee’s website shares information with Google and its affiliate DoubleClick and with the analytics company Webtrends. This means that, when a person visits the Committee’s website, it sends browser information about their visit to each one of those third parties. More information about how this works is available at <https://newsroom.fb.com/news/2018/04/data-off-facebook/>.

During the week prior to April 16, 2018, on sites that use Facebook services, the Like button appeared on 8.4 million websites, the Share button on 931,000 websites covering 275 million webpages, and there were 2.2 million Facebook pixels installed on websites.

Question 2. What percentage of all internet sites have Facebook tracking software?

See Response to Question 1.

Question 3. Do you track users even when they are logged out from Facebook?

When people visit apps or websites that feature our technologies—like the Facebook Like or Comment button—our servers automatically log (i) standard browser or app records of the fact that a particular device or user visited the website or app (this connection to Facebook’s servers occurs automatically when a person visits a website or app that contains our technologies, such as a Like button, and is an inherent function of Internet design); and (ii) any additional information the publisher of the app or website chooses to share with Facebook about the person’s activities on that site (such as the fact that a purchase was made on the site). (See <https://www.facebook.com/policies/cookies>). This is a standard feature of the Internet, and most websites and apps share this same information with multiple different third-parties whenever people visit their website or app. For example, the Senate Commerce Committee’s website shares information with Google and its affiliate DoubleClick and with the analytics company Webtrends. This means that, when a person visits the Committee’s website, it sends browser information about their visit to each one of those third parties. More information about how this works is available at <https://newsroom.fb.com/news/2018/04/data-off-facebook/>.

When the person visiting a website featuring Facebook’s tools is not a registered Facebook user, Facebook does not have information identifying that individual, and it does not create profiles for this individual.

We use the browser and app logs that apps and websites send to us in the following ways for non-Facebook users. First, these logs are critical to protecting the security of Facebook and to detecting or preventing fake account access. For example, if a browser has visited hundreds of sites in the last five minutes, that’s a sign the device might be a bot, which would be an important signal of a potentially inauthentic account if that browser then attempted to register for an account. Second, we aggregate those logs to provide summaries and insights to websites and apps about how many people visit or use their product, or use specific features like our Like button—but without providing any information about a specific person. We do not create profiles for non-Facebook users, nor do we use browser and app logs for non-Facebook users to show targeted ads from our advertisers to them or otherwise seek to personalize the content they see. However, we may take the opportunity to show a general ad that is unrelated to the attributes of the person or an ad encouraging the non-user to sign up for Facebook.

When the individual is a Facebook user, we are also able to use this information to personalize their experiences on Facebook, whether or not they are logged out, but we will not target ads to users relying on this information unless the user allows this in their privacy settings. We do not sell or share this information with third parties.

Question 4. Do you collect data on people who have chosen not to use Facebook?

See Response to Question 3.

Question 5. How is this data used?

See Response to Question 3.

Question 6. Does it inform a user's "interests" on Facebook?

See Response to Question 3.

Question 7. If it does inform a user's "interests", was any of the data collected passively from users while they were browsing sites outside of Facebook passed to Cambridge Analytica?

No. Kogan's app did not have access to advertising interests data or browser logs.

Question 8. When the option or opportunity was previously available for folks to get the user data of individuals' friends, what was the total pool of data points one could obtain of friends, or was it all the exact same?

In April 2014, we announced that we would more tightly restrict our platform 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. A small number of developers asked for and were granted short-term extensions beyond the one-year transition period, the longest of which lasted several months. These extensions ended several years ago. A transition period of this kind is standard when platforms implement significant changes to their technology base and was necessary here to avoid disrupting the experience of millions of people. New apps that launched after April 30, 2014 were required to use our more restrictive platform APIs, which incorporated several key new elements, including:

- Institution of a review and approval process, called App Review (also called Login Review), for any app seeking to operate on the new platform that would request access to data beyond the user's own public profile, email address, and a list of friends of the user who had installed and authorized the same app;
- Generally preventing new apps on the new platform from accessing friends data without review; and
- Providing users with even more granular controls over their permissions as to what categories of their data an app operating on the new platform could access.

Our investigation is ongoing and as part of it we are taking a close look at applications that had access to friends data under Graph API v.1.0 before we made technical changes to our platform to change this access.

Question 9. Why did you change the policy of getting access to friends back in 2015?

See Response to Question 8.

Quality Assurance - Policy changes within the company

Question 1. What kind of privacy review is required to make a change to the Facebook platform?

See Response to Question 1 under “Data Protection on Facebook.”

Question 2. Is this review of platform changes mandatory? If so, when did that level of review become mandatory?

See Response to Question 1 under “Data Protection on Facebook.”

Question 3. Before that level of review was required, what checks were in place to ensure that new features wouldn’t adversely impact users’ privacy?

See Response to Question 1 under “Data Protection on Facebook.”

Question 4. What level of employee seniority was required of employees to approve a launch of such a privacy-impacting feature? For example, have you ever let an intern make changes that impact people’s privacy?

See Response to Question 1 under “Data Protection on Facebook.”

The Cambridge Analytica Data

Question 1. Given the confessions made in undercover clips, and the means by which Cambridge Analytica obtained and used Facebook data, would you ever allow them broad access to your platform’s user data again?

No. Facebook banned Cambridge Analytica from our service. We understand that the company is now defunct.

Question 2. Do you believe they have violated the Federal Trade Commission Act and its broad prohibition against “unfair and deceptive acts and practices” by misrepresenting the terms of their Facebook app?

Facebook has not violated the Federal Trade Commission Act. Facebook is not in a position to determine whether third-party app developers violated the Act and leaves that determination to the FTC, although we can confirm that misrepresenting the terms of an app to users is a violation of Facebook’s developer policies.

Question 3. Previously, would you request an app developer or academic researcher outline any contractual or other association with outside entities - such as foreign nationals or states, or other potentially dangerous private operations? Are you doing so now?

