

**OPEN HEARING ON FOREIGN INFLUENCE
OPERATIONS' USE OF SOCIAL MEDIA PLATFORMS
(COMPANY WITNESSES)**

HEARING
BEFORE THE
SELECT COMMITTEE ON INTELLIGENCE
OF THE
UNITED STATES SENATE
ONE HUNDRED FIFTEENTH CONGRESS
SECOND SESSION

WEDNESDAY, SEPTEMBER 5, 2018

Printed for the use of the Select Committee on Intelligence



Available via the World Wide Web: <http://www.govinfo.gov>

U.S. GOVERNMENT PUBLISHING OFFICE

31-350 PDF

WASHINGTON : 2019

SELECT COMMITTEE ON INTELLIGENCE

[Established by S. Res. 400, 94th Cong., 2d Sess.]

RICHARD BARR, North Carolina, *Chairman*

MARK R. WARNER, Virginia, *Vice Chairman*

JAMES E. RISCH, Idaho

MARCO RUBIO, Florida

SUSAN COLLINS, Maine

ROY BLUNT, Missouri

JAMES LANKFORD, Oklahoma

TOM COTTON, Arkansas

JOHN CORNYN, Texas

DIANNE FEINSTEIN, California

RON WYDEN, Oregon

MARTIN HEINRICH, New Mexico

ANGUS KING, Maine

JOE MANCHIN III, West Virginia

KAMALA HARRIS, California

MITCH McCONNELL, Kentucky, *Ex Officio*

CHUCK SCHUMER, New York, *Ex Officio*

JAMES INHOFE, Oklahoma, *Ex Officio*

JACK REED, Rhode Island, *Ex Officio*

CHRIS JOYNER, *Staff Director*

MICHAEL CASEY, *Minority Staff Director*

KELSEY STROUD BAILEY, *Chief Clerk*

CONTENTS

SEPTEMBER 5, 2018

OPENING STATEMENTS

Burr, Hon. Richard, Chairman, a U.S. Senator from North Carolina	1
Warner, Mark R., Vice Chairman, a U.S. Senator from Virginia	3

WITNESSES

Sandberg, Sheryl, Chief Operating Officer, Facebook	6
Prepared statement	9
Dorsey, Jack, Chief Executive Officer, Twitter, Inc.	19
Prepared statement	21

SUPPLEMENTAL MATERIAL

Responses to Questions for the Record by:	
Sheryl Sandberg	68
Jack Dorsey	133

**OPEN HEARING ON FOREIGN INFLUENCE
OPERATIONS' USE OF SOCIAL MEDIA
PLATFORMS (COMPANY WITNESSES)**

WEDNESDAY, SEPTEMBER 5, 2018

U.S. SENATE,
SELECT COMMITTEE ON INTELLIGENCE,
Washington, DC.

The Committee met, pursuant to notice, at 9:32 a.m., in Room G-50, Dirksen Senate Office Building, Hon. Richard Burr (Chairman of the Committee) presiding.

Present: Senators Burr, Warner, Risch, Rubio, Collins, Blunt, Lankford, Cotton, Wyden, Heinrich, King, Manchin, Harris, and Reed.

**OPENING STATEMENT OF HON. RICHARD BURR, CHAIRMAN, A
U.S. SENATOR FROM NORTH CAROLINA**

Chairman BURR. I'd like to call the hearing to order. And I'd like to welcome our witnesses today: Jack Dorsey, chief executive officer at Twitter—Jack, welcome—and Sheryl Sandberg, chief operating officer at Facebook. I thank both of you for being here with us this morning.

Before I make my remarks, I want to say a few words about our colleague, our friend, and committee ex officio member Senator John McCain.

John could be blunt, and he could be direct, but when it came to committing himself to a cause that he believed in, John McCain was without equal. This Senate, this deliberative body, with its history and its traditions, will survive the passing of John McCain, but there can be no denying that the place is a little smaller without him. We will continue to do the important work we do here with passion, resolve, and a sense of purpose born from moral conviction. John would want that. In fact, he would insist on it from each of us.

My friends, if I can borrow the phrase: Arizona's loss is our loss, and our loss is America's loss. John McCain will be dearly missed, and as you can see, we have set his spot on the dais today.

Jack, Sheryl—as a committee, we've learned more about social media over the last 18 months than I suspect most of us ever thought we would in a lifetime. We've learned about social media's boundless potential for good and its ability to enable thoughtful and engaged interactions on a global scale.

But we've also learned about how vulnerable social media is to corruption and misuse. The very worst examples of this are abso-

lutely chilling and a threat to our democracy: the founding ideal of different people from different beliefs and ideas all living peacefully under a single flag. The committee takes this issue very seriously and we appreciate the fact that Facebook and Twitter are represented here this morning with an equivalent and appropriate measure of seriousness.

The purpose of today's hearing is to discuss the role that social media plays in the execution of foreign influence operations. In the past, we've used terms like misinformation and divisive content to describe this activity.

Now as we go into our fourth and final hearing on this subject, I think it's important that we be precise and candid with our language, because that's what the significance of this threat demands. We need to be precise about the foreign actors we're talking about, we need to be precise about the consequences of not acting, and we need to be candid about where responsibility for solving this problem lies.

Two weeks ago your companies announced a series of successful disruptions that resulted in the removal of 652 Facebook pages, groups, and accounts, and 284 Twitter accounts based on their violating your company's standards of coordinated manipulation and inauthentic behavior. Google's own internal security teams did commendable work disrupting this influence operation and we would have valued the opportunity to speak with them at the appropriate level of corporate representation. Nevertheless, their efforts should be acknowledged.

In a departure from what we've all gotten a little accustomed to, this activity didn't come from Russia. It came from Iran. My instinct is to applaud the diligence of your security teams and credit you with taking the problem very seriously.

But I'm not sure your success is the big story here. As I understand it, a third-party security team was crucial to identifying the scope of the Iranian activity. And even more concerning is that more foreign countries are now trying to use your products to shape and manipulate American political sentiment as an instrument of statecraft.

Jack, I was pleased when informed about your efforts to improve conversational health at Twitter. I think that kind of initiative can do a lot to improve the transparency of public discourse on your platform, and foreign influence operations thrive without transparency.

Sheryl, I fully support Facebook's hiring of the right security experts, building the necessary technologies and collaborating across law enforcement, commercial, cybersecurity, and social media company lines.

I think the observation that no one company can fight this on their own is spot on. Unfortunately, what I described as a national security vulnerability and an unacceptable risk back in November remains unaddressed. That risk and vulnerability was highlighted yet two weeks ago. Without question, positive things are happening. The collaboration, dedication, and resources and demonstrated willingness to work with us are critical and valued by every member of this committee.

It takes courage to call out a state actor and your companies have done that. But clearly this problem is not going away. I'm not even sure it's trending in the right direction. I will go back to what I said up front: we need to be candid about responsibility, and by that, I mean both the responsibility we have to one another—from one side of this dais to the other—as participants in this public policy discussion. And more importantly our shared responsibility to the American people.

Technology always moves faster than regulation, and to be frank, the products and services that enable social media don't fit neatly into the consumer safety or regulatory constructs of the past. The old definitions that used to differentiate a content publisher from a content facilitator are just not helpful here. I think that ambiguity has given rise to something of a convenient identity crisis, whereby judgments about what is and isn't allowable on social media are too episodic, too reactive, and too unrestricted. People are affected by the information your platforms channel to them. That channeling isn't passive or random. It's a function of brilliant algorithms and an incentive structure that prizes engagement. None of that is under attack here.

What is under attack is the idea that business as usual is good enough. The information your platform disseminates changes minds and hardens opinions. It helps people make sense of the world. When you control that or you influence a little of it, you're in a position to win wars without firing a shot. That's how serious this is.

We've identified the problem. Now it's time to identify the solution. Sheryl and Jack, I'm glad you decided to appear and your willingness to be part of the solution. I'm disappointed Google decided against sending out the right senior-level executive to participate in what I truly expect to be a productive discussion.

If the answer is regulation, let's have an honest dialogue about what that looks like. If the key is more resources or legislation that facilitates information sharing and government cooperation, let's get it out there. If it's national security policies that punish the kind of information and influence operations we're talking about this morning, to the point that they aren't even considered in foreign capitals, then let's acknowledge that. But whatever the answer is, we've got to do this collaboratively and we've got to do it now. That's our responsibility to the American people.

I'll offer a closing point. This is for the witnesses and the members alike. There are no unsolvable problems. There is only the will to do what needs to be done—or its absence.

With that, I turn to the Vice Chairman for any comments.

**OPENING STATEMENT OF HON. MARK R. WARNER, VICE
CHAIRMAN, A U.S. SENATOR FROM VIRGINIA**

Vice Chairman WARNER. Thank you, Mr. Chairman. And let me first of all echo your comments about our colleague and friend, John McCain. I hope we all take his advice to continue to put country first.

Welcome to the witnesses. Mr. Chairman has pointed out that today is an important public discussion. I am pleased that both Facebook and Twitter have sent their company's top leadership to

address some of the critical public policy challenges. I look forward to a constructive engagement.

I'd say, though, that I am deeply disappointed that Google, one of the most influential digital platforms in the world, chose not to send its own top corporate leadership to engage this committee. Because I know our members have a series of difficult questions about structural vulnerabilities on a number of Google's platforms that we will need answers for: from Google Search, which continues to have problems surfacing absurd conspiracies; to YouTube, where Russian-backed disinformation agents promoted hundreds of divisive videos; to Gmail, where state-sponsored operatives attempted countless hacking attempts. Google has an immense responsibility in this space.

Given its size and influence, I would have thought that leadership at Google would have wanted to demonstrate how seriously it takes these challenges and actually take a leadership role in this important discussion. Unfortunately, they didn't choose to make that decision. But for the two companies that have chosen to constructively engage and to publicly answer some difficult and challenging questions, again, thank you.

Now, it would be an understatement to say that much has changed in the aftermath of the 2016 campaign. With the benefit of hindsight, it's obvious that serious mistakes were made by both Facebook and Twitter. You, like the Federal Government, were caught flat-footed by the brazen attacks on our election.

Even after the election, you were reluctant to admit there was a problem. I think in many ways it was pressure that was brought to bear by this committee that led Facebook, Twitter, and yes, Google to uncover the malicious activities of the Russian-backed internet Research Agency activities on each of your platforms.

Now each of you have come a long way with respect to recognizing the threat. We've seen important action by your companies to make political advertising more transparent—and we discussed this yesterday—by complying with the terms Senator Klobuchar and I put forward in the Honest Ads Act. In addition, as the Chairman mentioned, since last September you have identified and removed some bad actors from your platforms.

The bad news, I'm afraid, is that there's still a lot of work to do, and I'm skeptical that ultimately you'll be able to truly address this challenge on your own. I believe Congress is going to have to act.

First, on the disinformation front: Russia has not stopped. Russian-linked information warfare exists today. Just recently, we saw the two of you take action to take down suspected Russian operations. We also know Microsoft uncovered Russian attempts to hack political organizations and potentially several political campaigns.

The Russians also continue to infiltrate and manipulate American social media to hijack our national conversation. Again, you've gotten better, and I'm pleased to see that you've begun to take action, but also the Russians are getting better as well. They have now become harder to track. Worse, now that the Russian playbook is out there, other adversaries, as we saw recently, like Iran, have joined the fray.

But foreign-based disinformation campaigns represent just a fraction of the challenge before you. In the same way that bots, trolls, fake pages, algorithmic gaming can be used to spread fake news, these same tools can be used to assist financial stock pumping fraud, to create filter bubbles and alternative realities, to incite ethnic and racial violence, and countless other misuses.

Imagine the challenge and damage to the markets if Ford's communications from the Fed Chairman were leaked online. Or consider the price of a Fortune 500 company's stock if a dishonest short seller was able to spread false information about the company's CEO or the effects of its products rapidly online.

Russian disinformation has revealed a dark underbelly of the entire online ecosystem, and this threatens to cheapen American discourse, weaken privacy, erode truth, and undermine our democracy on a previously unimagined scale. Worse, this is only going to get harder as we move into artificial intelligence, use of Deepfake technology.

During the 2016 election campaign, the Russians demonstrated how bad actors can effectively marry offensive cyber operations, including hacking, with information operations. I'm afraid that we're on the cusp of a new generation of exploitation, potentially harnessing hacked personal information, to enable tailored and targeted disinformation in social engineering efforts. That future should concern us all.

As someone who was involved in the tech industry for more than 20 years, I respect what this industry represents, and I don't envy the significant technical and policy challenges you face. But the size and reach of your platforms demand that we as policy makers do our job to ensure proper oversight, transparency, and protection for American users and our democratic institutions.

The era of the Wild West in social media is coming to an end. Where we go from here, though, is an open question. These are complicated technological challenges, and Congress has at times demonstrated that it still has some homework to do. I do think this committee has done more to understand the threat to our democracy posed by social media than any others, and I want to commend my colleagues on this committee for tackling this challenge in a bipartisan way.

As has been mentioned, this is our fourth public hearing on the subject, and we've met behind closed doors countless times with third-party researchers, with government officials, and with each of the platforms. We've done the work, and we're positioned to continue to lead in this space.

Again, as the Chairman has already indicated, today's hearing is not about gotcha questions or scoring political points. Our goal today is to begin to shape actual policy solutions which will help us tackle this challenge.

Now, I've put forth some ideas that I'd like to get your constructive thoughts on. For instance, don't your users have a right to know when they're interacting with bots on your platform? Isn't there a public interest in insuring more anonymized data is available to help researchers and academics identify the potential problems and misuse? Why are your terms of service so difficult to find and nearly impossible to read, much less understand? Why

shouldn't we adopt ideas like data portability, data minimization, or first-party consent? And after witnessing numerous episodes of misuse, what further accountability should there be with respect to the flawed advertising model that you utilize?

Now these are just some of our ideas. We have received a lot of positive feedback on some of these ideas from both experts and users. We've also been accused of trying to bring about the death of the internet. I'm anxious to hear your views on our proposals and suggestions your teams can bring to the table on this front.

We have to be able to find smart, thoughtful policy solutions that get us somewhere beyond the status quo, without applying ham-handed 20th-century solutions to 21st-century problems. At the same time, we should be mindful to adopt policies that do not simply entrench the existing dominant platforms.

These are not just challenges for our politics or our democracy. These threats can affect our economy, our financial system, and other parts of our lives. I'm hopeful that we can get there. I'm confident in American ingenuity. And I'm optimistic that Congress led by this committee in a bipartisan fashion can move this conversation forward.

I look forward to the discussion and appreciate the hearing being called. Thank you, Mr. Chairman.

Chairman BURR. I thank the Vice Chairman. At this time, I'd like to swear in our witnesses. If I could ask both of you to stand and raise your right hand?

Do you solemnly swear to give this committee the truth, the full truth and nothing but the truth so help you God?

[The witnesses answered in the affirmative.]

Please be seated. Ms. Sandberg, I'd like to recognize you first and then Mr. Dorsey for any opening statement you'd like to make. The floor is yours.

**STATEMENT OF SHERYL SANDBERG, CHIEF OPERATING
OFFICER, FACEBOOK**

Ms. SANDBERG. Thank you. Chairman Burr, Vice Chairman Warner, and members of this select committee, thank you for giving me the opportunity to speak with you today. My written testimony goes into more detail about the actions we're taking to prevent election interference on Facebook. But I wanted to start by explaining how seriously we take these issues and talk about some of the steps we're taking.

Free and fair elections are the foundation of any democracy. As Americans, they are part of our national identity and that's why it's incumbent upon all of us to do all we can to protect our democratic process. That includes Facebook. At its best, Facebook plays a positive role in our democracy, enabling representatives to connect with their constituents, reminding people to register and to vote, and giving people a place to freely express their opinions about the issues that matter to them.

However, we've also seen what can happen when our service is abused. As a bipartisan report from this committee said, Russia used social media as part of, and I quote: a comprehensive and multi-faceted campaign to sow discord, undermine democratic institutions and interfere in U.S. elections and those of our allies.

We were too slow to spot this and too slow to act. That is on us. This interference was completely unacceptable. It violated the values of our company and of the country we love. Actions taken show how determined we are to do everything we can do to stop this from happening.

The threat we face is not new. America has always confronted attacks from determined, well-funded opponents who want to undermine our democracy. What is new is the tactics they are using. To stay ahead, we all need to work together, as Chairman Burr said: government, law enforcement, industry and experts from civil society. And that is why I'm grateful for the work this committee is doing.

At Facebook, we're investing in security for the long term. As our defenses improve, bad actors learn and improve too, and that's why security is never a finished job. We have more than doubled the number of people we have working in safety and security and we now have over 20,000 people and we are able to view reports in 50 languages, 24 hours a day.

Better machine learning and artificial intelligence have enabled us to be more proactive in finding abuse. In the first three months of 2018 alone, over 85 percent of the violent content we took down or added warning labels to was identified by our technology before it was reported. These are expensive investments, but that will not stop us because we know they are critical.

Our first line of defense is finding and shutting down fake accounts, the source of much of the inauthentic activity we see on Facebook. Authenticity matters because people need to trust that the content they're seeing is valid and they need to trust the connections they make. We are now blocking millions of attempts to register false accounts each and every day.

We're making progress on fake news. We're getting rid of the economic incentives to create it and we're limiting the distribution it gets on Facebook. We demote articles rated by third-party fact-checkers as false. We warn people who have shared them or who are about to share them, and we show them related articles to give them more facts.

We've also taken strong steps to prevent abuse and increase transparency in advertising. Today on Facebook, you can go to any page and see all the ads that page is running, even if they wouldn't be shown to you. For political and issue ads, you can also see who paid for the ads, how much was spent, and the demographics of the people who saw them.

We're also going to require people running large pages with large audiences in the United States to go through an authorization process and confirm their identity. These steps won't stop everyone who's trying to game the system, but they will make it a lot harder.

As these past few weeks and months have shown, this work is starting to pay off. In July, we removed 32 pages and accounts involved in coordinated, inauthentic behavior. In August, we removed 650 pages and accounts that originated in Iran, as well as additional pages and accounts from Russia. And just last week, we took down 58 pages and accounts from Myanmar, many of which were posing as news organizations.

We are focused, as I know you are, on the upcoming U.S. mid-terms and on elections around the world. Our efforts in recent elections from Germany, to Italy, to Mexico, to the Alabama special Senate election, show us that the investments we are making are yielding results. We also know, as Chairman Burr said, that we cannot stop interference by ourselves. We're working with outside experts, industry, partners and governments, including law enforcement, to share information about threats and prevent abuse.

We're getting better at finding and stopping our opponents, from financially motivated troll farms to sophisticated military intelligence operations. We don't have access to the intelligence governments have access to, so we don't always know exactly who is behind these attacks or their motives, and that's why we will continue working closely with law enforcement.

Chairman Burr, I want to thank you for your leadership. Vice Chairman Warner, I want to thank you for your white paper, which has so many ideas on how we can work together to strengthen our defense. Senators, let me be clear, we are more determined than our opponents and we will keep fighting.

When bad actors try to use our site, we will block them. When content violates our policies, we will take it down. And when our opponents use new techniques, we will share them so we can strengthen our collective efforts.

Everyone here today knows that this is an arms race, and that means we need to be ever more vigilant. As Chairman Burr has noted, nothing less than the integrity of our democratic institutions, processes, and ideals is at stake. We agree, and we will work with all of you to meet this challenge.

Thank you.

[The prepared statement of Ms. Sandberg follows:]

**HEARING BEFORE THE UNITED STATES SENATE SELECT COMMITTEE ON
INTELLIGENCE**

September 5, 2018

Testimony of Sheryl Sandberg
Chief Operating Officer, Facebook

I. INTRODUCTION

Chairman Burr, Vice Chairman Warner, and Members of the Select Committee on Intelligence, thank you for the invitation to participate in today's hearing on Foreign Influence Operations' Use of Social Media Platforms.

I appreciate the opportunity to explain how seriously Facebook takes the issue of election interference and update you on the steps we're taking to prevent it.