In November 2013, when Kogan’s app first became active on the platform, apps generally could be launched on the Facebook Platform without affirmative review or approval by Facebook. Kogan’s app used the Facebook Login service, which allowed users to utilize their Facebook credentials to authenticate themselves to third-party services. Facebook Login and Facebook’s Graph API also allowed Kogan’s app to request permission from its users to access certain categories of data that users had entered into their Facebook profiles, as well as certain data their friends had shared with them, if enabled by these friends’ privacy settings.

The App Review process introduced in 2014 requires developers who create an app that asks for more than certain basic user information from installers 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 users’ 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 in the process of investigating every app that had access to a large amount of information before we changed our Platform in 2014. Where we have concerns, we will conduct interviews, make requests for information (RFI)—which ask a series of detailed questions about the app and the data it has access to—and perform audits that may include on-site inspections. If we determine that there has been improper use of data, we will ban those developers and notify everyone affected. Facebook is launching the Data Abuse Bounty to reward people who report any misuse of data by app developers. The Data Abuse Bounty, inspired by the existing bug bounty program that we use to uncover and address security issues, will help us identify violations of our policies.

Further, Facebook’s Platform Policy makes clear to app developers the relevant requirements regarding users’ privacy that apply to apps operating on the Platform, including the requirements to give users choice and control, and to respect user privacy. Application developers explicitly agree to Facebook’s Statement of Rights and Responsibilities and Platform Policy when they set up their Facebook accounts. The Platform Policy imposes a variety of obligations on app developers regarding the features, functionality, data collection and usage, and content for apps on the Platform, as well as Facebook’s right to take enforcement action if an application violates the Platform Policy.

Prior to the introduction of App review in 2014, the Facebook Platform Policy, included provisions to the following effect:

- **Give People Control: Section 2(8):** Delete all of a person’s data you have received from us (including friend data) if that person asks you to
- **Protect Data: Section 3(3):** Only use friend data (including friends list) in the person’s experience in your app.
- **Protect Data: Section 3(10):** Don’t transfer any data you receive from us (including anonymous, aggregate, or derived data) to any ad network, data broker or other advertising or monetization-related service.

- **Login: Section 7(4):** Request only the data and publishing permission your app needs.

The Platform Policy also outlined the actions Facebook could take for violations of the policy:

- **Things You Should Know: Section 6(8):** We can audit your app to ensure it is safe and does not violate our terms. If requested, you must provide us with proof that your app complies with our terms.
- **Things You Should Know: Section 6(15):** We may enforce against your app or web site if we conclude that your app violated our terms or is negatively impacting the Platform. We may or may not notify you in advance.
- **Things You Should Know: Section 6(16):** Enforcement is both automated and manual, and can include disabling your app, restricting you and your app's access to Platform functionality, requiring that you delete data, terminating agreements with you or any other action we deem appropriate.

Question 4. Do you know exactly how much Kogan profited from the data he provided to Cambridge Analytica and any other entities?

GSR certified to Facebook that it received payments totaling approximately 750,000 GBP from SCL for services relating to Kogan's modeling and use of data gathered by his app. The certification also stated that Kogan used the proceeds to operate GSR. Recently, Kogan has stated publicly that the above payment came from SCL. Kogan has also recently testified to the UK Parliament that GSR received additional payments not reflected in his certification to Facebook.

Question 5. From your understanding, was Kogan on payroll with Cambridge Analytica when he ran the personality app on Facebook?

Kogan has testified that he was not on Cambridge Analytica's payroll when he shared data and provided services to Cambridge Analytica. Rather, Kogan testified that he owned GSR, which entered into an agreement with Cambridge Analytica to provide it with services relating to certain Facebook data.

Question 6. Did Facebook make any attempt to pro-actively contact the 87 million users you say had their data harvested by Cambridge Analytica in the more than two years after you were alerted to the breach? If not, why not?

When Facebook learned about Kogan's breach of Facebook's data use policies in December 2015, we took immediate action. We retained an outside firm to assist in investigating Kogan's actions, to demand that Kogan and each party he had shared data with delete the data and any derivatives of the data, and to obtain certifications that they had done so. Because Kogan's app could no longer collect most categories of data due to changes in Facebook's platform, the company's highest priority at that time was ensuring deletion of the data that

Kogan may have accessed before these changes took place. With the benefit of hindsight, we wish we had notified people whose information may have been impacted. Facebook has since notified all people potentially impacted with a detailed notice at the top of their News Feed.

Question 7. Why did Facebook hire Joseph Chancellor, who was the business partner of Aleksandr Kogan, around the same time as the Guardian article alerted you to the violation of your policies?

Mr. Chancellor is a quantitative researcher on the User Experience Research team at Facebook, whose work focuses on aspects of virtual reality. We are investigating Mr. Chancellor's prior work with Kogan through counsel.

Question 8. Why do you continue to employ him to this day?

See Response to Question 7.

Question 9. Did any of the Facebook employees who were embedded with the Trump presidential campaign have any sense that they were helping target ads with data that was obtained through these disreputable means?

While our investigation is ongoing, our review indicates that Facebook employees did not identify any issues involving the improper use of Facebook data in the course of their interactions with Cambridge Analytica during the 2016 US Presidential campaign. No one from Facebook was assigned full time to the Trump campaign.

Question 10. Is there no way any red flags would have arisen from how either good the targeting data was, or the way they were using it?

We expect that advertisers will use targeted advertising, and many political campaigns use custom audiences. The fact that a campaign used a custom audience and the performance of that audience would not normally be an indicator of any wrongdoing.

Question 11. To your knowledge, what foreign actors or entities may have accessed the same level of data that Cambridge Analytica has utilized?

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.

Russia

Facebook has downplayed the reach of Russian advertising during the 2016 election.

But the company's main business model is based on giving ads and posts prominence in the feeds of well-targeted users.

Question 1. Has Facebook performed any analyses that looks at smaller groups of people and how effective those ads were against targeted groups? If so, can Facebook share that information?

We learned from press accounts and statements by congressional leaders that Russian actors might have tried to interfere in the 2016 election by exploiting Facebook’s ad tools. This is not something we had seen before, and so we started an investigation. We found that fake accounts associated with the IRA spent approximately \$100,000 on around 3,500 Facebook and Instagram ads between June 2015 and August 2017. Our analysis also showed that these accounts used these ads to promote the roughly 120 Facebook Pages they had set up, which in turn posted more than 80,000 pieces of content between January 2015 and August 2017. The Facebook accounts that appeared tied to the IRA violated our policies because they came from a set of coordinated, inauthentic accounts. We shut these accounts down and began trying to understand how they misused our platform. We shared the ads we discovered with Congress, in a manner that is consistent with our obligations to protect user information, to help government authorities complete the vitally important work of assessing what happened in the 2016 election.

Question 2. Do your company’s records show that Russia-backed ads and posts reached a higher number of people in certain states or regions of the United States?

Approximately 25 percent of the ads that we’ve identified and turned over to the committee were geographically targeted to a region smaller than the United States. The ads (along with the targeting information) are publicly available at <https://democrats-intelligence.house.gov/facebook-ads/social-media-advertisements.htm>.

Question 3. If so, how responsive were Facebook users in those targeted regions to the Russian posts and ads?

Below is an overview of our analysis to date of the IRA’s ads:

- Impressions (an “impression” is how we count the number of times something is on screen, for example this can be the number of times something was on screen in a person’s News Feed)
 - 44 percent of total ad impressions were before the US election on November 8, 2016.
 - 56 percent of total ad impressions were after the election
- Reach (the number of people who saw a story at least once):
 - We estimate 11.4 million people in the US saw at least one of these ads between 2015 and 2017.
- Ads with zero impressions:

- Roughly 25 percent of the ads were never shown to anyone. That's because advertising auctions are designed so that ads reach people based on relevance, and certain ads may not reach anyone as a result.
- Amount spent on ads:
 - For 50 percent of the ads, less than \$3 was spent.
 - For 99 percent of the ads, less than \$1,000 was spent.
 - Many of the ads were paid for in Russian currency, though currency alone is a weak signal for suspicious activity.
- Content of ads:
 - Most of the ads appear to focus on divisive social and political messages across the ideological spectrum, touching on topics from LGBT matters to race issues to immigration to gun rights.
 - A number of the ads encourage people to follow Pages on these issues, which in turn produced posts on similarly charged subjects.

We estimate that roughly 29 million people were served content in their News Feeds directly from the IRA's 80,000 posts over the two years. Posts from these Pages were also shared, liked, and followed by people on Facebook, and, as a result, three times more people may have been exposed to a story that originated from the Russian operation. Our best estimate is that approximately 126 million people may have been served content from a Page associated with the IRA at some point during the two-year period. This equals about four-thousandths of one percent (0.004%) of content in News Feed, or approximately 1 out of 23,000 pieces of content. While our data on Instagram is less complete, we believe another 16 million saw the IRA's Instagram posts starting in October 2016. Prior to that time, when our data is less incomplete, we believe another 4 million people may have seen this content.

Question 4. When did anyone at Facebook become aware that Russians or other foreign nationals were running ads in connection with the election?

See Response to Question 1.

Question 5. What happened with that information and what was done?

See Response to Question 1.

FEC

Question 1. Has anyone raised or approached you about potential infractions of any election laws that obtaining or using Facebook's data might be linked to including Cambridge Analytica's use of Facebook data?

We have a compliance team that trains our sales representatives to comply with all federal election law requirements in this area. We also have processes designed to identify inauthentic and suspicious activity and we also maintain a sanctions compliance program to screen advertisers and paid app developers. Facebook's denied party screening protocol involves checking paid app developers and advertisers against applicable denied party listings. Those screened remain in an on-going monitoring portfolio and are screened against changes to applicable denied party listings. Moreover, our payments subsidiaries file Suspicious Activity Reports on developers of certain apps as appropriate. However, like other offline and online companies, Facebook has limited insight into the use of shell corporations or other sophisticated structures that may disguise the true buyer. In addition, the general challenge of attributing online activities to specific governments or organizations is widely recognized in the intelligence and law enforcement communities.

Question 2. We are now learning that there is reason to believe that Cambridge Analytica and its foreign national employees participated in the decision making of its U.S. political committee clients, possibly in violation of our campaign finance law.⁸ What steps will you take to determine whether the companies behind political or issue ads posted on Facebook are not in violation of our campaign finance laws?

See Response to Question 1.

Question 3. Will you undergo this examination before these ads are allowed to be posted on your platform?

See Response to Question 1.

Technological Capabilities or Limitations to Protecting Data

Question 1. Is it fair to say that not only were you not vigilant in following up or tracking those who have assessed Facebook's data, but that you have no technical solutions to track data activity once it's outside your network, such as specialty whether it's properly deleted?

We use a variety of tools to enforce Facebook policies against violating parties, including developers. We review tens of thousands of apps per year and regularly disapprove noncompliant apps as part of our proactive review process. We also use tools like cease and desist letters, account suspensions, letter agreements, and civil litigation. For example, since 2006, Facebook has sent over 1,150 cease-and-desist letters to over 1,600 targets. In 2017, we took action against about 370,000 apps, ranging from imposing certain restrictions to removal of the app from the platform. Moreover, we have required parties who have procured our data without authorization to delete that data. We have invested significant resources in these efforts. Facebook is presently investigating apps that had access to large amounts of information before we changed our platform policies in 2014 to significantly reduce the data apps could access. To date, around 200 apps (from a handful of developers: Kogan, AIQ, Cube You, the Cambridge Psychometrics

⁸ <http://fortune.com/2018/03/26/watchdog-alleges-cambridge-analytica-violated-election-law/>

Center, and myPersonality) have been suspended pending a thorough investigation into whether they did in fact misuse any data.

Question 2. Or at least without a formal deep audit?

See Response to Question 1.

Question 3. What are the specific aspects of a formal audit, including the technical capabilities?

With respect to our audit of all apps that had access to large amounts of information before we changed our platform policies in 2014, where we have concerns that data may have been shared outside the app in violation of our policies, we will conduct interviews, make requests for information (RFI)—which ask a series of detailed questions about the app and the data it has access to—and perform audits that may include on-site inspections.