As this Committee's bipartisan report states, in January 2017, the CIA, NSA, and FBI "revealed key elements of a comprehensive and multifaceted Russian campaign against the United States." The Committee's subsequent investigation "has exposed a far more extensive Russian effort to manipulate social media outlets to sow discord and to interfere in the 2016 election and American society," as well as additional examples of Russia's attempts to "interfere in U.S. elections and those of our allies."

We were too slow to spot this and too slow to act. That's on us. This interference was completely unacceptable. It violated the values of our company and of the country we love.

The actions we've taken in response—beginning with the steps Facebook's General Counsel, Colin Stretch, outlined to this Committee last year—show our determination to do everything we can to stop this kind of interference from happening.

We're investing heavily in people and technology to keep our community safe and keep our service secure. This includes using artificial intelligence to help find bad content and locate bad actors. We're shutting down fake accounts and reducing the spread of false news. We've put in place new ad transparency policies, ad content restrictions, and documentation requirements for political ad buyers. We're getting better at anticipating risks and taking a broader view of our responsibilities. And we're working closely with law enforcement and our industry peers to share information and make progress together.

This work is starting to pay off. We're getting better at finding and combating our adversaries, from financially motivated troll farms to sophisticated military intelligence operations. We've removed hundreds of Pages and accounts involved in coordinated inauthentic behavior—meaning they misled others about who they were and what they were doing.

The threat we face is not new. America has always confronted attacks from opponents who wish to undermine our democracy. What is new are the tactics they use. That means it's going to take

everyone—including industry, governments, and experts from civil society—working together to stay ahead.

At its best, Facebook plays a positive role in our democratic process—and we know we have a responsibility to protect that process on our service. We’re investing for the long term because security is never a finished job. Our adversaries are determined, creative, and well-funded. But we are even more determined—and we will continue to fight back.

II. ASSESSING PAST RUSSIAN ATTEMPTS TO INFLUENCE ELECTIONS

As Facebook’s General Counsel emphasized in his November 2017 testimony before this Committee, our security team has been aware of traditional Russian cyber threats, such as hacking and malware, for many years. Before Election Day in November 2016, we detected and mitigated several threats from actors with ties to Russia. This included activity by APT28, a group that the U.S. government has publicly linked to Russian military intelligence services.

Although our primary focus was on these traditional threats, we also saw some new behavior—namely, the creation of fake personas that were then used to seed stolen information to journalists. Some of these fake personas were also linked to a Facebook Page called DC Leaks, which publicized an off-platform website of the same name that hosted stolen information. This activity violated our policies, and we removed the DC Leaks accounts.

After the election, we continued to investigate these new threats. We found that the Internet Research Agency (IRA), a Russian entity located in St. Petersburg, Russia, had used coordinated networks of fake Pages and accounts to interfere in the election: promoting or attacking candidates and causes, creating distrust in political institutions, and spreading discord. Our investigation demonstrated that the IRA did this by using both organic content and Facebook’s advertising tools.

We found that some 470 fake Pages and accounts associated with the IRA spent approximately \$100,000 on about 3,500 Facebook and Instagram ads between June 2015 and August 2017. Our analysis showed that these accounts used these ads to promote roughly 120 Facebook Pages that they had set up, which had posted more than 80,000 pieces of content between January 2015 and August 2017. We shut down the accounts and Pages we identified at the time that were still active. The Instagram accounts we deleted had posted about 120,000 pieces of content.

In April of this year, we took down more than 270 additional Pages and accounts controlled by the IRA that primarily targeted people living in Russia and Russian speakers around the world, including in countries neighboring Russia, such as Azerbaijan, Uzbekistan, and Ukraine. Some of the Pages we removed belonged to Russian news organizations that we determined were surreptitiously controlled by the IRA.

We continue to monitor our service for abuse and share information with law enforcement and others in our industry about these threats. Our understanding of overall Russian activity in 2016 is limited because we do not have access to the information or investigative tools that the U.S. government and this Committee have. We look forward to your final report and expect that your

findings and the information you share will help us further protect Facebook and those who use our service.

III. COMBATING FOREIGN ELECTION INTERFERENCE AND ADVANCING ELECTION INTEGRITY

We've made important changes and investments to improve our ability to detect and stop foreign election interference and strengthen the security of our platform. We have more than doubled the number of people working on safety and security and now have over 20,000. We review reports in over 50 languages, 24 hours a day. Better machine learning technology and artificial intelligence have also enabled us to be much more proactive in identifying abuse. We're focused on:

Removing Fake Accounts. One of the main ways we identify and stop foreign actors is by proactively detecting and removing fake accounts, since they're the source of much of the interference we see.

- We use both automated and manual review to detect and deactivate fake accounts, and we are taking steps to strengthen both. These systems analyze distinctive account characteristics and prioritize signals that are more difficult for bad actors to disguise.
- We block millions of attempts to register fake accounts every day. Globally, we disabled 1.27 billion fake accounts from October 2017 to March 2018. By using technology like machine learning, artificial intelligence, and computer vision, we can proactively detect more bad actors and take action more quickly.
- We're also investing heavily to keep bad content off our services. For example, we took down 836 million pieces of spam in the first quarter of 2018—much of it before it was reported to us.

Preventing Coordinated Inauthentic Behavior. Our Community Standards prohibit coordinated inauthentic behavior, which is when multiple accounts—including both fake and authentic accounts—work together to mislead people. This behavior is not allowed because we don't want organizations or individuals creating networks of accounts that misinform people about who they are or what they're doing.

- In July, we took down 283 Pages and accounts in Brazil that were using fake accounts to share disinformation ahead of the country's October elections.
- In July, we also removed 32 Pages and accounts from Facebook and Instagram because they were involved in coordinated inauthentic behavior. We're still investigating, but some of the activity is consistent with what we saw from the IRA before and after the 2016 elections, and we've found evidence of some connections between these accounts and IRA accounts we disabled last year. But there are differences, too. It's clear that whoever set up these accounts went to greater lengths to obscure their true identities than the IRA did in 2016.

- In August, we removed over 650 Pages and accounts from Facebook and Instagram that originated in Iran, as well as more Pages and accounts that can be linked to sources that the U.S. government has previously identified as Russian military intelligence services.
- We also took down over 50 Pages and accounts from Facebook in Myanmar for engaging in coordinated inauthentic behavior. We discovered that they used seemingly independent news and opinion Pages to covertly push the messages of the Myanmar military.

Although inauthentic actors continue to look for new ways to mislead people, we're taking steps to make this harder for them.

Tackling False News. We're working to stop the spread of false news. We partner with third-party fact-checking organizations to limit the spread of articles they rate as false, and we disrupt the economic incentives for traffickers of misinformation. We also invest in news literacy programs and work to inform people by providing more context on the stories they see.

- We have partnerships with independent third-party fact-checkers in 17 countries. Stories they rate as false 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 are also beginning to use machine learning to help identify and demote foreign Pages that are likely to spread financially-motivated hoaxes to people in other countries.
- We know that misinformation can be associated with harm, especially in places like Myanmar and Sri Lanka. In these cases, we are implementing a policy that allows us to remove misinformation that has the potential to contribute to imminent violence or physical harm.
- We're currently testing fact-checking for photos and videos in nine countries. This includes identifying visuals that have been manipulated (e.g., a video that is edited to show something that did not really happen) or taken out of context (e.g., a photo from a previous tragedy associated with a different, present day conflict).
- We know how important it is to empower people to decide for themselves what to read, trust, and share. We invest in promoting news literacy and provide people with more context around the news they see. For example, if third-party fact-checkers write articles providing more information about a news story, we show those articles immediately below the story. We've also started showing people more information about articles—such as the publisher's Wikipedia entry, related articles on the same topic, and information about how the article has been shared on Facebook.
- We notify people and Page Admins if they try to share a story, or have shared one in the past, that's been determined by third-party fact-checkers to be false.
- We're learning from academics, increasing our work with third-party fact-checkers and talking to other organizations about how we can work together.

- We are also working to detect false news on the state and local level. Ahead of the 2018 U.S. midterm elections, we're working with the Associated Press to use their reporters in all 50 states to identify and debunk false and misleading stories.

Increasing Ad Transparency. We've taken strong steps to prevent abuse and increase transparency in advertising.

- *Political Advertisements.* All politics and issue ads on Facebook and Instagram in the U.S. must be clearly labeled with a "Paid for by" disclosure at the top of the ad so people can see who is paying for them. This is especially important when the Page name doesn't match the name of the company or person funding the ad. We have also added new requirements for advertisers:
 - Any person who wants to run one of these ads must upload an identification document and confirm their identity. They also must prove they live in the U.S. by providing a residential mailing address. We then mail a letter with a code that the person must provide to us in order to become authorized to run ads with political content.
 - When people click on the "Paid for by" label, they'll be taken to an archive with more information. They will be able to see the ad campaign budget associated with an individual ad; how many people saw it; and the age, location and gender of the people who were shown the ad. The archive can be reached at <https://www.facebook.com/ads/archive>. People on Facebook visiting the archive can see and search ads with political or issue content an advertiser has run in the U.S. for up to seven years.
 - Enforcement of the new features and this policy, available at https://www.facebook.com/policies/ads/restricted_content/political, began in the United States on May 24, 2018.
- *View Active Ads.* Everyone can now see the ads every Page is currently running. People can log into Facebook, visit any Page, and select "Info and Ads." They will see ad creative and copy and can flag anything suspicious by clicking on "Report ad."
- *More Page Information.* People around the world can also learn more about Pages, even if they don't advertise. For example, they can see any recent name changes and the date the Page was created. We're also going to require people that run Pages with large audiences in the U.S. to go through an authorization process and confirm their location. We're going to make sure their Pages display more information, including the location of the people running the Page. This will make it much harder for people to run Pages using fake accounts, or to grow virally and spread misinformation or divisive content that way.

These steps by themselves won't stop all bad actors trying to game the system, but they will make it harder for them to succeed—and they will help prevent people from advertising in obscurity. Whenever we introduce new policies, we won't always get everything right, even in the long term. Election interference is a problem that's bigger than any one company, which is

why we support the Honest Ads Act. The changes we have made are consistent with the Act's objectives and the standards, and we're committed to working with Congress to help raise the bar for all political advertising online.

Preventing Foreign Interference Around the World. We're deploying new tools and teams to identify threats and support the electoral process in the run-up to specific elections.

- In Germany, we worked closely with the authorities to support election security. In Italy, we asked independent fact-checkers to go hunting for false stories. And ahead of the recent Mexican elections, we partnered with Google and others to fund an independent fact-checking organization, "Verificado 2018"; placed full-page ads in leading papers under the title "Tips to Detect Fake News"; and took down thousands of Pages, Groups, and accounts in Mexico and across Latin America because they were part of a broader network of coordinated behavior.
- We tested one of the tools we used to spot foreign interference during the Alabama Senate election and have since used it in other elections around the world.
- We also ran public service announcements about false news in 25 countries, including in advance of French, Kenyan, German, Italian, Turkish, Irish, and Mexican elections.

Maintaining Compliance Controls. We've created a strong program to ensure compliance with our legal obligations and support our efforts to prevent foreign interference and support election integrity.

- *Enforcing Compliance with Federal Law.* Facebook's compliance team maintains a Political Activities and Lobbying Policy that is available to all employees. This Policy is covered in our Code of Conduct training for all employees and includes guidelines to ensure compliance with the Federal Election Campaign Act.
- *Suspicious Activity Reporting.* We have processes designed to identify inauthentic and suspicious activity, and we maintain a sanctions compliance program to screen advertisers, partners, vendors, and others using our payment products. Our payments subsidiaries file Suspicious Activity Reports on developers of certain apps and take other steps as appropriate, including denying such apps access to our platforms.

Promoting Civic Engagement. Facebook helps representatives connect with their constituents, and helps people register to vote and learn more about the issues that matter to them. We believe we have a responsibility to build tools that support this civic engagement, and we provide them to the world for free.

- *Access to Information.* We're building products that make it easier for people to find information about where candidates and political parties stand on the issues they care about.
 - We launched the Issue Tab, which allows politicians' Pages to provide short, unfiltered statements in their own words about issues that are important to them

and their constituents. This was used in the run-up to the recent election in Mexico.

- We also introduced Ballot, which allows people to see who's running for office at different levels of government, visit the candidates' Pages to learn more about them, and compare the candidates' perspectives on issues.
- *Reminders to Register and Vote.* We are encouraging people who are eligible to register to vote, reminding people of deadlines and connecting them with non-partisan resources.
 - We've run voting registration reminders in the run-up to national elections in the U.S. and several other countries. We're also launching voter registration drives during the U.S. primaries in all states that require voter registration.
 - We show messages at the top of News Feed on Election Day in 66 countries reminding people to vote and helping them find their polling place. We also show these reminders for state, county, and municipal elections in the U.S.
 - Efforts like these helped more than 2 million people get registered to vote in the 2016 U.S. elections.
- *Supporting Independent Research.* We recently announced a new election research commission, named Social Science One, to provide independent, credible research about the role of social media in elections and in democracy.

IV. COMBATING TARGETED HACKING AND DATA COLLECTION

Alongside our work on elections, we're also strengthening our defenses against a broader set of threats.

Facebook has a security team dedicated to understanding how bad actors attack individuals and networks, building defenses against such attacks, and reacting quickly to mitigate potential damage. We also have a working group dedicated to detecting and mitigating attacks against high-profile users. In April 2017, we published a report on information operations, including targeted data collection.

Over the last several years, nation states and non-state actors have increased attacks against individuals' personal accounts—both email and social media—to steal information from them and the organizations with which they are affiliated. This includes attacks that use Facebook for reconnaissance and the delivery of malicious content, such as links to phishing sites and malware, and attacks meant to take over the accounts of targeted individuals.

We have detected and stopped multiple attacks aimed at U.S. and foreign interests. We notify individuals and the appropriate government authorities when these attempts are detected and share what we learn about the techniques and tools used with law enforcement and with our industry partners.

We have also implemented additional measures to protect people who are likely to be targeted in times of heightened cyber activity, including elections, periods of conflict or political turmoil, and other high-profile events:

- Building AI systems to detect and stop attempts to send malicious content;
- Providing customizable security and privacy features, including two-factor authentication options and marketing to encourage people to adopt them;
- Sending notifications to individuals if they have been targeted by sophisticated attackers, with custom recommendations depending on the threat model;
- Sending proactive notifications to people who have not yet been targeted, but may be at risk based on the behavior of particular malicious actors;
- Deploying AI systems to monitor login patterns and detect the signs of a successful account takeover campaign;
- When possible, communicating directly with likely targets and providing them with instructions on how to secure their account; and
- Where appropriate, working with government bodies responsible for elections to notify and educate people who may be at greater risk.

We are also aware that threat actors seek to use social media to target military personnel, and we have built new capabilities specifically to handle this category of threat:

- We've partnered with Blue Star Families and USAA to create an [online safety guide](#) for service members and their families.
- We recently released a [video PSA](#) to help people identify and report military scams.
- We train and advise military officials on best practices for maintaining secure accounts and Pages, including setting up [two-factor authentication](#) and managing [Page Roles](#).

We believe that our adversaries will continue to attempt operations that include both traditional techniques and online disinformation.

V. COOPERATION WITH GOVERNMENT ENTITIES, INDUSTRY, AND CIVIL SOCIETY

Because cyber threats constantly evolve, we all need to work together: industry, government, and experts from civil society. It's especially critical for companies and government to cooperate. We have worked successfully with the DOJ, the FBI, and other law enforcement agencies to address a wide variety of threats to our platform, and we are actively engaged with DHS and the FBI's new Foreign Influence Task Force focused on election integrity.

Our security team regularly conducts internal reviews to monitor for state-sponsored threats. We do not publicly disclose the elements of these reviews for security reasons, but they include monitoring and assessing thousands of account details, such as location information and connections to others on our platform. We are committed to keeping law enforcement apprised of these efforts.

Additionally, as part of official investigations, government officials sometimes request data about people who use Facebook. We have an easily accessible online portal and 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.

We are also working with the broader community to identify and combat threats. One example is our partnership with the Atlantic Council's Digital Forensic Research Lab, which is providing us with real-time updates on emerging threats and disinformation campaigns around the world. They assisted in our work around the Mexico election, our recent takedown of a financially motivated "like" farm in Brazil, and the accounts we recently disabled for coordinated inauthentic behavior here in the U.S.

We also partner with cybersecurity firms. In July, FireEye contacted us about a network of Pages and accounts originating from Iran that engaged in coordinated inauthentic behavior. Based on that tip, we started an investigation and identified and removed additional accounts and Pages from the network.

We share information about threats with a number of other tech companies to help combat those threats more effectively and recently organized several meetings with industry participants to more specifically discuss election protection efforts.

We also participate in discussions with governments around the world at key events such as the Munich Security Conference and CyCon, which is organized by the NATO Cooperative Cyber Defense Centre of Excellence.

We know we can't stop interference by ourselves. We don't have all the investigative tools that the government has, and we can't always attribute attacks or identify motives. But we will continue to work closely with law enforcement around the world and do everything we can to stop foreign election interference wherever it occurs on our platform.

We want to thank Chairman Burr for his leadership on this issue, and Vice Chairman Warner for his recent white paper and his ideas about strengthening election security online. We look forward to continuing our work with this Committee.

VI. CONCLUSION

What happened in the 2016 election cycle was unacceptable. Any attempt to use our platform to interfere in elections runs counter to everything Facebook stands for. People come to Facebook every day to have authentic conversations and to share, not to be deceived or misled.

We are learning from what happened, and we are improving. When we find bad actors, we will block them. When we find content that violates our policies, we will take it down. And when our attackers use new techniques, we'll share them to improve our collective defense. We are even more determined than our adversaries, and we will continue to fight back.

This is an arms race, and that means we need to be ever more vigilant. As Chairman Burr has noted, "Nothing less than the integrity of our democratic institutions, processes and ideals is at stake." We agree, and we are determined to meet this challenge.

Chairman BURR. Thank you, Ms. Sandberg. Mr. Dorsey, the floor is yours.

**STATEMENT OF JACK DORSEY, CHIEF EXECUTIVE OFFICER,
TWITTER, INC.**

Mr. DORSEY. Thank you Chairman Burr, Vice Chairman Warner and the committee for the opportunity—for the opportunity to speak on behalf of Twitter to the American people. I look forward to our conversation about the work we're doing to help protect the integrity of U.S. elections and elections around the world.

I am someone of very few words and typically pretty shy, but I realize how important it is to speak up now. If it's OK with all of you I'd like to read you something I personally wrote as I considered these issues. I'm also going to tweet this out now.

First, I want to step back and share our view of Twitter's role in the world. We believe many people use Twitter as a digital public square. They gather from all around the world to see what's happening and have a conversation about what they see. In any public space you will find inspired ideas and you'll find lies and deception—people who want to help others and unify, and people who want to hurt others and themselves, and divide.

What separates a physical and digital public space is greater accessibility and velocity. We're extremely proud of helping to increase the accessibility and velocity of a simple, free, and open exchange. We believe people would learn faster by being exposed to a wide range of opinions and ideas, and it helps make our Nation and the world feel a little bit smaller. We aren't proud of how that free and open exchange has been weaponized and used to distract and divide people and our Nation. We found ourselves unprepared and ill-equipped for the immensity of the problems that we have acknowledged: abuse, harassment, troll armies, propaganda through bots and human coordination, misinformation campaigns, and divisive filter bubbles. That's not a healthy public square. Worse, a relatively small number of bad faith actors were able to game Twitter to have an outsized impact.

Our interests are aligned with the American people and this committee. If we don't find scalable solutions to the problems we're now seeing, we lose our business and we continue to threaten the original privilege and liberty we were given to create Twitter in the first place.

We weren't expecting any of this when we created Twitter over 12 years ago. We acknowledge the real world negative consequences of what happened and we take the full responsibility to fix it. We can't do this alone and that's why this conversation is important and why I am here.