Question 4. And still with an audit, can you clarify what level of detail you have or could find misuse from someone?

See Response to Question 3.

It's being reported, and opined by others in your field, including former employees of yours, that it's notoriously difficult to track down and secure personal information once it has been unleashed.

Question 1. So that makes it all the more important to be vigilant on the front end, no?

In April 2014, we announced that we would more tightly restrict our platform 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. A small number of developers asked for and were granted short-term extensions beyond the one-year transition period, the longest of which lasted several months. These extensions ended several years ago. A transition period of this kind is standard when platforms implement significant changes to their technology base and was necessary here to avoid disrupting the experience of millions of people. New apps that launched after April 30, 2014 were required to use our more restrictive platform APIs. We required apps seeking additional categories of data to undergo proactive review by our internal teams. We rejected more than half of the apps seeking these permissions, including the second version of Kogan's app.

We review apps to ensure that the requested permissions clearly improve the user experience and that the data obtained is tied to an experience within the app. We conduct a variety of manual and automated checks of applications on the platform for Policy compliance, as well as random sampling. 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.

Recently, we announced a number of additional steps we're taking to address concerns raised by Kogan's app.

- **Review our platform.** We will investigate all apps that had access to large amounts of data before the platform changes we announced in 2014, and we will audit any app where we identify suspicious activity. If we identify misuses of data, we'll take immediate action, including banning the app from our platform and pursuing legal action if appropriate.
- **Tell people about data misuse.** We will tell people about apps that have misused their data. This includes building a way for people to know if their data might have been accessed via the app. Moving forward, if we remove an app for misusing data, we will tell everyone who used it.
- **Turn off access for unused apps.** If someone has not used an app within the last three months, we will turn off the app's access to their data.
- **Restrict Facebook Login data.** We are changing Login, so that the only data that an app can request without app review will include name, profile photo, and email address. Requesting any other data will require approval from Facebook. We will 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've permitted those apps to use. But we're making it easier for people to see what apps they use and the information they have shared with those apps.
- **Reward people who find vulnerabilities.** We launched the Data Abuse Bounty program so that people can report to us any misuses of data by app developers.
- **Update our policies.** We have updated our terms and Data Policy to explain how we use data and how data is shared with app developers.

Question 2. How much do you anticipate Facebook will be investing in your investigations or audits into app developers, and others who have had access to user data?

How much value would you estimate that Facebook has lost through this latest string of controversies, and the Cambridge Analytica data security issue?

We are investing so much in security that our costs will increase significantly. But we want to be clear about what our priority is: protecting our community is more important than maximizing our profits.

Question 3. And how much personally to do suspect you've lost?

See Response to Question 2.

Question 4. What personal data of yours, or say your wife's, is available or exploitable on any of the platforms you run?

Mark Zuckerberg's data was among the data that was shared with Kogan's app, which may have been improperly shared with Cambridge Analytica.

Question 5. Seems like millions, or even billions, spent earlier and being proactively protective would, or could have, saved tens of billions overall, wouldn't you agree?

See Response to Question 1.

Question 6. Do you think there's enough accountability at all levels within Facebook, including for yourself, Ms. Sandberg, others in senior positions?

As our CEO Mark Zuckerberg has said, when you are building something unprecedented like Facebook, there are going to be mistakes. What people should hold us accountable for is learning from the mistakes and continually doing better—and, at the end of the day, making sure that we're building things that people like and that make their lives better.

The [Washington Post has reported](#) that Mr. Kogan says that none of the data that was taken for research purposes in 2013 was provided to Cambridge Analytica. He says that after he began working with Cambridge Analytica, he sent out a new survey to Facebook users, with new terms of service that allowed for broad uses of the data. That new survey app collected data from nearly 300,000 Facebook users and captured data on 30 million of their friends. He says he has deleted all the data that he obtained from Facebook.

Question 1. Can Facebook prove all of this as fact or fiction?

On December 11, 2015, *The Guardian* published an article reporting that Kogan and his company, GSR, may have passed information his app had obtained from Facebook users to SCL Elections Ltd. (SCL)/Cambridge Analytica. By doing so, Kogan and his company violated Facebook's Platform Policies, which explicitly prohibited selling or licensing 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 his app from our platform and launched an investigation into these allegations. Kogan signed a certification declaring that he had deleted all data that he obtained through his app and obtained certifications of deletion from others he had shared data with, including Cambridge Analytica. In March 2018, new allegations surfaced that Cambridge Analytica may not have deleted data as it had represented. Our investigation of these matters is ongoing.

Facebook's Definition and Regulatory Positions

Question 1. Do you believe you are an actual media entity now?

Facebook does not create or edit the content that users share on its Platform, although we do take steps to arrange, rank and distribute that content to those who are most likely to be

interested in it, or to remove objectionable content from our service. These activities are protected functions under Communications Decency Act Section 230 and the First Amendment.

Question 2. Are you solely a tech company?

We are, first and foremost, a technology company. Facebook does not create or edit the content that our users published on our platform. While we seek to be a platform for a broad range of ideas, we do moderate content in good faith according to published community standards in order to keep users on the platform safe, reduce objectionable content, and to make sure users participate on the platform responsibly.

Question 3. When it comes to news posts and political advertising, why should Facebook get a regulatory exemption that traditional media doesn't get?

Facebook is committed to transparency for all ads, including ads with political 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 about elections or issues that are being debated across the country. In order to be authorized by Facebook, advertisers will need to confirm their identity and location. Furthermore, all political ads will include a disclosure in their election-related ads, which reads: "Paid for by," and when users click on this disclosure they will be able to see details about the advertiser. Users will also be able to see an explanation of why they saw the particular ad. This is similar to the disclosure included on political TV advertisements.

Facebook with Law Enforcement

Question 1. How wide is the use and specific collection of social media data with law enforcement, say in a given year? (FBI, CBP, ICE)

As part of official investigations, government officials sometimes request data about people who use Facebook. We have strict processes in place to handle these government requests, and we disclose account records in accordance with our terms of service and applicable law. We also have law enforcement response teams available around the clock to respond to emergency requests. As part of our ongoing effort to share information about the requests we have received from governments around the world, Facebook regularly produces a Transparency Report about government requests to Facebook.

Question 2. Have you seen an increase in such request under the current Administration?

See Response to Question 1.

Question 3. Or has there been a variation in the type or aggressiveness of these requests over the same time?

See Response to Question 1.

Social Media Addiction

Obvious the social media revolution has brought in a number of addition issues into play that we in Congress need to consider, from platforms for terrorist organizations and hate groups, to censorship and online addiction. And that is something I wanted to inquire about.

***Question 1.* I know it was raised by one member during your hearing, but do you fund any research on the issue of potential social media addiction, and if not, would you consider funding independent third-party research in this area?**

Facebook employs social psychologists, social scientists, and sociologists, and collaborates with top scholars to better understand well-being. Facebook has also pledged \$1 million towards research to better understand the relationship between media technologies, youth development and well-being. Facebook is teaming up with experts in the field to look at the impact of mobile technology and social media on kids and teens, as well as how to better support them as they transition through different stages of life. Facebook is committed to bringing people together and supporting well-being through meaningful interactions on Facebook.

Senate Committee on Commerce, Science, and Transportation

Hearing Follow-up Questions

Senator Capito

Please send someone to the opioid meeting.

Thank you for highlighting this important issue. Yes, we will work with the Administration to send a Facebook representative. We are committed to doing our part in combating the opioid crisis and look forward to a continued dialogue with you.

Senator Moran

How can a bug bounty deal with reporting the sharing of data?

The Data Abuse Bounty Program, inspired by the existing Bug Bounty Program, helps us identify violations of our policies by requesting narrative descriptions of violations from individuals with direct and personal knowledge of events. The Data Abuse Bounty will reward people with first-hand knowledge and proof of cases where a Facebook platform app collects and transfers people's data to another party to be sold, stolen, or used for scams or political influence. We'll review all legitimate reports and respond as quickly as possible when we identify a credible threat to people's information. If we confirm data abuse, we will shut down the offending app and, if necessary, take legal action against the company selling or buying the data. We'll pay a bounty to the person who reported the issue or allow them to donate their bounty to a charity, and we'll also alert those we believe to be affected. We also encourage our users to report to us content that they find concerning or that results in a bad experience, as well as other content that may violate our policies. We review these reports and take action on abuse, like removing content and disabling accounts.

Senator Baldwin

Do you know whether Aleksandr Kogan sold any of the data he collected to anyone other than Cambridge Analytica?

Kogan represented to us that he provided data to SCL, Eunoia Technologies (a company founded by Christopher Wylie), and a researcher at the Toronto Laboratory for Social Neuroscience at the University of Toronto. He represented to Facebook that he only received payment from SCL/Cambridge Analytica.

How much do you know or have you tried to find out how Cambridge Analytica used the data while they had it before you believed they deleted it?

On December 11, 2015, *The Guardian* published an article reporting that Kogan and his company, GSR, may have passed information his app had obtained from Facebook users to SCL Elections Ltd. (SCL)/Cambridge Analytica. By doing so, Kogan and his company violated Facebook's Platform Policies, which explicitly prohibited selling or licensing 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 his app from our platform and launched an investigation into these allegations. Kogan signed a certification declaring that he had deleted all data that he obtained through his app and obtained certifications of deletion from others he had shared data with, including Cambridge Analytica. In March 2018, new allegations surfaced that Cambridge Analytica may not have deleted data as it had represented. Our investigation of these matters is ongoing.

I find some encouragement in the steps you have outlined today to provide greater transparency regarding political ads. I want to get further information on how you can be confident that you have excluded entities based outside of the United States.

Pursuant to the new transparency measures Facebook is launching, all advertisers who want to run ads with political content targeted at the US will have to confirm their identity and location by providing either a US driver's license or passport, last four digits of their social security number, and a residential mailing address. In addition, 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.

Senator Cruz

The predicate for Section 230 immunity under the CDA is that you're a neutral public forum. Do you consider yourself a neutral public forum or are you engaged in political speech, which is your right under the First Amendment?

We are, first and foremost, a technology company. Facebook does not create or edit the content that our users published on our platform. While we seek to be a platform for a broad range of ideas, we do moderate content in good faith according to published community standards in order to keep users on the platform safe, reduce objectionable content and to make sure users participate on the platform responsibly.

Section 230 of the Communications Decency Act provides that “[N]o provider or user of an interactive computer service shall be treated as the publisher or speaker of any information provided by another information content provider.” Outside of certain specific exceptions, this means that online platforms that host content posted by others are generally not liable for the speech of their users, and, indeed, Section 230 explicitly provides that a platform that chooses to moderate content on its service based on its own standards does not incur liability on the basis of that decision. Specifically, 47 U.S.C. 230(c)(2) provides, in relevant part, that “[N]o provider or user of an interactive computer service shall be held liable on account of—(A) any action voluntarily taken in good faith to restrict access to or availability of material that the provider or user considers to be obscene, lewd, lascivious, filthy, excessively violent, harassing, or otherwise objectionable, whether or not such material is constitutionally protected.”

Senator Cantwell

During the 2016 campaign, Cambridge Analytica worked with the Trump campaign to refine tactics. Were Facebook employees involved in that?

During the 2016 election cycle, Facebook worked with campaigns to optimize their use of the platform, including helping them understand various ad formats and providing other best practices guidance on use of the platform.

Senator Gardner

To clarify one of the comments made about deleting accounts from Facebook, in your user agreement it says when you delete IP content, if deleted in a manner similar to emptying recycle bin on a computer, you understand it may persist in backup copies for a reasonable period of time. How long is that?

In general, when a user deletes their account, we delete things they have posted, such as their photos and status updates, and they won't be able to recover that information later. (Information that others have shared about them isn't part of their account and won't be deleted.)