We've made significant progress recently on tactical solutions like identification of many forms of manipulation intending to artificially amplify information, more transparency around who buys ads and how they are targeted, and challenging suspicious logins and account creation. We've seen positive results from our work. We're now removing over 200 percent more accounts for violating our policies. We're identifying and challenging 8 to 10 million suspicious accounts every week, and we're thwarting over a half million accounts from logging in to Twitter every single day.

We've learned from 2016, and more recently from other nations' elections, how to protect the integrity of elections: better tools, stronger policy, and new partnerships are already in place. We intend to understand the efficacy of these measures to continue to get better, but we all have to think a lot bigger than decades past, today. We must ask the question, what is Twitter incentivizing people to do, or not do, and why? The answers will lead to tectonic shifts in Twitter and how our industry operates. Required changes won't be fast or easy.

Today we're committing to the people and this committee to do that work and do it openly. We're here to contribute to a healthy public square, not compete to have the only one. We know that's the only way our business thrives and helps us all defend against these new threats.

In closing, when I think of my work, I think of my mom and dad in St. Louis, a Democrat and a Republican. For them, Twitter has always been a source of joy, a source of learning, and a source of connection to something bigger than themselves. They're proud of me, proud of Twitter, and proud of what made it all possible. What made it possible was the fact that I was born into a Nation built by the people for the benefit of the people—where I could work hard to make something happen which was bigger than me. I treasure that and will do everything in my power to protect it from harm. Thank you.

[The prepared statement of Mr. Dorsey follows:]

United States Senate Select Committee on Intelligence

**Testimony of Jack Dorsey
Chief Executive Officer
Twitter, Inc.**

September 5, 2018

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

I am grateful for the opportunity to appear here today.

The purpose of Twitter is to serve the public conversation. We serve our global audience by focusing on the needs of the people who use our service, and we put them first in every step we take. We want to be a global town square, where people from around the world come together in an open and free exchange of ideas. We must be a trusted and healthy place that supports free and open democratic debate.

Twitter is committed to improving the collective health, openness, and civility of public conversation on our platform. Twitter's is built and measured by how we help encourage more healthy debate, conversations, and critical thinking. Conversely, abuse, malicious automation, and manipulation detracts from it. We are committing Twitter to hold ourselves publicly accountable towards progress.

The public conversation occurring on Twitter is never more important than during elections, the cornerstone of our democracy. Our service shows the world what is happening, democratizes access to information and—at its best—provides people insights into a diversity of perspectives on critical issues; all in real-time. We work with commitment and passion to do right by the people who use Twitter and the broader public. Any attempts to undermine the integrity of our service is antithetical to our fundamental rights and undermines the core tenets of freedom of expression, the value upon which our company is based. This issue affects all of us and is one that we care deeply about as individuals, both inside and outside the company.

We appreciate the continued partnership with the Committee, and we share your concern about malicious foreign efforts to manipulate and divide people in the United States and throughout the world. We have implemented significant improvements since we last appeared before the Committee in November, and we will continue to undertake important steps in the coming months and years.

I look forward to sharing our work with the members of this Committee and listening to your recommendations on how best to increase the health of our platform and its role in our democracy from manipulation by hostile foreign actors.

From Twitter's perspective, this threat is not limited solely to elections or politics. Instead, we view it as a challenge to the fundamental health of our platform, and by extension, to

the global public conversation that Twitter serves. We commit to continuing to confront that challenge together.

I. RUSSIAN INTERFERENCE IN THE 2016 ELECTION AND LESSONS LEARNED

Twitter continues to engage in intensive efforts to identify and combat state-sponsored hostile attempts to abuse social media for manipulative and divisive purposes. We now possess a deeper understanding of both the scope and tactics used by malicious actors to manipulate our platform and sow division across Twitter more broadly. Our efforts enable Twitter to fight this threat while maintaining the integrity of peoples' experience on the service and supporting the health of conversation on our platform. Our work on this issue is not done, nor will it ever be. The threat we face requires extensive partnership and collaboration with our government partners and industry peers. We each possess information the other does not have, and our combined information is more powerful in combating these threats together.

A. Retrospective Review

Last fall, we conducted a comprehensive retrospective review of platform activity related to the 2016 election. To better understand the nature of the threat and ways to address future attempts at manipulation, we examined activity on the platform during a 10-week period preceding and immediately following the 2016 election (September 1, 2016 to November 15, 2016). We focused on identifying accounts that were automated, linked to Russia, trying to get unearned attention, and Tweeting election-related content, and we compared activity by those accounts to the overall activity on the platform. We reported the results of that analysis in November 2017, and we updated the Committee in January 2018 about the findings from our ongoing review.

As we reported in January 2018, we identified 50,258 automated accounts that were Russian-linked and Tweeting election-related content, representing less than two one-hundredths of a percent (0.016%) of the total accounts on Twitter at the time. Of all election-related Tweets that occurred on Twitter during that period, these malicious accounts constituted approximately one percent (1.00%), totaling 2.12 million Tweets. Additionally, in the aggregate, automated, Russian-linked, election-related Tweets from these malicious accounts generated significantly fewer impressions (*i.e.*, views by others on Twitter) relative to their volume on the platform. Additional information on the accounts associated with the Internet Research Agency is included below.

Twitter is committed to ensuring that promoted accounts and paid advertisements are free from hostile foreign influence. In connection with the work we did in the fall, we conducted a comprehensive analysis of accounts that promoted election-related Tweets on the platform throughout 2016 in the form of paid ads. We reviewed nearly 6,500 accounts and our findings showed that approximately one-tenth of one-percent—only nine of the total number of accounts—were Tweeting election-related content and linked to Russia. The two most active accounts out of those nine were affiliated with Russia Today (“RT”), which Twitter subsequently

barred from advertising on Twitter. And Twitter is donating the \$1.9 million that RT spent globally on advertising to academic research into election and civic engagement.

B. Insights from Our Review

Although the volume of malicious election-related activity that we could link to Russia was relatively small, we strongly believe that any such activity on Twitter is unacceptable. We remain vigilant about identifying and eliminating abuse on the platform perpetrated by hostile foreign actors, and we will continue to invest in resources and leverage our technological capabilities to do so. Twitter's main focus is promoting healthy public discourse through protection of the democratic process. Tied to this is our commitment to providing tools for journalism to flourish by creating and maintaining a platform that helps to provide people with high-quality, authentic information in a healthy and safe environment.

We also recognize that, as a private company, there are threats that we cannot understand and address alone. We must continue to work together with our elected officials, government partners, industry peers, outside experts, and other stakeholders so that the American people and the global community can understand the full context in which these threats arise.

II. IMPROVEMENTS TO TWITTER

We have made the health of Twitter our top priority, and our efforts will be measured by how we help encourage more healthy debate, conversations, and critical thinking on the platform. Conversely, abuse, automation, and manipulation will detract from the health of our platform. Twitter recently developed and launched more than 30 policy and product changes designed to foster information integrity and protect the people who use our service from abuse and malicious automation. Twitter has made a number of improvements specifically in preparation for the 2018 election, described below.

A. Combating Malicious Automation and Protecting Conversation Health

Using the insights from our retrospective review, Twitter continues to develop the detection tools and systems needed to combat malicious automation on our platform. Twitter has refined its detection systems. Twitter prioritizes identifying suspicious account activity, such as exceptionally high-volume Tweeting with the same hashtag or mentioning the same @handle without a reply from the account being addressed, and requires an individual using the platform to confirm control. Twitter has also increased its use of challenges intended to catch automated accounts, such as reCAPTCHAs, that require users to identify portions of an image or type in words displayed on screen, and password reset requests that protect potentially compromised accounts. Twitter is also in the process of implementing mandatory email or cell phone verification for all new accounts.

Our efforts have been effective. Due to technology and process improvements, we are now removing 214% more accounts year-over-year for violating our platform manipulation policies. For example, over the course of the last several months, our systems identified and challenged between 8.5 million and 10 million accounts each week suspected of misusing

automation or producing spam. Spam can be generally described as unsolicited, repeated actions that negatively impact other people. This includes many forms of automated account interactions and behaviors as well as attempts to mislead or deceive people. This constitutes more than three times the 3.2 million we were catching in September 2017. We thwart 530,000 suspicious logins a day, approximately double the amount of logins that we detected a year ago.

These technological improvements have brought about a corresponding reduction in the number of spam reports from people on Twitter, evidence to us that our systems' ability to automatically detect more malicious accounts and potential bad faith actors than they did in the past. We received approximately 25,000 such reports per day in March of this year; that number decreased to 17,000 in August.

We also removed locked accounts from people's follower counts, to ensure these figures are more reliable. Accounts are locked when our systems detect unusual activity and force a password change or other challenge. If the challenge has not been met or the password has not been changed within a month, the account is locked, barring it from sending Tweets, Retweets or liking posts from others.

B. Corporate Reorganization and Formation of a Dedicated Cross-Functional Analytical Team

Our improvements include important structural changes. I recently reorganized the structure of the company to allow our valued employees greater durability, agility, invention, and entrepreneurial drive. The reorganization simplified the way we work, and enabled all of us to focus on health of our platform.

In particular, we have created an internal cross-functional analytical team whose mission is to monitor site and platform integrity. Drawing on expertise across the company, the analytical team can respond immediately to escalations of inauthentic, malicious automated or human-coordinated activity on the platform. The team's work enables us to better understand the nature of the malicious activity and mitigate it more quickly.

To supplement its own analyses, Twitter's analytical team also receives and responds to reports from across the company and from external third parties. The results from all of the team's analyses are shared with key stakeholders at Twitter and provide the basis for policy changes and product initiatives and removal of accounts.

The primary focus of the cross-functional analytical team is election readiness. Leading up to and during the 2018 election period, the team will examine, respond to, and escalate instances of suspected inauthentic, election-related coordinated activity in political conversation and conduct in-depth analyses of relevant Twitter data.

C. Political Conversations Dashboard

Our cross-functional team has developed a political conversations dashboard to evaluate the integrity of political conversations on the platform in the aggregate, focusing primarily (but

not exclusively) on elections in the United States in the near term. For example, this dashboard surfaces information about sudden shifts in sentiment around a specific conversation, suggesting a potential coordinated campaign of activity, as well as information about groups of potentially linked accounts that are posting about the same topic.

Through real-time review and detection of anomalous and potentially malicious automated or human-coordinated activity, the team will work to identify and address any attempts by bad faith actors to interfere with the electoral process, and will be better informed about where and how to deploy resources to proactively review potential malicious activity. Accounts will be escalated for review in real-time if exhibiting anomalous patterns of behavior. These efforts will significantly improve our ability to detect malicious automated and human-coordinated activity surrounding political content as well as the speed with which we address those issues.

D. Candidate Verification

Twitter serves the public conversation by promoting health and earning the trust of the people who use our service. We cannot succeed unless the American people have confidence in the integrity of the information found on the platform, especially with respect to information relevant to elections and the democratic process. To promote transparency and assist our stakeholders in identifying messages from elected officials and those who are running for office, we have made a concerted effort to verify all major party candidates for both federal and key state positions. Through verification – a blue badge that appears next to a person’s Twitter handle throughout the platform – we let people know that accounts of public interest are the authentic accounts (as opposed to impersonation or parody accounts).

E. Election Labels

In addition, we have developed a new U.S. election label to identify political candidates. The label includes information about the office the candidate is running for, the state the office is located in, and the district number, if applicable. Accounts of candidates who have qualified for the general election and who are running for governor or for the U.S. Senate or House of Representatives will display an icon of a government building. These new features are designed to instill confidence that the content people are viewing is reliable and accurately reflects candidates’ and elected officials’ positions and opinions.

F. Advertising and Promoted Content

As we learned from our 2016 retrospective review and the important work of your Committee, bad faith actors have attempted to influence the electoral process by propagating paid content on the platform, including political advertisements and promoted Tweets. As we reported in the fall, we have devoted considerable resources to increasing transparency and promoting accountability in the ads served to Twitter customers.

Twitter first implemented an updated Political Campaigning Policy to provide clearer guidance about how we define political content and who can promote-political content on our

platform. Under the revised policy, advertisers who wish to target the United States with federal political campaigning advertisements are required to self-identify as such and certify that they are located within the United States. Foreign nationals will not be permitted to serve political ads to individuals who identify as being located in the United States.

Twitter accounts that wish to target the U.S. with federal political campaigning advertisements must also comply with a strict set of requirements. Among other things, the account's profile photo, header photo, and website must be identical to the individual's or organization's online presence. In addition, the advertiser must take steps to verify that the address used to serve advertisements with content related to a federal political campaign is genuine.

To further increase transparency and better educate those who access promoted content, accounts serving ads with content related to a federal political campaign will now be visually identified and contain a disclaimer. This feature will allow people to more easily identify federal political campaign advertisements, quickly identify the identity of the account funding the advertisement, and immediately tell whether it was authorized by the candidate.

In June, we launched the Ads Transparency Center, which is open to everyone on Twitter and the general public, and currently focuses on electioneering communications. Twitter requires extensive information disclosures of any account involved in federal electioneering communications and provides specific information to the public via the Ads Transparency Center, including:

- Purchases made by a specific account;
- All past and current ads served on the platform for a specific account;
- Targeting criteria and results for each advertisement;
- Number of views each advertisement received; and
- Certain billing information associated with the account.

These are meaningful steps that will enhance the Twitter experience and protect the health of political conversations on the platform.

In addition, we recently announced the next phase of our efforts to provide transparency with the launch of a U.S.-specific Issue Ads Policy and certification process. The new policy impacts advertisements that refer to an election or a clearly identified candidate or advertisements that advocate for legislative issues of national importance. To provide people with additional information about individuals or organizations promoting issue ads, Twitter has established a process that verifies an advertiser's identity and location within the United States. These advertisements will also be included in the Ads Transparency Center. We are also

examining how to adopt political campaigning and issue ads policies globally. We remain committed to continuing to improve and invest resources in this space.

G. Engagement with Key Stakeholders

Information sharing and collaboration are critical to Twitter's success in preventing hostile foreign actors from disrupting meaningful political conversations on the platform. We recognize the value of inputs we receive from our industry peers about hostile foreign actors. We have shared and remain committed to sharing information across platforms to better understand and address the threat of hostile foreign interference with the electoral process. On August 24, 2018, Twitter hosted our industry peers to discuss data sharing about hostile foreign actors regarding 2018 election security.

We also have well-established relationships with law enforcement agencies active in this arena, including the Federal Bureau of Investigation Foreign Influence Task Force and the Department of Homeland Security's Election Security Task Force. We look forward to continued cooperation with them on these issues, as only they have access to information critical to our joint efforts to stop bad faith actors.

Additionally, to further promote information sharing and to tap into the experience and expertise of active stakeholders, we recently updated a Partner Support Portal. Our goal is to expedite our response to reports from people active in the election arena. This includes election support organizations, U.S.-based research organizations, and universities and academics who study the spread of misinformation in the media. Reports from accounts within this select group are expedited and can be actioned promptly.

Consistent with our longstanding commitment to serving the public conversation, we partnered with experts at the University of Oxford and Leiden University to better evaluate our work on conversation health, focusing on informational echo chambers and unhealthy discourse on Twitter. This collaboration will also enable us to study how exposure to a variety of perspectives and opinions serves to reduce overall prejudice and discrimination. While looking at political discussions, these projects do not focus on any particular ideological group and the outcomes will be published in full in due course for further discussion.

Last October, Twitter barred advertising from Russia Today and Sputnik, both of which the U.S. Intelligence Community determined to have interfered with the election on behalf of the Russian government. We also devoted the \$1.9 million these accounts spent on the platform to research. The first recipients of those funds include the Kofi Annan Foundation's Global Commission on Elections, Democracy, and Security, the Atlantic Council, the EU DisinfoLab and the Reporters Committee for Press Freedom.

We also collaborate with a number of non-governmental organizations that are focused on voter registration, civic engagement, and media literacy, including RockTheVote, Democracy Works, TurboVote Challenge, HeadCount, DoSomething, and Ballotpedia.

H. Additional Safety Measures for Accessing Public Tweet Data

To further address malicious automation and abuse on the platform, we have also recently updated our developer policies, which govern the access and use of public Tweet data made available to developers and other third parties through our application programming interfaces (“APIs”).

We recognize that access to that data could be manipulated, so we have taken steps to prevent the use of our APIs for products and services that are abusive or that disrupt the health of conversations. Those to whom we grant access to our APIs are prohibited from using the data to manipulate conversations or otherwise abuse the data. Between April and June 2018 alone we removed more than 143,000 applications that we determined to be in violation of our developer policies. Most violated our policies against producing spam via APIs. And we continue to invest in and improve our detection tools to stop misuse of public Twitter data.

In July 2018, we introduced a new measure designed to increase developers’ accountability for applications that create and engage with Twitter content and accounts. Twitter now reviews and conducts compliance checks of all developers’ stated use of the data that they wish to access. We have also added new protections aimed to prevent the registration of low quality and spam-generating applications. We believe that these additional steps will help protect the integrity of our platform.

III. RECENT ACTIVITY ON THE PLATFORM

Twitter continues to see bad faith actors continue their attempts to manipulate and divide people on Twitter. Two such examples include recent activity related to new malicious activity by the Russian Internet Research Agency and malicious accounts located in Iran.

A. Malicious Accounts Affiliated with the Russian Internet Research Agency

Twitter has seen recent activity on the platform affiliated with the Russian Internet Research Agency. As we reported to the Committee in January 2018, we continue to identify accounts that we believe may be linked to the Internet Research Agency (“IRA”). As of today, we have suspended a total of 3,843 accounts we believe are linked to the IRA. And we continue to build on our contextual understanding of these accounts to improve our ability to find and suspend this activity as quickly as possible in the future, particularly as groups such as the IRA evolve their practices in response to suspension efforts across the industry.

As an example of Twitter’s ongoing efforts, Twitter identified 18 accounts in March 2018 we believe to be linked to the Internet Research Agency uncovered by our ongoing additional reviews. These accounts were created and registered after the 2016 election. These accounts used false identifies purporting to be Americans, and created personas focused on divisive social and political issues. The accounts represented both sides of the political spectrum. We continue to work with our law enforcement partners on this investigation.

B. Malicious Accounts Located in Iran

In August 2018, we were notified by an industry peer about possible malicious activity on their platform. After receiving information from them, we began an investigation on our platform to build out our understanding of these networks. We immediately notified law enforcement on this matter as soon as we discovered malicious activity.

We initially identified accounts based on indicators such as phone numbers and email addresses. Some of these accounts appeared to pretend to be U.S. person and discuss U.S. social commentary. In most cases, the accounts that appeared to suggest a U.S. affiliation or target U.S. person were created after the 2016 election. These accounts were in violation of our platform manipulation policies, and were engaged in coordinated activity intended to propagate messages artificially across accounts.

These accounts appear to be located in Iran. This is indicated by, for example, accounts related by an Iranian mobile carrier or phone number or Iranian email address on the account. Although Twitter is blocked in Iran, we may see people active on our service via a virtual private network, or VPN.

We suspended 770 accounts for violating Twitter policies. Fewer than 100 of the 770 suspended accounts claimed to be located in the U.S. and many of these were sharing divisive social commentary. On average, these 100 accounts Tweeted 867 times, were followed by 1,268 accounts, and were less than a year old. One advertiser ran \$30 in ads in 2017. Those ads did not target the U.S. and the billing address was located outside of Iran. We will remain engaged with law enforcement on this issue.

Twitter has been in close contact with our industry peers on this matter and received detailed information from them about the malicious accounts located with Iran, which has assisted us in our investigation, and we have shared our own details and work with other companies. We expect this process will continue and that the industry can continue to build on this effort and assist with this ongoing investigation.

* * *

Our core mission is to serve the public conversation. It is why we exist. We must promote and maintain the health of that conversation. The people who use our service must have confidence in the integrity of the information found on the platform, especially with respect to information relevant to elections and the democratic process. In taking the steps I have outlined above, we continue our efforts to address those threats posed by hostile foreign governments and foster an environment conducive to healthy, meaningful conversations on our platform. This work is essential, and today's hearing better equips us to confront this new threat to our platform and democracies across the globe.

I look forward to answering your questions.

Chairman BURR. Jack, thank you very much for that testimony and I might add that the Vice Chairman and I commented as you grow older, you will find a need for a bigger device to go to your notes on than that small one. We have a hard time with the small devices.

For members, we will do seven minute question rounds today. For planning purposes, we will break at approximately 10:45 for five minutes just to let our witnesses stretch and take a breath. And we will limit today's hearing to one round. We'll try to accommodate any members that might be caught in the Judiciary Committee but want to try to get back, but I know that they've got their own challenges. With that, I would recognize myself for seven minutes.

This question is to both of you. How would you define social media for this committee and more importantly for the American people? And I will start with you, Ms. Sandberg.

Ms. SANDBERG. Social media enables you to share what you want to share when you want to share it, without asking permission from anyone. And that's how we meet our mission, which is giving people a voice. And I think what's more important than just the content people share, is the connections they make. Social media enables people to celebrate their birthdays. In the last year, people have raised \$300 million on Facebook on birthday funders for non-profits they care about. Safety check: Millions of people in the worst circumstances of their lives have let their loved ones know they're safe. And small businesses to grow. All around the country I meet with small businesses, from a woman making dresses in her living room and selling them on Instagram, to a local plumber, who are able to find their customers on Facebook and then able to grow and hire people and live their American dream.

Chairman BURR. Jack.

Mr. DORSEY. I believe it's really important to—to understand how the people see it. And we believe that the people use Twitter as they would a public square and they often have the same expectations that they would have of any public space. For our part, we see our platform as hosting and serving conversations. Those conversations are in the public. We think there's a lot of benefit to those conversations being in the public, but there's obviously a lot of risks as well.

We see that news and entertainment are actually byproducts of public conversation. And we see our role as helping to not only serve that public conversation so that everyone can benefit, even if they don't have a Twitter account, but also to increase the health of that conversation as well. And in order to do that, we need to be able to measure it. We need to understand what healthy participation looks like in a public square, and we need to amplify that. And more importantly, we need to question a lot of the fundamentals that we started with 12 years ago in the form of incentives. When people use our product every single day—when they open our app up—what are we incentivizing them to do? Not telling them what to do, but what are we actually incentivizing them to do? And that certainly speaks to the buttons that we have in our service, all the way to our business model.

Chairman BURR. Ms. Sandberg, this question is for you. One root problem that we see is that users don't truly understand the types of data that are being collected on and off your platform. How is that data shared with advertisers or others to deliver targeted advertising and what vetting, if any, do you do on targeted advertising to prevent hostile actors from targeting your users for their products?

Ms. SANDBERG. Senator, it's a really important question because it goes to the heart of our service. We sell ads and we use information that people share with us or share with third-party sites to make those ads relevant to them. But privacy and advertising are not at odds. In fact, they go together. When people share information with us, we do not give it to advertisers without their permission. We never sell data. And they have control over the information we use.

Chairman BURR. Again for both of you, and I'll start with you, Mr. Dorsey. What's your company's ability to collaborate with other social media companies in this space?

Mr. DORSEY. We have a real openness to this and we have established a more regular cadence with our industry peers. We do believe that we have an opportunity to not only create more transparency with an eye towards more accountability, but also a more open way of working and a way of working that, for instance, allows for a review period by the public on how we think about our policies.

But more so, taking some of the lessons that we have learned and benefited from in the open-source software space to actually think about developing our policies, our enforcement, and also our products going forward. We've been experimenting a little bit with this recently, but we would like to be a company that is not only hosting an open conversation but is also participating in that open conversation. So, we're more than open to more collaboration, and not just with our industry peers but with scholars, academics, and also our government partners.

Chairman BURR. Thank you.

Ms. Sandberg.

Ms. SANDBERG. I think our collaboration has greatly increased. We've always worked closely with law enforcement and we continue to do that and particularly the FBI's new task force. We've always shared information with other companies but I think we are doing better and we can continue to do better.

Mr. Chairman, you noted in your opening remarks that some of the tips we got came from a private security firm. In our mind that's the system working. Our opponents are very well-funded. They are very organized, and we are going to get those tips from law enforcement, from each other, from private firms. And the faster we can collaborate, the faster we share those tips with each other, the stronger our collective defenses will be.

Chairman BURR. Last question from the Chair—again for both of you and I'll go in reverse—you first, Ms. Sandberg. If a foreign-influence campaign is detected among your platforms, is there a defined process by which other platforms are alerted to the campaign that you've discovered?

Ms. SANDBERG. Our security teams have been in close contact and so right now when we find something, we are reaching out to our companies—other companies to do it and working more closely together.

We've been talking about how, I think, there's still room for improvement there. I think we can do more to formalize the process. We've had a series of meetings and I think we're going to continue to work and we can do better.

Chairman BURR. Mr. Dorsey.

Mr. DORSEY. This is not something we want to compete on. We hosted our peer companies at our offices just in the past two weeks on this very topic and helping to increase our cadence of meeting and also what we can share. If there were an occurrence, we would immediately look to alert our peer companies and this committee and our government law enforcement partners.

Chairman BURR. Thank you for that. Let me just say in closing that I hope both of you, if you see impediments that exist in your ability to notify or to collaborate as it relates to nefarious actors, that you'll certainly make this committee aware in cases where we can help. With that, Vice Chairman.

Vice Chairman WARNER. Thank you, Mr. Chairman. As I indicated in my opening statement, I hope we can move forward on the policy discussion, so I'd like to get your thoughts on some of the ideas I and others have suggested, and I want to start with you, Mr. Dorsey.

I think after some initial false starts, it does really appear that you have committed to a shift in your company's culture with respect to the safety and security on your platform. Obviously, I have been impressed by some of the increasing efforts you've taken. A question I have, though, is that obviously on your platform there are a lot of automated accounts or bots, and there's nothing inherently good or bad about an automated account. As a matter of fact, there are certain very good things that come out of some of these automated accounts. But, do you believe that an individual Twitter user should have the right to know when he or she is being contacted, whether that contact is initiated by a human being or a bot?

Mr. DORSEY. I do believe that first and foremost, anyone using Twitter has the right to more context around not only the accounts that they're seeing, but also the information.

Vice Chairman WARNER. Would that go as far as actually having a policy on your platform indicating—I wouldn't ask you to take them down—but at least allowing the user to know whether that contact was initiated by a human being versus a machine?

Mr. DORSEY. As far as we can detect them. We can certainly label and add context to accounts that come through our API. Where it becomes a lot trickier is where automation is actually scripting our website to look like a human actor. So as far as we can label—and we can identify these automations—we can label them, and I think that is useful context and it's an idea that we have been considering over the past few months. It's really a question of the implementation, but we are interested in it and we are going to do something along those lines.

Vice Chairman WARNER. It's not going to solve the problem, but I do think giving that indication to users would allow them then

perhaps to make a little more judgment. Because we had, for example, back in early August, we had a panel of experts, and they were saying that some of the content—in terms of political content, I'm not talking about total tweets—but total political content was 25 to 30 to 1 on the far left and far right generated by either foreign actors or automated accounts. And my question is: Doesn't that volume on the extremes drown out real conversation and political conversation amongst Americans, regardless of where they fall on the political spectrum?

Mr. DORSEY. It does, in the shared areas of Twitter. So there are two main categories of usage in Twitter. One, is the people you follow, and those Tweets end up in your timeline. Two, are the more common shared spaces, like Search, Trends, and also Replies. That's where anyone could interject themselves, and that's where we see the most gaming of our systems, and that's where we've also made the most progress in terms of identifying these patterns and shutting them down before they spread too far. That is independent of our work on automation, because we're seeing the same patterns through human coordination as well.

Vice Chairman WARNER. I appreciate your comments about the willingness to notify a user whether it's a human being or a machine contacting you. I also think that there's room for improvement on some of the high volume Twitter accounts, to really do a little bit of extra examination.

Ms. Sandberg, let me move to you. Obviously, in a digital economy, I think data increasingly represents the single greatest asset you have. Obviously it's a part of the advertising model that you've created.

But I think most users are actually pretty much in the dark about how much data is actually being collected on them, what it's actually worth. I think as we've seen from other fields, like health care, the fact that we have such a lack of price transparency really makes health care reform really challenging.

I think some of that lack of price transparency and value within social media also exists, so I'd like to first of all ask, does a Facebook user have a right to know what information you are collecting about that user?

Ms. SANDBERG. Yes, and we really agree with you that people who use Facebook should understand what information is being used, how it's used, and the controls they have. We've worked hard to simplify this. We've put out things like privacy shortcuts, which show you all your settings in one place, and something called download your information, where you can download all of your information in a portable way and be able to take it with you and see what it is.

Vice Chairman WARNER. I understand, and I think you're making progress there, but again, if a user has that information, he or she may not know the value. Wouldn't it be actually helpful to your user to actually be able to then put some valuation on the data you're collecting from the user and publish that in a way so that people actually know what their information is worth?

Ms. SANDBERG. Mr. Vice Chairman, I think this is one of the proposals you laid out in your white paper, and like all of this, you know, we don't think it's a question of whether regulation—we

think it's a question of the right regulation that supports users, is transparent, and doesn't squash innovation. And we're happy to work with you on the proposal.

Vice Chairman WARNER. Well, I just think it's that more price transparency is always better, and I think this would be something that would help users sort through. There was another question that we've talked in the past about: Is there anything, even with a willing user, are there any rights or details about an individual user that they should not be able to give up or consent to having used?

Ms. SANDBERG. I'm sorry, I don't understand the question.

Vice Chairman WARNER. My question is this: At some point, are there certain pieces of personalized information that a user shouldn't be able to voluntarily give to an enterprise like yours or Twitter?

Ms. SANDBERG. I think there are, and I think there are many ways users have control over what they do. I also think there are probably corner cases of law enforcement holds or security matters where information is critically important.

Vice Chairman WARNER. I just wonder whether—just a question of whether you can consent away all of your rights—ought to be something we ought to have a discussion on. I've only got a few more seconds.

Let me ask, Ms. Sandberg, you made mention in your opening testimony the fact that sometimes political actors are using the platforms really to incite violence. I think you made at least some mention of Myanmar, where we've obviously seen a great tragedy take place there, where hundreds of thousands of Rohingya Muslims are fleeing in many ways. The U.N. High Commissioner has said that fake accounts on Facebook have incited that violence.

Do you believe that Facebook has both a moral obligation and potentially even a legal obligation to take down accounts that are actually incentivizing violence?

Ms. SANDBERG. I strongly believe that. In the case of what's happened in Myanmar, it's devastating, and we're taking aggressive steps and we know we need to do more. Probably the most important thing we've done is ramped up our ability to review reports in Burmese.

Vice Chairman WARNER. I appreciate your comment that Facebook would have both a moral and legal obligation, so sorting through what that would look like so that if there were other platforms that weren't being as responsible, there ought to be some sanctions. So I look forward to working with you on that issue as well.

Thank you, Mr. Chairman.

Chairman BURR. Senator Risch.

Senator RISCH. Thank you. Thank you both for being here today. This is, I think, the third hearing we've held over the last year or so—fourth—the Chairman says the fourth—that we've had on this issue.

I think the problem is really well laid out. We've spent hours and hours and hours talking about this and what the issues are and what the problems—I'm still not hearing what—very specifically how we're getting after this. I know there're some things being

done. I tend to agree with you that no matter what's done, as long as these platforms are there, there's going to be people finding their way into it to do bad things. And obviously, everybody wants to get that reduced as much as possible.

And I'm glad to hear that you and the entire industry are trying to do something about this. The entity up here that I serve in, there are lots of people that would love to help you run your organizations through what we call the regulatory process. That isn't all of them, obviously, and hopefully it isn't even a majority of them, but there will be—and you've already seen efforts in that regard—but you're going to have to do things yourselves to try to get around this so that we don't have the horrible things happen that spawn that type of regulation.

I want to drill down a little bit. In each of your companies, who sets these standards or the description of what a coordinated manipulation or inauthentic behavior is? What entity do you have in each of your companies who make these determinations?

Ms. Sandberg, let me start with you.

Ms. SANDBERG. Our policy team is setting those, and our security team is finding them. And coordinated inauthentic behavior means behavior on our site that's inauthentic, so people are not representing themselves to be who they are to be. And coordinated means they are coordinating it, and they can be coordinating with authentic actors and coordinating with inauthentic actors. Both are unacceptable.

Senator RISCH. When the team is sitting there meeting, is there generally a unanimity amongst them on something—a fact situation comes in front of them. Is this something that is easy to recognize—people are unanimous about it—or do you wind up with debates as to whether or not a certain platform should be shut down?

Ms. SANDBERG. I think on a lot of issues we face like hate speech, there's broad debate. When it comes to what is an inauthentic actor, which is a fake account posing as someone, they're hard to find. But once we find them, we know what they are.

Senator RISCH. And what about—the Chairman referred to standards in his opening statement. Who sets these standards, the same committee?

Ms. SANDBERG. The same group of people.

Senator RISCH. And are they published, so that a user can look at that? Well, give me some examples of standards that are unacceptable.

Ms. SANDBERG. In the coordinated inauthentic behavior or in general?

Senator RISCH. In general.

Ms. SANDBERG. Yes, so we publish our community standards comprehensively. And what that does is define what's permitted on Facebook and what's not permitted on Facebook. So some examples are, bullying is not permitted, hate is not permitted, language that leads to violence is not permitted, and this is published in detail publicly.

Senator RISCH. Mr. Dorsey, where's your company on these things?

Mr. DORSEY. So, we have a team called Trust and Safety who is responsible for designing and writing these policies that reports up

to our lead of legal and safety, and—and our compliance teams which report directly to me.

Senator RISCH. I'd like to ask both of you: One of the things this committee wrestles with frequently when it comes to privacy issues and those kinds of things is the difference between a U.S. citizen and a non-U.S. citizen. And under U.S. law, they can be treated differently under different circumstances.

Do your companies make any distinction between a U.S. citizen versus a non-U.S. citizen? And I guess, now I'm more focusing in on the kind of behavior we saw where elections are attempted to be manipulated and—and that sort of thing. Ms. Sandberg, let's start with you. Does your company make a distinction as they're weighing the activity of certain actors?

Ms. SANDBERG. So for political and issue ads, we are now going through a verification process. And in order to run those in the United States, people have to verify that they are legally able to do that. So that's one area where we would distinguish.

Senator RISCH. And what does that mean, legally able to do that? If a citizen of another country, any other country, decides they want to say something about a U.S. election, are they disqualified from doing that with your company?

Ms. SANDBERG. In the free content—so what their posts are to their friends and family or publicly—people are allowed to talk about any issues in any country, as long as they're not crossing over into the areas we discussed that aren't allowed, like hate and bullying. In advertising, in U.S. elections, you have to be a U.S. citizen.

Senator RISCH. Mr. Dorsey.

Mr. DORSEY. We have very similar policies and we do segment them by advertising and also the more organic social creation of content as well.

We don't always have an understanding of where an account is located. We have to infer this oftentimes. And this is where we do get a lot of help from our law enforcement partners. It is not only to understand where some of these threats are coming from, but also the intent. And the faster that we get that information, the faster that we can act.

Senator RISCH. One of the concerns that I have—and I appreciate that explanation—but what we've seen on this committee, and have actually seen in other contexts, is that in today's world it is so easy to either employ or even impersonate a U.S. citizen to do something in a given context. Do you have difficulties in that regard?

Ms. SANDBERG. Well, finding inauthentic behavior is a challenge and I think you're seeing us put real resources to bear. This is why we're investing so heavily in people and technology. This is why we're investing in programs like verification.

I think the other step we're taking here is around transparency. So being able to see if people bought political ads, where they're located, being able to see who's running a page; these are steps we think are really important for helping us find what—to your point—can be very difficult things to find.

Senator RISCH. Mr. Dorsey, briefly.

Mr. DORSEY. We've decided to focus a lot more on the behavioral patterns that we're seeing across the network. While we can't al-

ways recognize in real-time where someone might be coming from or if they were—if they are representing someone who does not exist, we can see common patterns of behavior and utilizing the network to spread their information.

So we have been building a lot of our machine learning and deep learning technology to recognize these patterns and shut them down before they spread too quickly. And then, also, link them to other accounts that demonstrate similar patterns. And we've gotten a lot more leverage out of that in terms of scalability than working on systems to identify whether it's a fake profile or not.

Senator RISCH. Interesting, thank you.

Chairman BURR. Senator Wyden.

Senator WYDEN. Thank you, Mr. Chairman.

Mr. Chairman, I want to thank you and Senator Warner for your kind comments about John McCain. And what is not often remembered is John McCain wrote some of the really important rules of the road for the internet when he was Chairman of the Commerce Committee. And it was always bipartisan, so I very much appreciate both of you mentioning our wonderful friend, John McCain.

And Ms. Sandberg, Mr. Dorsey, welcome and I've enjoyed visiting with you. Let me go right to the question that is foremost on my mind, and that is consumer privacy as a national security issue.

Technology companies like yours hold vast amounts of very private information about millions of Americans. The prospect of that data being shared with shady businesses, hackers, and foreign governments is a massive privacy and national security concern. Russians keep looking for more sophisticated ways of attacking our democracy.

Personal data reveals not just your personal and political leanings, but what you buy, even who you date. My view is personal data is now the weapon of choice for political influence campaigns. And we must not make it easier for our adversaries to seize these weapons and use them against us.

So I'd like to see if we could do a yes or no on this. And I wrote it because I think we can. My view is, from this point on, beefing up protections and controls on personal privacy must be a national security priority. I'd like a yes or no, Ms. Sandberg.

Ms. SANDBERG. Yes.

Senator WYDEN. Mr. Dorsey.

Mr. DORSEY. Yes.

Senator WYDEN. Okay. Let me turn now to a question based on a lot of analysis my office has done and you all have talked to us about. We have reviewed Facebook privacy audits required by the 2011 consent agreement after your company was found to use unfair and deceptive practices.

One section of the audits deals with how Facebook shared the personal information of Americans with smart phone manufacturers. These included the Chinese companies Huawei and ZTE. I found portions of this audit very troubling and the findings could affect many Americans. I believe, Ms. Sandberg, the American people deserve to see this information. Will you commit this morning to making public the portion of your audits that relate to Facebook's partnerships with smart phone manufacturers?

Ms. SANDBERG. Senator, I really appreciate the question and the chance to clarify this issue because it's really important. With regards to the audits, our third-party auditor, PWC, does audits on a rolling basis every two years, but they're continual. They are given to us. We have shared them with the FTC voluntarily and we will continue to do that.

I can't commit right in this moment to making that public because a lot of that has sensitive information which could help people game the system, but we will certainly work with you to see what disclosures would be prudent. But—

Senator WYDEN. Let's do this. Because that's a constructive answer and I've got other things I've got to cover. I'm just going to assume you will work with this. We understand the question of redaction on sensitive national security matters.

Can you get back to me within a week with respect to how Facebook will handle what I think is troubling information?

Ms. SANDBERG. We're going to get back to you as quickly as possible. We can definitely prioritize this request. So we'll do it as fast as we can depending on the volume of requests everyone has.

Senator WYDEN. Thank you. And look, so you all know where I'm going with this. To me, protecting data privacy has to be a higher tier issue in terms of national security. It's going to be the foundation of the legislation that I've talked to both of you about. So that's why I feel strongly and I think your answer is constructive and I hope we can get that quickly.

What I also want to get to with you, Ms. Sandberg, is the issue of micro targeting to discourage voting. This is one of the most powerful tools in the propaganda arsenal. Going after individual Americans with ads and really lasering in on the ability to affect political campaigns. It's certainly been used in the past with the Russians to discourage minority Americans from voting. Would Facebook's current policies prohibit using micro targeting to discourage voting?

Ms. SANDBERG. Senator, we feel very strongly about this. There is a long history in this country of trying to suppress civil rights and voting rights and that activity has no place on Facebook. Discriminatory advertising has no place on Facebook.

Senator WYDEN. So what are you doing to prohibit this micro-targeting? I mean what about ads that share false information about the date of the election or the location of a polling place or ads that tell people they can vote with a text message from their phone. You have said that it's unacceptable to target minorities and others, but I really need to drill down more deeply in knowing, because I think this is a primary—we can get bipartisan agreement on. What do you do to deal with micro targeting?