There are some limited exceptions to these policies: For instance, information can be accessed and preserved for an extended period when it is the subject of a legal request or obligation, governmental investigation, or investigations of possible violations of our terms or policies, or otherwise to prevent harm. We also retain information from accounts disabled for terms violations for at least a year to prevent repeat abuse or other term violations.

Can you get a warrant to join a Page? Pretending you're a separate user to track it? Can the Government just do that? The FBI? Anybody?

As part of official investigations, government officials sometimes request data about people who use Facebook. We have strict processes in place to handle these government requests, and we disclose account records in accordance with our terms of service and applicable law. We also have law enforcement response teams available around the clock to respond to emergency requests. As part of our ongoing effort to share information about the requests we have received from governments around the world, Facebook regularly produces a Transparency Report about government requests to Facebook. Additionally, Facebook's policy (available at https://www.facebook.com/help/112146705538576?helpref=faq_contentand) expressly prohibits impersonation. This policy applies to everyone, including law enforcement, and Facebook will disable any account that we believe is inauthentic.

Senator Wicker

Does Facebook allow minors (13-17) to opt in to share their call and text history?

Call and text history logging is part of an opt-in feature that lets people import contact information to help them connect with people they know on Facebook and Messenger. We introduced the call and text history component of this feature for Android users several years ago, and currently offer it in Messenger and Facebook Lite, a lightweight version of Facebook, on Android.

We've reviewed this feature to confirm that Facebook does not collect the content of messages—and will delete all logs older than one year. In the future, the client will only upload to our servers the information needed to offer this feature—not broader data such as the time of calls. We do allow people from 13 to 17 to opt into this service. However, we do take other steps to protect teens on Facebook and Messenger:

- We provide education before allowing teens to post publicly.
- We don't show search results based on specific profile data (high school, birthday/age, and hometown, or current city) of teens to unconnected adults when the adults search on Facebook.
- Unconnected adults can't message minors who are 13-17.
- We have age limits for advertisements. For example, ads for dating sites, financial services and other products or services are gated to users under 18.
- We've also helped many teenagers with information about bullying prevention campaigns and online safety tips.

There have been reports that Facebook can track users' internet browsing activity even after that user has logged off of the Facebook platform. Can you confirm whether or not this is true? Would you also let us know how Facebook discloses to its users that engaging in this type of tracking gives us that result of tracking between devices?

When people visit apps or websites that feature our technologies—like the Facebook Like or Comment button—our servers automatically log (i) standard browser or app records of the fact that a particular device or user visited the website or app (this connection to Facebook's servers occurs automatically when a person visits a website or app that contains our technologies, such as a Like button, and is an inherent function of Internet design); and (ii) any additional information the publisher of the app or website chooses to share with Facebook about the person's activities on that site (such as the fact that a purchase was made on the site). This is a standard feature of the Internet, and most websites and apps share this same information with multiple different third-parties whenever people visit their website or app. For example, the Senate Commerce Committee's website shares information with Google and its affiliate DoubleClick and with the analytics company Webtrends. This means that, when a person visits the Committee's website, it sends browser information about their visit to each one of those third

parties. More information about how this works is available at <https://newsroom.fb.com/news/2018/04/data-off-facebook/>.

When the person visiting a website featuring Facebook's tools is not a registered Facebook user, Facebook does not have information identifying that individual, and it does not create profiles for this individual.

We use the browser and app logs that apps and websites send to us—described above—in the following ways for non-Facebook users. First, these logs are critical to protecting the security of Facebook and to detecting or preventing fake account access. For example, if a browser has visited hundreds of sites in the last five minutes, that's a sign the device might be a bot, which would be an important signal of a potentially inauthentic account if that browser then attempted to register for an account. Second, we aggregate those logs to provide summaries and insights to websites and apps about how many people visit or use their product or use specific features like our Like button—but without providing any information about a specific person. We do not create profiles for non-Facebook users, nor do we use browser and app logs for non-Facebook users to show targeted ads from our advertisers to them or otherwise seek to personalize the content they see. However, we may take the opportunity to show a general ad that is unrelated to the attributes of the person or an ad encouraging the non-user to sign up for Facebook.

When the individual is a Facebook user, we are also able to use this information to personalize their experiences on Facebook, whether or not they are logged out, but we will not target ads to users relying on this information unless the user allows this in their privacy settings. We do not sell or share this information with third-parties.

Senator Blunt

Do you track non-Facebook data from devices on which they have used Facebook, even if they are logged off of Facebook or the device is offline? So you don't have bundled permissions for how I can agree to what devices I may use that you may have contact with? Do you bundle that permission, or am I able to individually say what I'm willing for you to watch and what I don't want you to watch?

When people visit apps or websites that feature our technologies—like the Facebook Like or Comment button—our servers automatically log (i) standard browser or app records of the fact that a particular device or user visited the website or app (this connection to Facebook's servers occurs automatically when a person visits a website or app that contains our technologies, such as a Like button, and is an inherent function of Internet design); and (ii) any additional information the publisher of the app or website chooses to share with Facebook about the person's activities on that site (such as the fact that a purchase was made on the site). This is a standard feature of the Internet, and most websites and apps share this same information with multiple different third-parties whenever people visit their website or app. For example, the Senate Commerce Committee's website shares information with Google and its affiliate DoubleClick and with the analytics company Webtrends. This means that, when a person visits the Committee's website, it sends browser information about their visit to each one of those third parties. More information about how this works is available at <https://newsroom.fb.com/news/2018/04/data-off-facebook/>.

When the person visiting a website featuring Facebook's tools is not a registered Facebook user, Facebook does not have information identifying that individual, and it does not create profiles for this individual.

We use the browser and app logs that apps and websites send to us—described above—in the following ways for non-Facebook users. First, these logs are critical to protecting the security of Facebook and to detecting or preventing fake account access. For example, if a browser has visited hundreds of sites in the last five minutes, that's a sign the device might be a bot, which would be an important signal of a potentially inauthentic account if that browser then attempted to register for an account. Second, we aggregate those logs to provide summaries and insights to websites and apps about how many people visit or use their product or use specific features like our Like button—but without providing any information about a specific person. We do not create profiles for non-Facebook users, nor do we use browser and app logs for non-Facebook users to show targeted ads from our advertisers to them or otherwise seek to personalize the content they see. However, we may take the opportunity to show a general ad that is unrelated to the attributes of the person or an ad encouraging the non-user to sign up for Facebook.