Ms. SANDBERG. So with everything when we're looking for abuse of our systems and things that are against our policies, we have a combination of people reviewing ads, and we have a combination of automated systems and machine learning that help us find things and take them down quickly.

Senator WYDEN. OK, I'll hold the record open for that. Could I have, say within a week, a written answer that would get into some of those specifics?

Ms. SANDBERG. We're going to get you answers to your questions as quickly and thoroughly as we can.

Senator WYDEN. Good. My last question deals with foreign governments aiding hoaxes and misinformation and I'd like to get both of you, in fact. Why don't you start with this Mr. Dorsey?

Do either of you or your companies have any indication that Iran, Russia, or their agents have supported, coordinated with, or attempted to amplify the reach of hoaxes?

Mr. Dorsey.

Mr. DORSEY. Of hoaxes?

Senator WYDEN. Yes.

Mr. DORSEY. We certainly have evidence to show that they have utilized our systems and gamed our systems to amplify information. I'm not sure in terms the definition of hoax in this case, but it is likely.

Senator WYDEN. Okay.

Ms. Sandberg.

Ms. SANDBERG. Just two weeks ago, we took down 650 pages and accounts from Iran. Some were tied to state-owned media and some of them were pretending to be free press, but they weren't free press. So it depends on how you define a hoax, but I think we're certainly seeing them use misinformation to campaign—

Senator WYDEN. My time is up. The only other area I'm going to want to explore with you is, we've got to deal with this back and forth between the private sector and the government. Very often, we ask you all about things you're doing and you say we need the government to also help us get to A, B, C, and then the government says the same thing about you. We'll want to explore that. Thank you Mr. Chairman for the extra time.

Chairman BURR. Senator Rubio.

Senator RUBIO. I want to thank you both for being here.

First of all, there's an empty chair next to you from Google. They're not here today and maybe it's because they're arrogant or maybe it's because there's a report that as of last night—this was posted at 3:36 yesterday—this group went on basically pretending to be Kremlin-linked trolls. They did everything. They used the details of the Internet Research Agency, which is a Kremlin-linked troll farm, and were able to buy ads online and place them on sites like CNN, CBS This Morning, HuffPost, The Daily Beast, so I'm sure they don't want to be here to answer these questions.

But I thank you both for being here. I was happy to read in your opening statement, Ms. Sandberg, that you talk about our democracy, our democratic process. You acknowledge responsibility for protecting our process. And you talked about our adversaries, clearly linking the company to the values and the importance of this country and I think in acknowledgment that your company would not exist were it not in the United States, because of the freedoms that we have.

Twitter didn't go as far, but you did describe yourself as a global town square—but you did say that you want to support free and open democratic debate. You did refer to our democracy and you did say that Twitter was built on the core tenet of freedom of expression, which is a very important core tenant.

Here is why this is relevant, because we're here today because we learned—and we've learned the hard way—that social media was largely seen as a tool for incredible good. Also, what makes it good can be manipulated by bad actors to do harm. And that's what happened. We have all learned that the hard way.

And so what we're asking you to do, and I think what you've agreed to do, is to use the powers that you have within your platforms to crack down on certain users who are hostile actors, who are using disinformation or misinformation or hate speech for the purposes of sowing discord, or interfering in our internal affairs—and that's a positive.

Here's the problem though: we have to start thinking about what happens when an authoritarian regime asks you to do that because their definition of disinformation or misinformation could actually be the truth. Their discord, or what they define as discord, would be things like defending human rights. Interfering in their internal affairs, they would define as advocating for democracy. And the reason why I think that answering that question is so important is because it's going to define what your companies are. Are your companies really built on these core values, or are they global companies, like all these other companies that come around here, who see their number one obligation to make money and therefore market access irrespective of the price they have to pay to do so?

So, for example, in 2016 the New York Times reported that Facebook was working on a program to restrict stories from showing up in newsfeeds based on the user's geography. The story implies—and I know that it hasn't been implemented—but it implies that that was being used in order to potentially try to get back into China, but any authoritarian government could try to use that tool.

Vietnam, by the way, where you do operate, has a new law beginning on 2019 January 1st that will require you to store user data inside the country and hand over that data, to the government, of users suspected of anti-state activity, including spreading news that may impede Hanoi or hurt the economy, for example, democracy activists.

Twitter has a policy of accommodating countries that have different ideas about the contours of freedom of expression by selectively blocking tweets and accounts. For example, one of the countries you complied with is Pakistan, who has asked you to block sites for blasphemy. The blasphemy—647 cases of blasphemy over a ten-year period from 1986 to 2007. Fifty percent of those cases were on non-Muslim Pakistanis—in a country three percent non-Muslim.

One high-profile case is Asia Bibi, who has been sentenced to death after a personal dispute over drinking water with a group of women. They accused her of insulting the prophet. She's arrested, imprisoned, sentenced to death. Not relevant to Twitter but relevant to the blasphemy laws that Pakistan has asked you to comply with.

Turkey has requested that you block over 12,000 accounts. Since 2014, you've blocked over 700. Many of them are journalists. One of them is an NBA player, Enes Kanter. Russia blocked almost 80 accounts as of last check. You complied with that. One of them was a pro-Ukrainian account in 2014.

And so here's why all of this is relevant. I guess the first question for Facebook is: These principles of our democracy—do you support them only in the United States or are these principles that you feel obligated to support around the world?

Ms. SANDBERG. We support these principles around the world. You mentioned Vietnam. We do not have servers in Vietnam. And with very minor exceptions of imminent threats that were happening, we've never turned over information to the Vietnamese government, including political information.

Senator RUBIO. And you never will?

Ms. SANDBERG. We would not.

Senator RUBIO. You would not agree to do so in order to operate?

Ms. SANDBERG. We would only operate in a country when we can do so in keeping with our values.

Senator RUBIO. And that would apply to China as well?

Ms. SANDBERG. That would apply to China as well.

Senator RUBIO. Thank you. And on Twitter, how is blocking the account of journalists or an NBA player in keeping with the core tenant of freedom of expression?

Mr. DORSEY. We enacted a policy some time ago to allow for per-country content takedown. Meaning that within the boundaries of that nation, the content would not be able to be seen but the rest of the world can see it. And that's important because the world can still have a conversation around what's happening in a market like Turkey. And also, we have evidence to show that a lot of citizens within Turkey access that content through proxies and whatnot, as well.

So, we do believe—and we have fought the government—the Turkish government—consistently around their requests and oftentimes won. Not in every case, but oftentimes have made some moves. So we would like to fight for every single person being able to speak freely and to see everything, but we have to realize that it's going to take some bridges to get there.

Senator RUBIO. Well, because a Twitter spokesman in response to a BuzzFeed article—I think about two years ago—here's the quote defending this policy. It said, "Many countries including the United States have laws that may apply to tweets and/or Twitter account content." And then you went on to say what you said, "On our continuing efforts to make services available to users everywhere et cetera." You would agree that there's no moral equivalency between what we're asking you to do here and what Turkey has asked you to do, or other countries have asked you to do, in that same realm?

Mr. DORSEY. We do have to comply with the laws that govern us within each one of these nations, but our ideals are similar and our desires—

Senator RUBIO. Whose ideals are similar? I'm sorry.

Mr. DORSEY. The company's.

Senator RUBIO. Are similar to who?

Mr. DORSEY. Similar to how we were founded and where we were founded in this country.

Senator RUBIO. I guess my point is, you're not arguing though that what we're asking you to do here—on this misinformation against foreign efforts to interfere in our elections—is the same as

what Turkey or other authoritarian regimes have asked you to do abroad, against political opponents of theirs. They're not morally equivalent, these two things?

Mr. DORSEY. Correct.

Senator RUBIO. Thank you.

Chairman BURR. The Chair will recognize Senator Heinrich for questions and then members should know that we will take a short recess, no more than five minutes, and then reconvene.

Senator Heinrich.

Senator HEINRICH. Thank you, Mr. Chair, and thank you both for being here. I think we've learned quite a bit over the course of the last couple of years. I think it would be an understatement to say that we were all caught flat-footed in 2016: social media platforms, the intelligence community, this committee, government as a whole.

Obviously, we want to learn from that and what I'd like to start with is to ask from each of you, since 2016 your platforms have been used throughout the course of a number of subsequent elections—elections in France, in Germany, and other Western allies across Europe.

What have you learned from those consequential elections after 2016 and how has that informed your current posture in terms of how you're gaining transparency into this activity? Go ahead.

Ms. SANDBERG. Senator, I think we've learned a lot and I think we're going to have to continue to learn because as we learn, our opponents learn, and we have to keep up. We're working on technology and investments in people making sure fake news is disseminated less on the platforms—transparency actions and taking down bad actors.

And we've seen everywhere, from Mexico to Brazil to other places around the world, these same techniques deployed differently and each time we see it, I think we get smarter. I think we see the new threat and I think we're able to connect the dots and prevent those threats going forward.

Senator HEINRICH. Mr. Dorsey.

Mr. DORSEY. We've also learned a lot from elections around the world, most recently the Mexican election. We have opened a new portal to cover that election, that allows any journalist or government law enforcement to actually report any suspicious behavior very quickly to us, so we can take more actions.

Otherwise, we have been investing in artificial intelligence and machine learning models to, again, recognize the patterns of behavior because we believe this is where the greatest leverage will come from, recognizing how people artificially amplify information and shutting it down before it spreads into the shared spaces of Twitter and more broadly into someone's replies to a tweet.

Senator HEINRICH. I want to get to the basic issue of whether our incentives in this case are aligned to deal with these challenges. If your users were to lose confidence in your platforms, in the authenticity of what you, Mr. Dorsey, called a public square—I might call it a digital public square—I assumed there would be very serious economic implications for your companies. Do you think the—the incentives have aligned for platform providers of all types in the

digital space, to want to get at these issues, and have a plan, and be able to respond in real time?

Ms. Sandberg and then you, Mr. Dorsey.

Ms. SANDBERG. Absolutely. Trust is the cornerstone of our business. People have to trust that what they see on Facebook is authentic. People have to trust that this is a positive force for democracy and the things they care about. And so this has been a huge issue for us and that's why we're here today and that's why we're going to keep working to get ahead of these threats and make sure we can minimize all of this activity.

Mr. DORSEY. Our incentives are aligned but I do believe it goes a lot deeper than just the alignment of our company incentives with this committee and the American people. I believe we need to question the fundamental incentives that are in our product today.

Every time someone opens up our service, every time someone opens up our app, we are implicitly incentivizing them to do something or not to do something. And that extends all the way to our business and those answers that we get from asking that question are going to create massive shifts in how Twitter operates and I also believe how our industry operates. So what worked 12 years ago does not work today—it hasn't evolved fast enough—but I think it is a layer—many, many, many, many layers deeper than the surface symptoms that we often find ourselves discussing.

Senator HEINRICH. Ms. Sandberg, you mentioned a number of things that would violate your standards, for example, hate speech, advocacy of violence. What about when you were dealing with real people, authentic users, intentionally spreading false information? And obviously there are huge free speech implications there. But, for example, what if a real person, a U.S. citizen, says that victims of the mass shootings were actually actors? Would that violate your standards and if the answer is no, how should we, and by we, I mean government and industry, deal with those very real challenges?

Ms. SANDBERG. Well let me start by saying I find claims like that personally, unbelievably upsetting. If you've been a victim or a parent of a victim, they deserve all our full support. And finding a line between what is hate speech and what is misinformation is very, very difficult, especially if you're dedicated to expressing free expression, and sometimes free expression is expressing things you strongly disagree with.

In the case of misinformation, what we do is we refer it to third-party fact-checkers. We don't think we should be the arbiter of what's true and what's false, and we think that's really important. Third-party fact-checkers then mark it as false. If it's marked as false, we dramatically decrease the distribution on our site. We warn you if you're about to share it. We warn you if you have shared it and, importantly, we show related articles next to that so people can see alternative facts.

The fundamental view is that bad speech can often be countered by good speech, and if someone says something is not true and they say it incorrectly, someone else has the opportunity to say, actually you're wrong. This is true and that's what we're working on through our systems.

Senator HEINRICH. I think one of the things we found in 2016 is that we didn't have the transparency and the literacy to do what you just pointed out there: to counter false speech with accurate speech to understand how this speech was propagating in the digital public space.

What more do you think we should be doing to simply make the public more literate about the fact that this information warfare is very real? It's going on all the time. It's not fake news. It's not a hoax. It's something we're all going to have to deal with, that our kids, even playing platforms like Pokemon Go, may have to—have to deal with as well.

Do either of you have a quick opinion on that? And then my time will be expired. I apologize, Mr. Chair.

Mr. DORSEY. I believe we need to point to where we see healthy participation and clearly mark what is healthy and what is unhealthy. And also realize that not everyone is going to choose healthy participation in the short term. But how do we encourage healthy participation in order to increase the reach and also increase the value of what they're giving to that digital public square.

Chairman BURR. This hearing stands in a recess subject to the call of the Chair.

[Whereupon the hearing recessed at 10:51 a.m. and reconvened at 11:01 a.m.]

Chairman BURR. I'd like to call the hearing back to order. The chair would recognize Senator Collins for questions.

Senator COLLINS. Thank you, Mr. Chairman. First let me thank you both for being here and also to express my outrage that your counterpart at Google is not at the table as well.

Mr. Dorsey, as of January of this year, Twitter has taken down more than 3,800 Russian IRA accounts that by Twitter's own estimate reached approximately 1.4 million people. One of those accounts purported to be under the control of the Tennessee GOP, although it was not. It was a Russian IRA account. It had more than 140,000 followers and would sometimes spread conspiracy theories and false claims of voter fraud.

My question to you is: Once you have taken down accounts that are linked to Russia, these impostor accounts, what do you do to notify the followers of those accounts that they have been following or engaged in accounts that originated in Russia, and are not what they appear to be?

Mr. DORSEY. Thank you for the question. We simply haven't done enough. So in this particular case, we didn't have enough communication going out in terms of what was seen and what was tweeted, and what people are falling into.

We do believe transparency is a big part of where we need the most work and improvement, and it's not just with our external communications, it's actually within the product and the service itself.

We need to meet people where they are, and if we determine that people are subject to any falsehoods or any manipulation of any sort, we do need to provide them the full context of that. And this is an area of improvement for us and something that we're going to be diligent to fix.

Senator COLLINS. I think this is critically important. If a follower just gets a message that says this Twitter account is no longer available, that does not alert the individual that he or she has been receiving messages—tweets—from a Russian entity whose goal is to undermine public confidence in elected officials and our democratic institutions.

So I really think we need something more than even the tombstone, or something else. We need to tell people that they were taken in or victims—innocent victims—of a foreign influence campaign.

Ms. Sandberg, let me ask you this same question. What is Facebook doing?

Ms. SANDBERG. We agree with you that people need to know, so we've been discussing these publicly, as well as in specific cases notifying people. So we notified people directly if they had liked—or had liked the original IRA accounts.

Most recently when there was an event that was going to be happening in Washington that inauthentic accounts—we notified all the people who either RSVP'd to that event, or who said they were interested in possibly going to that event.

Senator COLLINS. Thank you. That was the Night to Defeat the Right, or something like that, as I recall.

Mr. Dorsey, back to you. Clemson University researchers and others have shown that these Russian IRA accounts target specific leaders and social movements across the political spectrum. And again, the goal of the Russians, the Iranians—anyone else who is involved in this influence campaign—is to undermine the public's confidence in political leaders and weaken our democratic institutions and turn us against one another.

Well, I learned not from Twitter but from Clemson University that I was one of those targeted leaders and that there were 279 Russian-generated tweets that targeted me that had gone to as many as 363,000 followers. So why doesn't Twitter notify individuals like me that we have been targeted by foreign adversaries? I shouldn't find out from looking at Clemson University's database and working with their researchers. It seems to me that once you determine that, you should notify the people who are the targets.

Mr. DORSEY. I agree. It's unacceptable. And as I said earlier, we want to find ways to work more openly, not just with our peer companies but with researchers and universities and also law enforcement because they all bring a different perspective to our work, and can see our work in a very different light. And we are going to do—we're going to do our best to make sure that we catch everything and we inform people when it affects them. But, we are not going to catch everything. So it is useful to have an external partnership and work with them to make sure that we're delivering a message in a uniform manner where people actually are, without requiring them to find a new channel to get that information.

This is where a lot of our thinking is going and a lot of our work is going. But we recognize we need to communicate more directly where people are on our service, and we also recognize that we're not going to be able to catch everything alone, so we need to develop better partnerships in order to do that.

Senator COLLINS. I would close my questioning by encouraging both of you to work more closely with academia, with our government. The Clemson University researchers have done extraordinary work, but they have said that they've been provided data that is only within the last three years, which does not allow them to do the kind of analysis that they'd like to do and that's probably because of the new European Union privacy laws. But the EU has provided research exemptions. So I hope that you will commit to providing data that goes beyond that three year window to researchers who are looking into Russian influence efforts on your platforms. Thank you.

Chairman BURR. Senator Harris.

Senator HARRIS. Thank you, Mr. Chairman, for accommodating me. I'm in another hearing as you know. Good morning, and to the invisible witness, good morning to you. So I have a few questions for Ms. Sandberg. On November 2, 2017, your company's general counsel testified in front of this Intelligence Committee on Russian interference, and I asked a few questions.

I asked how much money did you make, and this is of the representative from both Facebook and Twitter—both of your general counsels were here. And I asked how much money did you make from legitimate advertising that ran alongside the Russian propaganda. The Twitter general counsel said, quote, "We haven't done the analysis but we'll follow-up with you and work on that." And the Facebook general counsel said the same is true for Facebook.

Again, I asked Facebook CEO Mark Zuckerberg on April 10, 2018, and he said that, quote, "Internet Research Agency, the Russian firm, ran about \$100,000 worth of ads." Following the hearing, I asked Facebook the same question in writing, and on June 8, 2018, we received a response that said, quote, "We believe the annual revenue that is attributable to inauthentic or false accounts is immaterial."

So my question is: What did you mean by immaterial? Because I'm a bit confused about the use of that term in this context.

Ms. SANDBERG. Thank you for the question.

Again we believe the total of the ad spending that we have found is about \$100,000. And so the question you're asking is with the inorganic content, I believe, what is the possible revenue we could have made? So here's the best way I can think of to estimate that, which is that we believe between 2015 and 2017, up to 150 million people may have seen the IRA ads or organic content in our service. And the way our service works is, ads don't run attached to any specific piece of content, but they're scattered throughout the content. This is equivalent to .004 percent of content in news feed and that was why they would say it was immaterial to our earnings.

But I really want to say that from our point of view, Senator Harris, any amount is too much.

Senator HARRIS. If I may, just so I'm clear about your response—so are you saying that then the revenue generated was .004 percent of your annual revenue? Of course that would not be immaterial.

Ms. SANDBERG. Again, the ads are not attached to any piece of content so—

Senator HARRIS. So what metric then? Just help me with that. What metric are you using to calculate the revenue that was generated, associated with those ads? And what is the dollar amount that is associated then with that metric?

Ms. SANDBERG. The reason we can't answer the question to your satisfaction is that ads are not—organic content—ads don't run with inorganic content on our service, so there is actually no way to firmly ascertain how much ads are attached to how much organic content. It's not how it works.

In trying to answer what percentage of the organic—

Senator HARRIS. But what percentage of the content on Facebook is inorganic?

Ms. SANDBERG. I don't have that specific answer, but we can come back to you with that.

Senator HARRIS. Would you say it's the majority?

Ms. SANDBERG. No. No.

Senator HARRIS. An insignificant amount? What percentage? You must know.

Ms. SANDBERG. If you ask about our inauthentic accounts on Facebook, we believe at any point in time it's 3 percent to 4 percent of accounts, but that's not the same answer as inorganic content because some accounts generate more content than others.

Senator HARRIS. I agree. So what percentage of your content is inorganic?

Ms. SANDBERG. Again, we don't know. I can follow up with the answer to that.

Senator HARRIS. Okay, please. That would be great. And then your company's business model is obviously—it's complex but benefits from increased user engagement and that results of course in increased revenue. So, simply put, the more people that use your platform, the more they are exposed to third-party ads, the more revenue you generate. Would you agree with that?

Ms. SANDBERG. Can you repeat? I just want to make sure I got it exactly right.

Senator HARRIS. So the more user engagement will result—and the more then that they are exposed to third-party ads—the more that will increase your revenue. So the more users that are on your platform—

Ms. SANDBERG. Yes. Yes. But only I think when they see really authentic content. Because I think in the short run and over the long run it doesn't benefit us to have anything inauthentic on our platform.

Senator HARRIS. That makes sense. In fact, the first quarter of 2018, the number of daily active users on Facebook rose 13 percent, I'm told. And corresponding ad revenue grew by half to \$11.79 billion. Does that sound correct to you?

Ms. SANDBERG. Sounds correct.

Senator HARRIS. And then would you agree that—I think it's an obvious point—that the more people that engage on the platform, the more potential there is for revenue generation for Facebook?

Ms. SANDBERG. Yes, Senator. But again, only when the content is authentic.

Senator HARRIS. I appreciate that point. And so a concern that many have is how you can reconcile an incentive to create and in-

crease your user engagement when the content that generates a lot of engagement is often inflammatory and hateful.

So, for example, Lisa-Maria Neudert, a researcher at Oxford and Internet Institute, says, quote, “The content that is the most misleading or conspiratorial, that’s what’s generating the most discussion and the most engagement, and that’s what the algorithm is designed to respond to.”

My concern is that according to Facebook’s community standards, you do not allow hate speech on Facebook. However, contrary to what we’ve seen, on June 28, 2017, a ProPublica report found that Facebook’s training materials instructed reviewers to delete hate speech targeting white men but not against black children because black children are not a protected class. Do you know anything about that, and can you talk to me about that?

Ms. SANDBERG. I do. And what that was, I think, a bad policy that’s been changed, but it wasn’t saying that black children—it was saying that children—it was saying that different groups weren’t looked at the same way, and we’ve fixed it.

Senator HARRIS. But isn’t that the concern with hate, period? That not everyone is looked at the same way?

Ms. SANDBERG. Well, hate speech is against our policies and we take strong measures to take it down. We also publish publicly what our hate speech standards are. We care tremendously about civil rights. We have worked very closely with civil rights groups to find hate speech on our platform and take it down.

Senator HARRIS. So when did you address that policy? I’m glad to hear you have. When was that addressed?

Ms. SANDBERG. When it came out—and again, that policy was a badly written, bad example, and not a real policy.

Senator HARRIS. The report that I’m aware of was from June of 2017. Was the policy changed after that report or before that report from ProPublica?

Ms. SANDBERG. I can get back to you on the specifics of when that would have happened.

Senator HARRIS. You’re not aware of when it happened?

Ms. SANDBERG. I don’t remember the exact date.

Senator HARRIS. Do you remember the year?

Ms. SANDBERG. Well, you just said it was 2017.

Senator HARRIS. So do you believe it was 2017 that the policy changed?

Ms. SANDBERG. It sounds like it was.

Senator HARRIS. Okay. And what is Facebook’s official stance on then hate speech regarding so-called, and legally defined, unprotected classes, such as children?

Ms. SANDBERG. Hate speech is not allowed on our platform and hate speech is, you know, important in every way. And we care a lot that our platform is a safe community. When people come to Facebook to share, they’re coming because they want to connect on the issues that matter to them.

Senator HARRIS. So, have you removed the requirement that you will only protect with your hate speech policy those classes of people that have been designated as protected classes in a legal context? Is that no longer the policy of Facebook?

Ms. SANDBERG. I know that our hate speech policies go beyond the legal classifications and they are all public and we can get back to you on any of that. It's all publicly available.

Senator HARRIS. Thank you so much. Thank you, Mr. Chairman. Chairman BURR. Senator BLUNT.

Senator BLUNT. Thank you, Chairman. Mr. Dorsey, Wired magazine last week had an article that said you'd admitted having to rethink fundamental aspects of Twitter. Would that be an accurate reflection of where you've been the last year?

Mr. DORSEY. Yes. We are rethinking the incentives that our service is giving to people.

Senator BLUNT. And what would be the biggest area where you're trying to rethink how you thought this was going to work out and the way it's turned out to be?

Mr. DORSEY. Well—and this is pretty far-reaching—so we're still in the process of doing this work, but when we created the service 12 years ago, we had this concept of followers. And we made the number of followers big and bold and a very simple but noticeable font.

And just that decision alone has incentivized people to want to grow that number, to increase that number. And the question we're now asking is, "Is that necessarily the right incentive? Is the number of followers you have really a proxy for how much you contribute to Twitter and to this digital public square?" And we don't believe it is. But that's just one question. The way we lay out our buttons on the bottom of every tweet in a reply and a retweet and a like, that also implies an incentive and a point of view that we're taking that we want to encourage people to do.

So as we think about serving the public conversation, as we think about our singular priority of increasing the health of that public conversation, we are not going to be able to do long-term work unless we are looking at the incentives that our product is telling people to do every single day.

Senator BLUNT. All right, that's helpful. Thank you. Senator Collins asked her last question—I didn't really quite get the answer to that question. But I think what she was asking is a question I had also, which was: In the interest of transparency and public education and looking at things available to researchers and policy makers, are you willing to archive suspended accounts so that people can look back at those? And would that be a period of, I think, three years was part of the question she asked. Give me a little better, more specific answer. You didn't have time to answer that, and I'd like you to have time to answer that.

Mr. DORSEY. We are looking at things like a transparency report. We put out a transparency report around terrorism, but we're looking at expanding that transparency report around suspensions of any account.

We are still coming up with the details of what this will look like and what it will include.

Senator BLUNT. As opposed to just a transparency report, are you willing to archive some of this where you may not be reporting on it at the time, but someone could look three years down the road and try to do an analysis of why that information was out there

the way it was and how it fit into your overall policy of taking whatever action you're taking?

Mr. DORSEY. I think it's a great idea to show the historical public record. We just need to understand what the legal implications are, and we can get back to you on that.

Senator BLUNT. Yes, I may come back with a question if I have time on legal implications, generally. I think for both of your companies, who have been pretty forward-leaning in the last couple of months as this conversation has moved pretty dramatically, the business implications, the liability implications of what we're asking you to do are pretty great.

Well, let me see if I can get a couple of Facebook questions in first. Ms. Sandberg, does Facebook differentiate between foreign and domestic influence operations when deciding whether to take down a page or remove an account from the platform?

Ms. SANDBERG. Our focus is on inauthenticity, so if something is inauthentic, whether it's trying to influence domestically or trying to influence on a foreign basis—and actually a lot more of the activity is domestic—we take it down.

Senator BLUNT. You take it down indiscriminately, whether it's a foreign influence or—or a domestic influence?

Ms. SANDBERG. And you saw that with the IRA. With the IRA accounts, the original ones for our election were targeted at the United States, but then there were another 270 accounts that were almost all targeted in Russia or at Russia—for Russian speakers and nearby languages. So a lot of those were domestic, and those are down.

Senator BLUNT. Well, it's been mentioned several times, and I think appropriately so, Google is not here today. But the two of you are, and Ms. Sandberg, again, just what seems like a long time ago, but only a few months, since Mr. Zuckerberg was here testifying before Congress. It seems like to me that Facebook has been pretty active in finding and taking down things that should not have been out there: the recent Iranian takedown, the Russian things that have been taken down.

Do you want to talk a little about what's the big challenge about being at the forefront of trying to figure this out from a business perspective or a liability perspective, either one? Then I'm going to come to Mr. Dorsey with the same question.

Ms. SANDBERG. Well I really appreciate what you said, because we have been investing very heavily in people, in our systems, in decreasing the dissemination of fake news, in transparency, and I think that's what you're seeing pay off.

I think we've all said, in the private meetings we had as well as this public discussion, that tighter coordination really helps us. If you look at our recent takedowns, some of it was information we found ourselves, some of it were hints we got from law enforcement, some of it is information we can share with other companies.

And so this is a big threat, and our opponents are going to keep getting better and we have to get better. We have to stay ahead. And the more we can all work together the better off we're going to be, and that's why I really appreciate the spirit with which this hearing this morning is taking place.

Senator BLUNT. And how does the takedown, the practice work, where legitimate accounts are sold then maybe—and repurposed by others? What are you looking at there as a challenge?

Ms. SANDBERG. So our policy is inauthenticity. If you are an inauthentic account, if you are pretending to be someone you're not, you come down. If you have touched the account of someone who is authentic, then we would leave the authentic account up, but in cases like I was answering with Senator Collins, if you are an authentic person who RSVP'd to an event that's not authentic, we would let you know.

Senator BLUNT. Okay, thank you for that. Okay, Mr. Dorsey, back to that other question. From a business and legal liability standpoint, what's the downside of being out there where you are now trying to every day implement policies that nobody's ever implemented before?

Mr. DORSEY. I think there are a number of short-term risks, but you know, we believe that the only way that we will grow and thrive as a company is by increasing the health of this digital public square that we're helping to build. We also benefit, as Sheryl mentioned, from tighter collaboration and tighter partnership. We've really strengthened our partnership with our government agencies since 2016.

There are a few areas that we would like to see more strength. We would like a more regular cadence of meetings with our law enforcement partnerships. We would love to understand the secular trends that they are aware of, and seeing in our pure companies or other mediums, or more broadly that would inform us about how to act much faster. And we would appreciate as much as we can consolidating to a single point of contact, so that we are not bouncing between multiple agencies to do our work.

So that is what we've found in attempting to do a lot of this new policy and work, in terms of partnership, but ultimately it comes back to: we need to build our technologies to recognize new patterns of behavior and new patterns of attack, and to understand what they actually mean, and then ideally get some help from our law enforcement partners to understand the intent and to understand the motivations behind it.

Senator BLUNT. Thank you, Mr. Dorsey. I'm sure my time is up. Thank you, Chairman.

Chairman BURR. Senator King.

Senator KING. Thank you, Mr. Chairman, and I want to also thank our witnesses. And thank you to your companies and your policy makers for making really great strides in the last year. As many of the people have talked about, we were all on our heels a year ago on this subject. And this has emerged as one of the most important parts of this committee's investigation.

I try to focus on what we're after here. And we're after the heart of democracy. Ms. Sandberg, you said the heart of democracy was free and fair elections. I would argue that the heart of free and fair elections is information. And that's really what we're talking about: getting information to people in a democratic setting. And also on all kinds of other topics, birthdays and everything else, but that's what we're talking about here.

There are three ways to defend ourselves it seems to me. One is better consumer discrimination about what they're seeing. The second is deterrence, which hasn't been mentioned here, that our adversaries need to understand that there's a price to be paid for trying to manipulate our society and our democracy. And the third is technical, and that's mostly what we've been talking about.

I had an experience, ironically, a couple of months before the 2016 election, meeting here in this building with a group of people from Lithuania, Estonia, and Latvia, who have been experiencing Russian interference with their elections and their propaganda, their information for years. And I said, "How do you defend yourself?" You can't unplug the internet. You can't turn off the TV station. The most interesting thing they said was, universally, the best defense is for the people to know it's happening.

And I would like from each of you some thoughts and hopefully a commitment to educating your users about the potential for abuse of the very medium that they're putting their trust in.

Ms. Sandberg.

Ms. SANDBERG. We really agree with you. And we've done this broadly and we're going to continue to do more. So we've worked on media literacy programs. We've worked on programs in public service announcements around the world that help people discern—this is real news, this is not—and help people be educated. I think one of the most important things we're doing is that once a piece of content has been rated as false by our third-party fact-checkers—if you're about to share it, we warn you right there. Hey this has been rated as false. And so, you are educated as you are about to take that critical step.

Senator KING. And Mr. Dorsey, I hope you're doing the same to educate your users as to the potential that they can be misled on your platform.

Mr. DORSEY. Yes. And to be frank, we haven't done a good job at this in the past. And I think the reason why is because we haven't met our customers where they are, in terms of actually when they're using the product and adding more context there.

We do benefit on Twitter that we have this amazing constituency of journalists globally using our service every single day, and they often, with a high degree of velocity, call out nonfactual information. We don't do a great job at giving them the best tools and context to do that work. And we think there's a lot of improvements we can make to amplify their content and their messaging so that people can see what is happening with that content.

Senator KING. If that can be amplified and underlined, it can become a self-healing process, whereby the response immediately responds to false or misleading information.

Deterrence, I'm not going to spend a lot of time on, except to say that many of us believe that one of the great gaps in our defenses against election interference and interference in our democracy is the fact that our adversaries feel no pain if they do so—that we have to develop a doctrine of cyber deterrence just as we have doctrines of military deterrence. And that's a gap, and that's something that we're working on both here and at Armed Services, other places.

Let me talk about the technical for a minute. How about feedback from users? And Ms. Sandberg, you testify that you have third-party fact-checkers. Also, would it be useful to have more in the way of ratings? And, you know, the eBay sellers—you have rating process and number of stars, and those kinds of things. Is there more you could do there to alert people as to the validity and the trustworthiness of what they're seeing?

Ms. SANDBERG. Senator, the most important determinant of what anyone sees on Facebook are decisions they make. So I choose my friends, you choose yours. I choose the news publications I follow, you choose yours. And that's why your news feed is so different from mine. And so, yes, if you don't want to follow someone, if you don't want to like a page, we encourage you to do that. We also make it very easy to unfollow on our site. So if I don't believe what you're saying anymore, I don't have to receive your—

Senator KING. But I'm talking about alerting a viewer or a reader to something that's come across on their newsfeed that has been found manifestly false or misleading: a banner, a note, a star.

Ms. SANDBERG. We do that through related articles. We note this has been rated as false, and here's a related article which would give you other facts that you could consider.

Senator KING. One of the things that we've been talking about here, and Senator Rubio has been a leader in discussing this, is what we call Deepfake, as I'm sure you're aware, the ability to manipulate video to the point where it basically conveys a reality that isn't real.

Is there a technological way that you can determine that a video has been manipulated in that way and tag it? So that people on Facebook, if they see a video that it'll be tagged: warning, this has been manipulated in a way that may be misleading. That's a question you may want to take under advisement. But it seems to me, again, this is an area—this is a new area that's going to get more and more serious, I'm afraid. And again, what I'm trying to do is give the consumer the maximum amount of information.

Ms. SANDBERG. We agree with you, Deepfakes is a new area and we know people are going to continue to find new ones. And as always, we're going to do a combination of investing in technology and investing in people so that people can see authentic information on our service.

Senator KING. As you're thinking about these cures, I hope you'll continuously come back to the idea that what we need to do is give people more information. I must say, I'm a little uncomfortable with where the line is between taking down misleading or fake information and taking down what someone else may consider legitimate information in the marketplace of ideas. Jefferson said we can tolerate error, as long as truth is left free to combat it. We have to be sure that we're not censoring. But at the same time, we're providing our customers, our users—your users with information that they can—the context, I think, is the word you use—they can have context for what it is that they're seeing.

I'd hate to see your platforms become political in the sense that you're censoring one side or the other of any given debate.

Mr. Dorsey.

Mr. DORSEY. So yes, we absolutely agree. As we are building a digital public square, we do believe expectations follow that. And that is a default to freedom of expression and opinion. And we need to understand when that default interferes with other fundamental human rights such as physical security or privacy. And what the adverse impact on those fundamental human rights are.

And I do believe that context does matter in this case. We had a case of voter suppression around 2016 that was tweeted out. And we are happy to say that organically, the number of impressions that were calling it out as fake were eight times that of the reach of the original tweet. That's not to say that we can rely on that, but asking the question how we make that more possible, and how we do it at velocity is the right one to ask.

Senator KING. That's the self-healing aspect. Thank you both very much. And if you have further thoughts as you're flying home, about technical ways you can increase the information available to your users through tags, ratings, stars, whatever, please share them with us and we'll look forward to working with you on this problem that is one that's important to our country. Thank you very much.

Chairman BURR. Senator Lankford.

Senator LANKFORD. Thank you, Mr. Chairman. I want to follow up on a statement that Senator King was mentioning as well about Deepfakes. That's something I've spoken to both of you about before in the past. It is a challenge for us and I would just reiterate some of the things that he was saying publicly. When it's the possibility and now the opportunity to be able to create video that looks strikingly real, but none of it is actually real—all of it is computer-generated—that is a very different day for video-sharing in the days ahead. And I know as you all have attacked issues like child pornography and other things on your platforms in the past, you all will aggressively go after these things. We're just telling you we're counting on it because Americans typically can trust what they see, and suddenly in video they can no longer trust what they see because the opportunity to be able to create video that's entirely different than anything in reality has now actually come. And so I appreciate your engagement on that.

And I want to talk to you a little bit, Mr. Dorsey, about following up some of the things that Senator Blunt had mentioned as well about suspended accounts. When you suspend an account, obviously there's information that's still there. Do you archive all of that information to be able to maintain for a suspended account that this is an account that we determine is either from a foreign actor or hostile actor or is inappropriate—not an authorized user? Is that something you hold that information, so you can maintain it?

Mr. DORSEY. I need to follow up with you on the exact details of our policies, but I believe we do, especially in regards to any law enforcement action.

Senator LANKFORD. Terrific. For Facebook, what is the practice when you suspend an account and say this is not an authorized user or we think this is a foreign or hostile user?

Ms. SANDBERG. If we have any suspicion that it's a foreign or hostile user, we would keep the information to be able to do further investigation.

Senator LANKFORD. So then the question is, is the investigation internal for you all? Or obviously if law enforcement subpoenas that and comes to you and says I have a subpoena to come get that information, that's a whole different issue. But is that something you do in your own investigation? Because as I'm sure you've seen in the past, some users will create a fake account or some sort of hostile account. That comes down, they'll create another one, and then there's some similarities in where they go and directions and relationships.

Do you maintain that data to be able to make sure that you're well prepared and educated for when they may come back to be aware of that again? For Twitter what is that, Mr. Dorsey?

Mr. DORSEY. So we do, do our own internal investigations and we are benefited every time our peers recognize something, and we do share that data so that we can check our own systems for similar vectors or similar accounts. And also work with law enforcement to understand the intent. If there is a request to allow an account to lay dormant by law enforcement, we will allow that to happen and work with them to make sure that we are tracking it accordingly.

Senator LANKFORD. Mr. Dorsey, the main thing I'm trying to identify though is, let's say it happened in 2017. You identify an account that you suspended and said this is your problem area or an unauthorized user, whatever it may be.

You take that account off, do you maintain that information? And so a year later if somebody comes back on with a similar profile you can still track it and say, this is the same as what we've seen before and it's going to take additional steps for you to get back on board or ways to be able to track their initial connections?

Mr. DORSEY. I'm sorry, yes. We do maintain that information and we have a ban evasion policy. So if someone is trying to evade a ban or suspension, no matter what the timeframe, we can take action on those accounts as well.

Senator LANKFORD. Okay.

Ms. Sandberg.

Ms. SANDBERG. If we have any suspicion that this would be engaged in foreign or domestic inauthentic activity or we have law enforcement interaction on it, we would keep that information.

Senator LANKFORD. Okay. Mr. Dorsey, you and I have spoken on this as well about data and the business model for both of you is obviously—it's a free platform for everyone to use—but obviously data and advertising and all those things are very helpful just in keeping your business open and keeping your employees paid. That's a given, and everyone understands that when they join that platform and that conversation. But for data in particular, how do you make sure that anyone who purchases into data or gets access to that uses it for its stated purpose, rather than using it to either sell to a third party or to open up as a shell company, and say they're using it for one purpose but they're actually using it for a foreign purpose or direction to be able to track real-time activity of Americans? How do you assure that companies that are purchasing

into that opportunity to have that data are actually fulfilling and using it as they stated they would?