When the individual is a Facebook user, we are also able to use this information to personalize their experiences on Facebook, whether or not they are logged out, but we will not target ads to users relying on this information unless the user allows this in their privacy settings. We do not sell or share this information with third-parties.

Do you track devices that an individual who uses Facebook has that is connected to the device that they use for their Facebook connection but not necessarily connected to Facebook?

Yes, Facebook's Data Policy specifically discloses that we associate information across different devices that people use to provide a consistent experience wherever they use Facebook.

Facebook's services inherently operate on a cross-device basis: understanding when people use our services across multiple devices helps us provide the same personalized experience wherever people use Facebook—for example, to ensure that a person's News Feed or profile contains the same content whether they access our services on their mobile phone or in a desktop computer's web browser.

In support of those and other purposes, we collect information from and about the computers, phones, connected TVs and other web-connected devices our users use that integrate with our Products, and we combine this information across a user's different devices. For example, we use information collected about a person's use of our Products on their phone to better personalize the content (including ads) or features they see when they use our Products on another device, such as their laptop or tablet, or to measure whether they took an action in response to an ad we showed them on their phone or on a different device.

Information we obtain from these devices includes:

- **Device attributes:** information such as the operating system, hardware and software versions, battery level, signal strength, available storage space, browser type, app and file names and types, and plugins.
- **Device operations:** information about operations and behaviors performed on the device, such as whether a window is foregrounded or backgrounded, or mouse movements (which can help distinguish humans from bots).
- **Identifiers:** unique identifiers, device IDs, and other identifiers, such as from games, apps or accounts people use, and Family Device IDs (or other identifiers unique to Facebook Company Products associated with the same device or account).
- **Device signals:** Bluetooth signals, and information about nearby Wi-Fi access points, beacons, and cell towers.
- **Data from device settings:** information a user allows us to receive through device settings they turn on, such as access to their GPS location, camera or photos.
- **Network and connections:** information such as the name of a user's mobile operator or ISP, language, time zone, mobile phone number, IP address, connection speed and, in some cases, information about other devices that are nearby or on their network, so we can do things like help them stream a video from their phone to their TV.
- **Cookie data:** data from cookies stored on a user's device, including cookie IDs and settings. More information is available at

<https://www.facebook.com/policies/cookies/> and <https://help.instagram.com/1896641480634370?ref=ig>.

Advertisers, app developers, and publishers can send us information through Facebook Business Tools they use, including our social plug-ins (such as the Like button), Facebook Login, our APIs and SDKs, or the Facebook pixel. These partners provide information about a person's activities off Facebook—including information about their device, websites they visit, purchases they make, the ads they see, and how they use their services—whether or not they have a Facebook account or are logged into Facebook. For example, a game developer could use our API to tell us what games a person plays, or a business could tell us about a purchase a person made in its store. We also receive information about a person's online and offline actions and purchases from third-party data providers who have the rights to provide us with that person's information.

We use the information we have to deliver our Products, including to personalize features and content (including a person's News Feed, Instagram Feed, Instagram Stories and ads) and make suggestions for a user (such as groups or events they may be interested in or topics they may want to follow) on and off our Products. To create personalized Products that are unique and relevant to them, we use their connections, preferences, interests and activities based on the data we collect and learn from them and others (including any data with special protections they choose to provide); how they use and interact with our Products; and the people, places, or things they're connected to and interested in on and off our Products.

For example, if people have shared their device locations with Facebook or checked into a specific restaurant, we can show them ads from an advertiser that wants to promote its services in their area or from the restaurant. We use location-related information—such as a person's current location, where they live, the places they like to go, and the businesses and people they're near—to provide, personalize and improve our Products, including ads, for them and others. Location-related information can be based on things like precise device location (if a user has allowed us to collect it), IP addresses, and information from their and others' use of Facebook Products (such as check-ins or events they attend). We store data until it is no longer necessary to provide our services and Facebook Products, or until a person's account is deleted—whichever comes first. This is a case-by-case determination that depends on things like the nature of the data, why it is collected and processed, and relevant legal or operational retention needs. 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 someone (information such as a person's name or email address that by itself can be used to contact them or identifies who they are) unless they give us permission. For example, we provide general demographic and interest information to advertisers (for example, that an ad was seen by a woman between the ages of 25 and 34 who lives in Madrid and likes software engineering) to help them better understand their audience. We also confirm which Facebook ads led users to make a purchase or take an action with an advertiser.

Senator Fischer

There have been some past reports that indicate that Facebook collects about 98 data categories. For those two billion active users. That's 192 billion data points that are being generated. I think at any time. From consumers globally. Do you store any?

Your question likely references a *Washington Post* article that purported to identify “98 data points that Facebook uses to target ads to you.” The article was based on the writer’s use of the tool that allows advertisers to select the audience that they want to see their ads. Anyone on Facebook can see the tool and browse the different audiences that advertisers can select.

The “data points” to which the article refers are not categories of information that we collect from everyone on Facebook. Rather, they reflect audiences into which at least some people on Facebook fall, based on the information they have provided and their activity. For example, the article lists “field of study” and “employer” as two of the “data points” that can be used to show ads to people. People can choose to provide information about their field of study and their employer in profile fields, and those who do may be eligible to see ads based on that information—unless they have used the controls in Ad Preferences that enable people to opt out of seeing ads based on that information. The same is true of the other items in the list of 98.

Further, 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, 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 to edit or delete these 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>.

Please note, however, that (as the article explains) many of these refer to “Partner Categories”—audiences that are offered by third-party data providers. We announced in April that we would stop offering this kind of targeting later this year.

Please also see our letter to you dated April 27, 2018.

Senator Heller

Can you tell me how many Nevadans were among the 87 million that received this notification?

A state-by-state breakdown is available at
<https://newsroom.fb.com/news/2018/04/restricting-data-access/>.