Mr. DORSEY. Well, there's a few things here. First and foremost, we're a little bit different than our peers and that all of our data is public by default. So when we sell data, what we're selling is speed and comprehensiveness. So you're actually purchasing either insights or a real-time streaming product. In order to purchase that you have to go through a very strict know-your-customer policy that we enact and then we audit every single year. If we have any indication that there is suspicious activity happening, that is an opportunity for us to reach out to law enforcement with the sole purpose of trying to understand the intent. That is the thing that we are not always going to be able to infer from us looking at the relationship.

You mentioned setting up companies that potentially are in front of governments. That is not information that we would necessarily have and that is where we are dependent upon the intelligence to inform us so that we can take stronger action.

Senator LANKFORD. So, how do you determine or what relationship—is it an initial relationship but there's not a follow up after that rapid access as you dictate on that? After that is determined, is there any way to check in on those companies to be able to make sure they're actually fulfilling their terms of service?

Mr. DORSEY. Absolutely. And we do it every year on a regular basis. But if we see anything suspicious at any point in time, we'll reach out directly.

Senator LANKFORD. Ms. Sandberg, tell me a little bit about WhatsApp? WhatsApp has been a feature of Facebook for a while. How is the encryption going on that? What's the relationship now with WhatsApp and what do you anticipate in the days ahead?

Ms. SANDBERG. We are strong believers in encryption. Encryption helps keep people safe. It secures our banking system, it secures the security of private messages, and consumers rely on it and depend on it. And so we're very committed to encryption in WhatsApp and continuing to protect the data and information of our users.

Senator LANKFORD. So that encryption is end-to-end at this point still on the WhatsApp platform?

Ms. SANDBERG. We'll get back to you on any technical details, but to my knowledge, it is.

Senator LANKFORD. Thank you. I yield back.

Chairman BURR. Senator Manchin.

Senator MANCHIN. Thank you, Mr. Chairman. And Ms. Sandberg and Mr. Dorsey, I want to thank both of you for being here. And I grew up in an age without computers and social media so I'm trying to get acclimated the best I can. I have seen how they've been used by my children and grandchildren and how much it helps connect people. I see an awful lot of good.

I also have concerns with internet and social media have been—how it's been used against us. And I think you're hearing concerns from all of my fellow colleagues up here. It's an attempt to divide Americans, change our way of life, change our democracy as we know it, and it can be very devastating.

In my little State of West Virginia—my beautiful little State of West Virginia, with all the wonderful people—has been hit ex-

tremely hard by illicit drugs and pharmaceutical opiates. According to the recent *Wired* article, Eileen Carey spent three years regularly reporting accounts illegally selling opiates on Instagram. And the practice was widespread on Facebook and Twitter, as well.

In many ways, the tools used by opiate dealers are similar to those adopted by other bad actors including Russia, target the vulnerable with ads that are easily circumventing the platforms, filters, and oversights, and using hashtags to gain attention of those interested. Last November, Facebook CEO Mark Zuckerberg said learning of the depths of the crisis was the biggest surprise and really saddening to see. But it still took months to take measures to correct the problem while other people were still dying.

According to the U.S. Code 230, formerly known as a Communications Decency Act of 1996, online service providers shall not be held civically liable for content that a third party posts on their platform, and they shall not be treated as a publisher or speaker of the content.

If we look at the example of drug overdose deaths, many prosecutors are increasingly treating the deaths as a homicide and looking to hold someone criminally accountable. There are now laws devised to hold drug dealers responsible for the death of victims using drugs they provided and, in some cases, they are charging friends, partners, siblings of the deceased.

So my question to both of you would be: I've heard of a report that details the way drug dealers continue to use your platforms for illegal drug sales. To what extent do you bear responsibility for the death of a drug user if they overdosed on drugs received through your platform?

Either one. I know it's a tough one.

Ms. SANDBERG. I'm happy to go.

Senator MANCHIN. Yes.

Ms. SANDBERG. This is really important to us. The opioid crisis has been devastating, and takes the lives of people in our country and around the world. It's firmly against our policies to buy or sell any pharmaceuticals on Facebook, and that includes the opioid drugs. We rely on a combination of machines and people reporting to take things down, and I think we've seen marked improvements.

We also took an additional step recently which is very important which is, we're requiring treatment centers who want to buy ads to be certified by a respected third party because another one of the problems has been that some treatment centers are actually doing harm, and so we're requiring certification before they can purchase ads and they can try to reach people for treatment.

Mr. DORSEY. This is also prohibited on our service and we do have a responsibility to fix it anytime we see it. And we are looking deeply at how this information spreads, and how the activity spreads so that we can shut it down before it spreads too far.

Senator MANCHIN. I know I asked a tough question. It was, do you all feel any responsibility because there has been a lot of people that have been affected, and a lot of people have died receiving information on how to obtain drugs through your all's platform?

So I would go another step further, just like we passed FOSTA and SESTA—FOSTA was the Fight Online Sex Trafficking Act, and stop enabling—and SESTA was the Stop Enabling Sex Traf-

fickers Act. We passed bills that held you liable and responsible. Don't you think we should do the same with opiate drugs and the way they're being used in your platform? Would you all support us doing that?

Mr. DORSEY. We're certainly open to dialogue around CDA and the evolutions of it. We benefit from a lot of the protections it gives in order for us in the first place to take actions on the content within our service. The only reason we're able to even speculate that we can increase more health in a public square is because of CDA 230. So we need to finely balance what those changes are and what that means.

Senator MANCHIN. Well, did it change your all's approach of how you use your platforms with the changing of Code 230?

Mr. DORSEY. We have to do that independent of changes to 230.

Ms. SANDBERG. These things are against our policies, and we want them off and we want to take all measures to get them off. The Safe Harbor of 230 has been very important in enabling companies like ours to do proactive enforcement, look for things proactively, without increasing our liability. And so, we'd want to work very closely on how this would be enacted.

Senator MANCHIN. Final question to both of you. Why are you not doing business in China?

Mr. DORSEY. We are blocked in China.

Ms. SANDBERG. We are as well.

Senator MANCHIN. You're blocked? For what reasons?

Ms. SANDBERG. The Chinese government has chosen not to allow our service in China. I think it happened on the same day.

Senator MANCHIN. Did you all not accept, basically, the terms of how you do business in China? Or you're just blocked from coming in to it? Or did you not agree? Did they give you a chance, or—? I'm saying other social platforms seem to be adapting and going in there.

I know a lot of our drugs—a lot of the fentanyl and all that—is coming from China, and we're trying to shut that down. But it was interesting to me that you all both have been blocked. And I would assume you didn't agree to their terms?

Mr. DORSEY. I don't know if there's any one particular decision point around understanding what the terms might be in our particular case. But when we were blocked, we decided that it wasn't a fight worth fighting right now, and we have other priorities.

Senator MANCHIN. Are you still looking to do business there?

Ms. SANDBERG. There was no particular time. You know, we've been open about the fact that our mission is to connect the world. And that means, it's hard to do that without connecting the world's largest population. But in order to go into China, we would have to be able to do so in keeping with our values. And that's not possible right now.

Senator MANCHIN. Thank you.

Thank you, Mr. Chairman.

Chairman BURR. Senator Cotton.

Senator COTTON. I want to commend both of you for your appearance here today, for what was no doubt going to be some uncomfortable questions. And I want to commend your companies for making you available. I wish I could say the same about Google.

I think both of you, and your companies, should wear it as a badge of honor that the Chinese Communist Party has blocked you from operating in their country. Perhaps Google didn't send a senior executive today because they've recently taken actions such as terminating cooperation that they had with the American military on programs like artificial intelligence that are designed not just to protect our troops and help them fight and win our country's wars, but to protect civilians as well. This is at the very same time that they continue to cooperate with the Chinese Communist Party on matters like artificial intelligence, or partner with Huawei and other Chinese telecom companies that are effectively arms of the Chinese Communist Party. And credible reports suggest that they are working to develop a new search engine that would satisfy the Chinese Communist Party's censorship standards after having disclaimed any intent to do so eight years ago.

Perhaps they didn't send a witness to answer these questions because there is no answer to those questions, and the silence we would hear right now from the Google chair would be reminiscent of a silence that that witness would provide.

So I just want to ask both of you, would your companies ever consider taking these kinds of actions that privilege a hostile foreign power over the United States and especially our men and women in uniform.

Ms. Sandberg.

Ms. SANDBERG. I'm not familiar with the specifics of this at all, but based on how you're asking the question, I don't believe so.

Mr. DORSEY. Also no.

Senator COTTON. So thank you for that answer. Mr. Dorsey, let's turn to Dataminr, which is one of the services that provides basically all of Twitter's data. The last time we had an executive from Twitter before this committee in an open setting, I asked about reports that Dataminr had recently ceased its cooperation with the Central Intelligence Agency, at the same time it continued to cooperate with Russia Today and other proxies of Russian intelligence services.

I have since seen reports that Dataminr no longer cooperates with Russia Today or any other proxy of Russian intelligence services. Is that correct?

Mr. DORSEY. That is correct.

Senator COTTON. Did you make that decision personally?

Mr. DORSEY. No, we have a long-standing term against utilizing public Twitter data for ongoing 24/7 surveillance.

Senator COTTON. And that's why you've decided to cease cooperation with the Russian government or proxies like Russia Today?

Mr. DORSEY. No. That's a different matter. This is in regards—

Senator COTTON. Could you explain why you ceased that cooperation then, or that relationship with, Russia Today and other Russian intelligence proxies?

Mr. DORSEY. When we learned of the link of Russia Today and Sputnik, we ceased to allow them to be an advertiser on the platform. We calculated the amount of advertising they did on the platform is \$1.9 million and we donated that to civil liberties nonprofits.

Senator COTTON. Would you now reconsider the decision to cease your cooperation with the Central Intelligence Agency or other American intelligence agencies?

Mr. DORSEY. We are always open to any legal process that an agency would present us, so we don't believe it necessary. This is a global policy around surveillance in general and real-time surveillance. I will state that all this information, because Twitter is public by default, is available to everyone by just going to our service.

Senator COTTON. You see a difference between cooperating with the United States government and the Russian government or the Chinese government?

Mr. DORSEY. Do I see a difference? I'm not sure what you mean.

Senator COTTON. Is Twitter an American company?

Mr. DORSEY. We are an American company.

Senator COTTON. Do you prefer to see America remain the world's dominant global superpower?

Mr. DORSEY. I prefer that we continue to help everywhere we serve and we are pushing towards that, but we need to be consistent about our terms of service and the reason why. And the reason why is we also have a right and a responsibility to protect the privacy of the people on Twitter from constant 24/7 surveillance. And we have other methods to enable any issues that an intelligence community might see, to subpoena and to give us a proper legal order, and we will work with them.

Senator COTTON. I have to say I disagree with any imperative to be consistent between the governments of China and Russia on the one hand and the government of the United States on the other hand. Or would you be consistent or even handed between the government of China and the government of Taiwan?

Mr. DORSEY. What I meant was a consistency of our terms of service. And of course there will always be exceptions, but we want to have those go through due legal process.

Senator COTTON. Let me turn to the actions you've taken about the 2016 election—both of your platforms—and specifically one action you haven't taken. You have removed several accounts as a result of your own investigations and I think some of this committee's work—and I commend your companies for that.

One set of accounts that remain on your platforms are WikiLeaks and Julian Assange. Secretary of State Mike Pompeo, when he was the director of the CIA, characterized WikiLeaks as a non-state hostile intelligence service. This committee has agreed with that assessment now for a couple years in a row, yet, both WikiLeaks, which propagated some of the leaked emails in the 2016 election from the Democrats, remain active on both Facebook and Twitter as does Julian Assange.

Ms. Sandberg, could you explain why Facebook continues to allow their accounts to be active?

Ms. SANDBERG. I'm not going to defend WikiLeaks and I'm not going to defend the actions of any page or actor on our platform. WikiLeaks has been public information. It's available broadly on other media and as such it doesn't violate our terms of service and it remains up on our site.

Senator COTTON. And Mr. Dorsey.

Mr. DORSEY. We also have not found any violation of our terms of service, but you know we are open as always to any law enforcement insight that would indicate a violation of our terms.

Senator COTTON. Thank you. My time has nearly expired. Again, I want to commend your companies for making you available and both of you for appearing. I would urge both of your companies, or any company like yours, to consider whether or not they want to be partners in the fight against our adversaries in places like Beijing and Moscow and Pyongyang and Tehran, as opposed to even-handed or neutral arbiters. Thank you.

Chairman BURR. Senator REED.

Senator REED. Well thank you, Mr. Chairman. Let me begin by thanking you and the Vice Chairman for recognizing my ex officio colleague Senator John McCain. We are both service academy graduates, so we don't know any Latin so we had various translations of ex officio. The one we liked best was real cool. So you were real cool, Mr. Chairman. Thank you.

Thank you both for being here. You have been organizing, based on your comments today, very diligently for the 2018 elections and trying to anticipate malign activities that we saw in 2016.

Have you seen the same type of coherence starting with Ms. Sandberg, from the Federal Government in terms of your ability to contact them to work with them?

Ms. SANDBERG. We've long had very good relationships with law enforcement. We've worked closely with DHS and FBI for a long time. And the FBI's new task force on this has been particularly helpful.

Senator REED. Mr. Dorsey, your comment?

Mr. DORSEY. We've also had really strong relationships with the government. We're always looking for opportunities to improve our partnership and I think if I were to list them out it would be a more regular cadence of meetings. It would be more proactive information about secular trends that they're seeing, not just on our platform, but other platforms and also in other channels and communication methods. And, finally, a consolidation of points to contact—more of a single point of contact. And we do have that consolidation for the 2018 elections, which we're really happy with.

Senator REED. Very good. One of the rules is to follow the money. And you've talked about how you, in terms of political advertising, have identified the citizenship of their advertisers but are you able to trace the monies? It's fairly easy to set up a corporation in the United States, and the money could all be coming from overseas even from some pernicious sources. Do you go that far Ms. Sandberg? And then Mr. Dorsey.

Ms. SANDBERG. Sir, you're right that there a lot of ways to try to game the system and so we are going to keep investing and trying to get ahead of any tactics our opponents would use, including that one.

Senator REED. Mr. Dorsey.

Mr. DORSEY. Sir, we do our best to understand the intent and where people are located and what's behind them, but this is where a strong partnership with government comes in. Because we will not always be able to infer agendas or intent or even location in some cases.

Senator REED. In the dialog that you've talked about with Lauren Forsman, is this one of those topics where you're asking them for information, or they're asking you and you're trying to follow the money, or have you seen any of that, or has it been sort of one of those issues that's just too hard to think about?

Mr. DORSEY. It's both. We have seen proactive outreach from the other side.

Senator REED. But that would be, I think, a critical issue in terms of governing the behavior campaigns, and I would hope that you would continue to work, and we would urge our colleagues in government to work with you, in that regard.

One of the issues, and I think Senator Warner and several others have brought it up, is the prevalence of bots. I'm not a technologist, but it seems to me that you could identify a bot's presence, that you could notify your consumers that 35 percent or 80 percent of these messages have been generated electronically. Is that feasible? And is that something you're doing?

Mr. DORSEY. It's a mixed answer right now. We are able to identify automations and activity coming through our API, and to Senator Warner's comments, we would be able to label that with context. But we are not necessarily as easily able to identify people who might be scripting our website, so making it look like it's an actual human or even the app—make it look like an actual human performing these actions. That becomes much more challenging and unclear.

So in consideration of labeling and context, we need to make sure that when people see that bot label, that they're assuming that everything it's not on is human. We need to make sure that there's a precision and accuracy as we label those things.

Senator REED. Wouldn't there be a value in beginning the labeling process, even with the heavy disclaimer that this identifies only a fraction of potential fictitious actors?

Mr. DORSEY. Yes, it's definitely an idea that we've been considering, especially this past year. It's really up to the implementation at this point.

Senator REED. Ms. Sandberg, your comments?

Ms. SANDBERG. This is one of the ideas I had an opportunity to discuss with Vice Chairman Warner yesterday in his office and is in his white paper, and we're committed to working with you on it.

Senator REED. Thank you. Let me just ask you a question. Going forward, I think we're going to come to a major debate within this country or in the whole world of who owns my data, which rapidly is becoming me. Is it a company like Facebook? Is it a company like Twitter? Which raises the question of do you believe that your users should have the right to control what you do with their data, either selectively, on an individual occurrence, or generically, or even simply purge it at some point? Do you believe that should be—

Ms. SANDBERG. Yes, very strongly. It's your information. You share it with us. If you want to delete it, we delete it. And if you want to take it with you, we enable you to download it and take it with you.

Senator REED. What about for those people who—I think many people—who in the hustle and bustle of everyday, that’s a very cumbersome process? Shouldn’t they be allowed to sort of have a check that says every two months delete it? Or delete it as soon as I put it in?

Ms. SANDBERG. Yes, and we’re working on some of those tools, and we’ve improved. We’ve made it easier to understand what information we have, how we’re getting it, and how we use it. And we’re going to continue to iterate here.

Senator REED. Mr. Dorsey, the same question.

Mr. DORSEY. We do believe people should have complete control over their—of their data. Again, Senator Warner brought up an interesting point earlier, which is—I don’t believe that there’s a real understanding of the exchange being made in terms of people performing activities on these services and services like Twitter, and how they can actually see that as an exchange—an exchange of value. And those are things I would love to think a lot more about, how do we make that more clear? And I think that goes back to the incentives conversation.

Senator REED. Thank you. Thank you, Mr. Chairman.

Chairman BURR. Thank you, Senator Reed, and I thank all the members for their questions and our panelists for their answers. I’m going to turn to the Vice Chairman for any last comments he might have.

Vice Chairman WARNER. One, I want to thank you both. I want to thank you for the spirit you brought to this, some of the suggestions—your responses to some of the suggestions. I wish our members were still here, because I think they all performed extraordinarily well.

I take away from this three or four quick points. One, very much appreciate, Mr. Dorsey, your acknowledgement that we ought to move towards—and I guess Ms. Sandberg echoed this as well—some ability to indicate to users whether they’re being contacted by a machine or a human being, recognizing there’s technical difficulties, and also acknowledging that just because it’s a bot that does not inherently mean it’s good or bad. It just must be a data point that an individual ought to have as they make determinations going forward.

I also really appreciated, Ms. Sandberg, your notion that not only should users have access to all of the information that you or others are collecting, but as we work through to this—how you monetize that and let users know the value of their data, I think that increased price transparency—and I was very grateful at your willingness to at least consider that, because I think that would go a long way towards making this exchange better understood by individuals.

Also, and I didn’t get a chance to really get into this at length, but you and I have had this conversation in the past around data portability. I don’t want to make the complete analogies—an old telecom guy—but when number portability came around, we got a lot more competition in the wireless industry and elsewhere. Data portability—I know you make it available right now—but in an easy, user-friendly format that can move from platform to platform,

I think would be extraordinarily important in terms of making sure that we continue to have competition in this space.

And then finally, I also appreciated your comment—I think we’re going to have more and more of these areas where manipulation may take place that actually incents violence. We both cited the horrible example of what’s happened with the Rohingya in Myanmar, but I appreciate your comment that you’ve said that Facebook ought to have both a moral and legal obligation if there are sites that are inciting violence and take those down. Getting from that idea into how we spell that all out will be a challenge, but I appreciate your willingness to work with me on it.

So Mr. Chairman, thank you for the fourth hearing on this. I think it was very, very important, and I hope our committee will continue to take the lead on these subjects.

Chairman BURR. I thank the Vice Chairman. I would ask both of you if there are any rules, such as antitrust, FTC regulations or guidelines, that are obstacles to collaboration between companies, I hope you’ll submit for the record where those obstacles are so that we can look at the appropriate steps that we could take as a committee to open those avenues up.

I want to thank both of you for appearing today and for your continued efforts to help find a solution to the challenging problem. This hearing represents the capstone of the fourth piece of the committee’s investigation into Russian interference in the 2016 elections. So far we’ve completed our inquiry into the attempted hack of State elections infrastructure, the intelligence community assessment on Russian activities in recent U.S. elections, the Obama Administration’s policy response to those operations.

With your testimony today at this, the fourth hearing we’ve held on social media, we heard the top-level perspective on how to address foreign influence operations on your platforms. When this committee began its investigation into Russian interference in the 2016 elections, neither Mark nor I fully appreciated how easily foreign actors could use social media to manipulate how Americans form their views.