How long do you keep a user's data? How long do you keep a user's data once they have left? If they choose to delete their account, how long do you keep their data?

In general, when a user deletes their account, we delete things they have posted, such as their photos and status updates, and they won't be able to recover that information later. (Information that others have shared about them isn't part of their account and won't be deleted.)

There are some limited exceptions to these policies: For instance, information can be accessed and preserved for an extended period when it is the subject of a legal request or obligation, governmental investigation, or investigations of possible violations of our terms or policies, or otherwise to prevent harm. We also retain information from accounts disabled for terms violations for at least a year to prevent repeat abuse or other term violations.

Senator Hassan

The other question I had, and it does not just apply to Facebook, is should the framework include financial penalties when large providers like Facebook are breached and privacy is compromised as a result? There is very little incentive for whether it is Facebook or Equifax to actually be abreast of protecting customer privacy and working for potential breaches or vulnerabilities in the system.

Protecting people's data is one of our most important responsibilities. We know that if people don't trust that their information will be safe on Facebook, they won't feel comfortable using our services.

We have every incentive to work as hard as we can to protect people's information, and we're committed to continuing our work to improve those protections.

Facebook is generally open to the idea of federal breach notification requirements, particularly legislation that would centralize reporting and ensure a consistent approach across the United States. For example, in Europe, the GDPR requires notification to a lead supervisory authority, rather than individual member states, in cases of a data breach. In the United States, however, there is no centralized notification scheme, and instead, reporting obligations vary widely across all 50 states. This complexity makes it harder to respond appropriately and swiftly to protect people in the event of a data breach. We believe this is an important issue and an area that is ripe for thoughtful regulation.

Facebook is generally not opposed to regulation but wants to ensure it is the right regulation. We are already regulated in many ways—for example, under the Federal Trade Commission Act—and we are subject to ongoing oversight by the FTC under the terms of a 2011 consent order. Facebook has inherent incentives to protect its customers' privacy and address breaches and vulnerabilities. Indeed, the recent discovery of misconduct by an app developer on the Facebook platform clearly hurt Facebook and made it harder for us to achieve our social mission. As such, Facebook is committed to protecting our platform from bad actors, ensuring we are able to continue our mission of giving people a voice and bringing them closer together. We are also actively building new technologies to help prevent abuse on its platform, including advanced AI tools to monitor and remove fake accounts. We have also significantly increased our investment in security, employing more than 15,000 individuals working solely on security and content review and planning to increase that number to over 20,000 by the end of the year. We have also strengthened our advertising policies, seeking to prevent discrimination while improving transparency.

Senator Klobuchar

Do you support a rule that would require you to notify your users of a breach within 72 hours?

Facebook is generally open to the idea of breach notification requirements, particularly legislation that would centralize reporting and ensure a consistent approach across the United States. For example, in Europe, the GDPR requires notification to a lead supervisory authority, rather than individual member states, in cases of a data breach. In the United States, however, there is no centralized notification scheme, and instead, reporting obligations vary widely across all 50 states. This complexity makes it harder to respond appropriately and swiftly to protect people in the event of a data breach. We believe this is an important issue and an area that is ripe for thoughtful regulation.

Senator Markey**Do you support a kids' privacy bill of rights where opt-in is the standard?**

Facebook is generally not opposed to regulation but wants to ensure it is the right regulation. The issues facing the industry are complex, multi-faceted, and affect an important part of peoples' lives. As such, Facebook is absolutely committed to working with regulators, like Congress, to craft the right regulations. Facebook would be happy to review any proposed legislation and provide comments.

Senator Young

Might we create stronger privacy rights for consumers through creating a stronger general property right regime online, say a law states that users own their online data or stronger opt in requirements on platforms like yours? If we're to do that, would you need to retool your model? If we're to adopt one of the two approaches?

Our Terms of Service confirm that people own the information they shared on Facebook. They entrust it to us to use it consistent with our Terms and Data Policy to provide meaningful and useful services to them. They have the ability to choose who can see it, delete it, or take it with them if they want to do so. We're also rolling out a new Privacy Shortcuts feature, which centralizes a broad range of choices that people have about how their information is used as a part of the Facebook service, and we're contacting people on our service to ask them to make choices about these issues as well.

Facebook already allows users to download a copy of their information from Facebook. This functionality, which we've offered for many years, includes numerous categories of data, including About Me, Account Status History, Apps, Chat, Follower, Following, Friends, Messages, Networks, Notes, and more. We recently launched improvements to our "Download Your Information" tool, including to give people choices about whether they want to download only certain types of information and about the format in which they want to receive the download, to make it easier for people to use their information once they've retrieved it.

Of course, the details of any new privacy legislation matter, and we would be pleased to discuss any specific proposals with you and your staff.

Senator Peters

Well, you bring up the principles because, as you are well aware, AI systems, especially in very complex environments when you have machine learning, it is sometimes very difficult to understand, as you mentioned, exactly how those decisions were arrived at. There are examples of how decisions are made on a discriminatory basis and that they can compound if you are not very careful about how that occurs. And so is your company—you mentioned principles. Is your company developing a set of principles that are going to guide that development? And would you provide details to us as to what those principles are and how they will help deal with this issue?

We are focused on both the technical and the ethical aspects of artificial intelligence. We believe these two should go hand-in-hand together in order to fulfill our commitment to being fair, transparent, and accountable in our development and use of AI. Facebook has AI teams working on developing the philosophical, as well as technical, foundations for this work. Facebook is also one of the co-founders and members of the Partnership on AI (PAI), a collaborative and multi-stakeholder organization established to study and formulate best practices on AI technologies, to advance the public's understanding of AI, and to serve as an open platform for discussion and engagement about AI and its influences on people and society. The thematic pillars that structure the work we're doing in the scope of the PAI—safety, fairness, transparency, and accountability—are the principles that we believe industry should follow and promote when building and deploying AI systems. The PAI's Fair, Transparent and Accountable AI Working Group is also working alongside industry, academia, and civil society to develop best practices around the development and fielding of fair, explainable, and accountable AI systems.