Like most technology, social media has the capacity to be used for good as intended, but also to advance agendas of those bent on manipulation and destruction. Given the amount of information companies like Google collect on each and every American, it is also too easy for bad actors to craft a message that appears tailored just for you.

The Russians undertook a structured influence campaign not against the American government but against the American people. Moscow saw the issues that talking heads yell about on cable news—race, religion, immigration, and sexual orientation—and they used those to sow discord and to foment chaos. They leveraged our social media to undermine our political system as well, but make no mistake, Russia neither leans left nor right. It simply seeks turmoil. A weak America is good for Russia.

I think it is also important to highlight that there is a very human component to all of this. No single algorithm can fix the problem. Social media is part of our daily lives. It serves as the family newsletter, a place to share life’s personal joys and sorrows,

a way to communicate one's status during a crisis, and everything in between.

Unfortunately, other states are now using the Russian playbook, as evidenced by the recently uncovered Iranian influence operations. We're at a critical inflection point. Will using social media to sow discord become an acceptable tool of statecraft? How many copycats will we see before we take this seriously and find solutions? Your companies must be at the forefront in combating those issues. You know your algorithms, your customers, and your data collection capabilities better than any government entity does—or should. Still, the burden is not entirely on your shoulders. Government, civil society, and the public will partner with you.

I'd like to take just a moment to thank our staff. They have worked diligently to uncover the scope of the problem. Their research has been thorough. Their efforts are seamlessly bipartisan and their drive to defend the public against foreign influence should make Americans watching today proud.

There is no clear and easy path forward. We understand the problem and it is a First Amendment issue. We cannot regulate around the First Amendment, but we also cannot ignore the challenge. I am confident that working together we can find a solution and a path forward that will only make us stronger, more connected, more prepared to face down those who seek to weaken our democracy.

For your participation in being part of the solution, we thank you immensely today.

This hearing is now adjourned.

[Whereupon, at 12:11 p.m., the hearing was adjourned.]

Supplemental Material

October 26, 2018

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

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

Thank you for your questions for the record from the September 5, 2018 Hearing titled Foreign Influence Operations' Use of Social Media Platforms. Per your request, attached are the answers for the record to your questions.

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

Sincerely,

Facebook, Inc.

facebook

Address: 1601 Willow Road
Menlo Park, CA 94025

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

[From Chairman Burr]

1. Aleksandr Kogan served as director of Global Science Research (GSR) where he used an app to harvest data from as many as 87 million Facebook users. Facebook has said publicly that Kogan claimed the data would only be used for academic purposes and then “lied to us” in passing the content to Cambridge Analytica.
- **Did Facebook data scientists co-author academic papers with GSR co-founders Aleksandr Kogan and Joseph Chancellor?**
 - **If yes, how does this reconcile with Facebook’s asserting a complete unawareness as to GSR’s user data harvesting practices?**

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

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

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

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

2. **On February 6, 2018, the day after the Senate Commerce and Judiciary hearing, Facebook terminated Joseph Chancellor’s employment. What were the circumstances of his termination?**
- **What was the hire date (month and year) of Joseph Chancellor, co-founder of GSR?**

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

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

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

[From Vice Chairman Warner]

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

- **What is the current status of Facebook's engagement with China?**
- **Do you have existing plans for attempting to enter the Chinese market? If yes, please describe.**
- **Are there any current discussions underway within Facebook about entering China? If yes, please describe.**

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

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

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

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

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

- **Give people a voice:** The one phrase in our mission that has never changed is “give people the power,” and one of the ways we do that is by giving people a voice. This means we err on the side of free expression—even when that means defending the right of people we deeply disagree with to say things that are controversial or offensive. Of course, there are limits. We don’t allow content that incites violence or attacks people, whether that’s terrorism or bullying or hate.
- **Build connection and community:** Our services help people connect more, and when they’re at their best, they also bring people closer together. That’s why this year we reworked News Feed to prioritize meaningful social interactions over passive consumption. In order for a community to be cohesive, it must share enough common ground—so while we give everyone a voice, we must make sure misinformation doesn’t spread virally and high quality, broadly trusted information is available to all.
- **Serve everyone:** Everyone deserves access to these tools. That’s why we operate in countries where we might lose money, why we work on Internet.org to spread connectivity to people who can’t even afford it, why our business model is ads—so our service can be free for everyone. And it’s also why, when a country passes a law limiting voice or that conflicts with one of our other principles, we fight to make sure the service remains available for as many people as possible.
- **Keep people safe and protect privacy:** People try to use our services for good and bad, and we have a responsibility promote the good and prevent harm. That’s why we have the initiatives on counterterrorism and self-harm. That’s why we have more than doubled the number of people working on safety and security and now have over 20,000. And that’s why, even though we care about giving people a voice, we take down a lot of content that is bullying, harassing, and attacking people.
- **Promote economic opportunity:** We talk a lot about the social aspect of community, but strong communities also provide people opportunity. Through our work helping small businesses grow, we aim to create more jobs and opportunity than any other company out there. Our services empower people, and our work supporting economic opportunity—whether it’s through Marketplace, Pages, WhatsApp, Messenger or Instagram—is a fundamental part of what we stand for.
- **Which of those values will you weigh when considering potential engagement in China?**

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

7. **Under the terms of your 20-year consent decree with the FTC, Facebook was required to establish a "comprehensive privacy policy" to undertake, among other things, "the identification of reasonably foreseeable, material risks, both internal and external, that could result in [Facebook's] unauthorized collection, use, or disclosure of covered information and an assessment of the sufficiency of any safeguards in place to control these risks." The consent decree says this should extend to "product design, development, and research."**
- **Does Facebook believe its failure to identify and address the privacy concerns of allowing data access to third-party applications like Aleksandr Kogan's Global Science Research (GSR) is consistent with the "reasonably foreseeable" language of the FTC consent decree?**

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

See Response to above Question.

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

- **Is that accurate?**

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

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

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

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

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

See Response to above Question.

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

See Response to above Question.

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

See Response to above Question.

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

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

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

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

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

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

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

See Response to above Question.

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

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

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

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

- **Did Facebook ever notify users that their data was being accessed in spite of this setting, and did you note this to the Federal Trade Commission?**
- **Approximately how many users did the applications that were granted this special access have, in total?**

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

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

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

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

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

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

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

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

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

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

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

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

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

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


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

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

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

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

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



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

- Is that accurate?
- Why didn’t you take action earlier?
- How much are you investing in addressing the misinformation and violence prevention issues in Myanmar?
- What about in other parts of the world where similar threats are possible?

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

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

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

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

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

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

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

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

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

See Response to above Question.

- 23. What is Facebook’s current policy on the posting or promotion of hacked emails on your platform?**

We prohibit any content that is claimed or confirmed to have come from a hacked source. In rare situations and on a case-by-case basis, we may choose to allow content that is newsworthy, significant, or important to the public interest even if it otherwise violates our policies. We do this only after weighing the public interest value of the content against the risk of real-world harm.

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

- **How is Facebook complying with GDPR?**

- **Are there protections that will flow to U.S. users as a result?**

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

- **Ads and Pages transparency.** Advertising should be transparent: users should be able to see all the ads an advertiser is currently running on Facebook, Instagram and Messenger. And for ads with political or issue content, we've created an archive that will hold ads with political or issue content for 7 years—including information about ad impressions and spend, as well as demographic data such as age, gender, and location. And people everywhere can see all the ads that Page is running on Facebook. We also announced in April that people who manage Pages with large numbers of followers will need to be verified. Those who manage large Pages that do not clear the process will no longer be able to post. This will make it much harder for people to administer a Page using a fake account, which is strictly against our policies. We will also show people additional context about Pages to help people have more information to evaluate their content. For example, you can see whether a Page has changed its name.
- **Verification and labeling.** Every advertiser will now need to confirm their ID and location before being able to run any ads with political or issue content in the US and certain other countries. All ads with political or issue content will also clearly state who paid for them.
- **Updating targeting.** We want ads on Facebook to be safe and civil. We thoroughly review the targeting criteria advertisers can use to ensure they are consistent with our principles. As a result, we removed nearly one-third of the targeting segments used by the IRA. We continue to allow some criteria that people may find controversial. But we do see businesses marketing things like historical books, documentaries, or television shows using them in legitimate ways.

- **Better technology.** Over the past year, we've gotten increasingly better at finding and disabling fake accounts. We now block millions of fake accounts each day as people try to create them—and before they've done any harm. This is thanks to improvements in machine learning and artificial intelligence, which can proactively identify suspicious behavior at a scale that was not possible before—without needing to look at the content itself.
- **Action to tackle fake news.** We are working hard to stop the spread of false news. We work with third-party fact-checking organizations to limit the spread of articles rated false. To reduce the spread of false news, we remove fake accounts and disrupt economic incentives for traffickers of misinformation. We also use various signals, including feedback from our community, to identify potential false news. In countries where we have partnerships with independent third-party fact-checkers, stories rated as false by those fact-checkers are shown lower in News Feed. If Pages or domains repeatedly create or share misinformation, we significantly reduce their distribution and remove their advertising rights. We also want to empower people to decide for themselves what to read, trust, and share. We promote news literacy and work to inform people with more context. For example, if third-party fact-checkers write articles about a news story, we show them immediately below the story in the Related Articles unit. We also notify people and Page Admins if they try to share a story, or have shared one in the past, that's been determined to be false. In addition to our own efforts, we're learning from academics, scaling our partnerships with third-party fact-checkers and talking to other organizations about how we can work together.
- **Significant investments in security.** As part of our larger company investment in the space, we have more than doubled the number of people working on safety and security and now have over 20,000. We expect these investments to impact our profitability. But the safety of people using Facebook needs to come before profit.
- **Industry collaboration.** Recently, we joined more than 60 global tech and security companies in signing a TechAccord pact to help improve security for everyone.
- **Information sharing and reporting channels.** In the 2017 German elections, we worked closely with the authorities there, including the Federal Office for Information Security (BSI). This gave them a dedicated reporting channel for security issues related to the federal elections.
- **Tracking 40+ elections.** We deployed new tools and teams to proactively identify threats in the run-up to specific elections. We first tested this effort during the Alabama Senate election, and have continued these efforts for elections around the globe, including the US midterms. Last year we used public service announcements to help inform people about fake news in 21 separate countries, including in advance of French, Kenyan and German elections.
- **Action against the Russia-based IRA.** In April, we removed 70 Facebook and 65 Instagram accounts—as well as 138 Facebook Pages—controlled by the IRA primarily targeted either at people living in Russia or Russian-speakers around the

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

[From Senator Feinstein]

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

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

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

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

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

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

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

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

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

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

- **Ads and Pages transparency.** Advertising should be transparent: users should be able to see all the ads an advertiser is currently running on Facebook, Instagram and Messenger. And for ads with political or issue content, we've created an archive that will hold ads with political or issue content for 7 years—including information about ad impressions and spend, as well as demographic data such as age, gender, and location. And people everywhere can see all the ads that Page is running on Facebook. We also announced in April that people who manage Pages with large numbers of followers will need to be verified. Those who manage large Pages that do not clear the process will no longer be able to post. This will make it much harder for people to administer a Page using a fake account, which is strictly against our policies. We will also show people additional context about Pages to help people have more information to evaluate their content. For example, you can see whether a Page has changed its name.
- **Verification and labeling.** Every advertiser will now need to confirm their ID and location before being able to run any ads with political or issue content in the US and certain other countries. All ads with political or issue content will also clearly state who paid for them.
- **Updating targeting.** We want ads on Facebook to be safe and civil. We thoroughly review the targeting criteria advertisers can use to ensure they are consistent with our

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

- **Better technology.** Over the past year, we've gotten increasingly better at finding and disabling fake accounts. We now block millions of fake accounts each day as people try to create them—and before they've done any harm. This is thanks to improvements in machine learning and artificial intelligence, which can proactively identify suspicious behavior at a scale that was not possible before—without needing to look at the content itself.
- **Action to tackle fake news.** We are working hard to stop the spread of false news. We work with third-party fact-checking organizations to limit the spread of articles rated false. To reduce the spread of false news, we remove fake accounts and disrupt economic incentives for traffickers of misinformation. We also use various signals, including feedback from our community, to identify potential false news. In countries where we have partnerships with independent third-party fact-checkers, stories rated as false by those fact-checkers are shown lower in News Feed. If Pages or domains repeatedly create or share misinformation, we significantly reduce their distribution and remove their advertising rights. We also want to empower people to decide for themselves what to read, trust, and share. We promote news literacy and work to inform people with more context. For example, if third-party fact-checkers write articles about a news story, we show them immediately below the story in the Related Articles unit. We also notify people and Page Admins if they try to share a story, or have shared one in the past, that's been determined to be false. In addition to our own efforts, we're learning from academics, scaling our partnerships with third-party fact-checkers and talking to other organizations about how we can work together.
- **Significant investments in security.** As part of our larger company investment in the space, we have more than doubled the number of people working on safety and security and now have over 20,000. We expect these investments to impact our profitability. But the safety of people using Facebook needs to come before profit.
- **Industry collaboration.** Recently, we joined more than 60 global tech and security companies in signing a TechAccord pact to help improve security for everyone.
- **Information sharing and reporting channels.** In the 2017 German elections, we worked closely with the authorities there, including the Federal Office for Information Security (BSI). This gave them a dedicated reporting channel for security issues related to the federal elections.
- **Tracking 40+ elections.** We deployed new tools and teams to proactively identify threats in the run-up to specific elections. We first tested this effort during the Alabama Senate election, and have continued these efforts for elections around the globe, including the US midterms. Last year we used public service announcements

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

- **Do you commit to notifying the public should Facebook identify other foreign information operations?**
- **Will Facebook commit to institutionalizing the alerting of users who have been exposed to foreign information operations?**

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

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

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

We agree that information sharing among companies and government is critical to combating constantly evolving cyber threats. We have been working with many others in the technology industry, including Google and Twitter, on this issue, building on our long history of working together on issues like child safety and counterterrorism. We also have a history of working successfully with the DOJ, the FBI, and other law enforcement to address a wide variety of threats to our platform, and we look forward to continuing to work with law enforcement and government on these issues.

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

We agree that information sharing among companies and government is critical to combating constantly evolving cyber threats. We have been working with many others in the technology industry, including Google and Twitter, on this issue, building on our long history of working together on issues like child safety and counterterrorism. We also have a history of working successfully with the DOJ, the FBI, and other law enforcement to address a wide variety of threats to our platform, and we look forward to continuing to work with law enforcement and government on these issues. We'd be happy to discuss these issues further with your staff.

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

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

- **Would a shared standard and the deployment of similar software used to detect spam and copyrighted material, facilitate the automated blocking of such content across all four platforms?**
- **In the interest of transparency, would you make this database open to the public or researchers to know which images are prohibited?**

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

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

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

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

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

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

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

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

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

[From Senator Wyden]

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

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

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

- In total, more than 290,000 accounts followed at least one of these Pages, the earliest of which was created in March 2017. The latest was created in May 2018.
- The most followed Facebook Pages were "Aztlan Warriors," "Black Elevation," "Mindful Being," and "Resisters." The remaining Pages had between zero and 10 followers, and the Instagram accounts had zero followers.

- There were more than 9,500 organic posts created by these accounts on Facebook, and one piece of content on Instagram.
- The 32 Pages and accounts ran about 150 ads for approximately \$11,000 on Facebook and Instagram, paid for in US and Canadian dollars. The first ad was created in April 2017, and the last was created in June 2018.
- The Pages created about 30 events since May 2017. About half had fewer than 100 accounts interested in attending. The largest had approximately 4,700 accounts interested in attending, and 1,400 users said that they would attend.

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

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

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

- 35. Facebook's statement noted that the administrators of the fake account "connected with admins from five legitimate Pages to co-host the event," and that Facebook "reached out to the admins of the five other Pages to update them on what happened."**
- **What is Facebook's policy in circumstances in which fake accounts have joined with legitimate, but unwitting American political actors in promoting events or causes?**

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

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

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

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

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

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

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

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

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

See Response to Question 36.

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

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

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

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

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

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

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

See Response to Question 39.

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

See Response to Question 38.

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

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

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

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

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

- **Ads and Pages transparency.** Advertising should be transparent: users should be able to see all the ads an advertiser is currently running on Facebook, Instagram and Messenger. And for ads with political or issue content, we've created an archive that will hold ads with political or issue content for 7 years—including information about ad impressions and spend, as well as demographic data such as age, gender, and location. And people everywhere can see all the ads that Page is running on Facebook. We also announced in April that people who manage Pages with large numbers of followers will need to be verified. Those who manage large Pages that do not clear the process will no longer be able to post. This will make it much harder for people to administer a Page using a fake account, which is strictly against our policies. We will also show people additional context about Pages to help people have more information to evaluate their content. For example, you can see whether a Page has changed its name.
- **Verification and labeling.** Every advertiser will now need to confirm their ID and location before being able to run any ads with political or issue content in the US and certain other countries. All ads with political or issue content will also clearly state who paid for them.
- **Updating targeting.** We want ads on Facebook to be safe and civil. We thoroughly review the targeting criteria advertisers can use to ensure they are consistent with our principles. As a result, we removed nearly one-third of the targeting segments used by the IRA. We continue to allow some criteria that people may find controversial. But we do see businesses marketing things like historical books, documentaries, or television shows using them in legitimate ways.
- **Better technology.** Over the past year, we've gotten increasingly better at finding and disabling fake accounts. We now block millions of fake accounts each day as people try to create them—and before they've done any harm. This is thanks to improvements in machine learning and artificial intelligence, which can proactively

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

- **Action to tackle fake news.** We are working hard to stop the spread of false news. We work with third-party fact-checking organizations to limit the spread of articles rated false. To reduce the spread of false news, we remove fake accounts and disrupt economic incentives for traffickers of misinformation. We also use various signals, including feedback from our community, to identify potential false news. In countries where we have partnerships with independent third-party fact-checkers, stories rated as false by those fact-checkers are shown lower in News Feed. If Pages or domains repeatedly create or share misinformation, we significantly reduce their distribution and remove their advertising rights. We also want to empower people to decide for themselves what to read, trust, and share. We promote news literacy and work to inform people with more context. For example, if third-party fact-checkers write articles about a news story, we show them immediately below the story in the Related Articles unit. We also notify people and Page Admins if they try to share a story, or have shared one in the past, that's been determined to be false. In addition to our own efforts, we're learning from academics, scaling our partnerships with third-party fact-checkers and talking to other organizations about how we can work together.
- **Significant investments in security.** As part of our larger company investment in the space, we have more than doubled the number of people working on safety and security and now have over 20,000. We expect these investments to impact our profitability. But the safety of people using Facebook needs to come before profit.
- **Industry collaboration.** Recently, we joined more than 60 global tech and security companies in signing a TechAccord pact to help improve security for everyone.
- **Information sharing and reporting channels.** In the 2017 German elections, we worked closely with the authorities there, including the Federal Office for Information Security (BSI). This gave them a dedicated reporting channel for security issues related to the federal elections.
- **Tracking 40+ elections.** We deployed new tools and teams to proactively identify threats in the run-up to specific elections. We first tested this effort during the Alabama Senate election, and have continued these efforts for elections around the globe, including the US midterms. Last year we used public service announcements to help inform people about fake news in 21 separate countries, including in advance of French, Kenyan and German elections.
- **Action against the Russia-based IRA.** In April, we removed 70 Facebook and 65 Instagram accounts—as well as 138 Facebook Pages—controlled by the IRA primarily targeted either at people living in Russia or Russian-speakers around the world including from neighboring countries like Azerbaijan, Uzbekistan, and Ukraine. The IRA has repeatedly used complex networks of inauthentic accounts to deceive and manipulate people in the US, Europe, and Russia—and we don't want them on Facebook anywhere in the world. In July, we removed 32 Pages and

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

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

- **How many of these accounts had some form of two-factor authentication enabled on their accounts?**
- **How many of these accounts were secured with a two-factor authentication security key?**

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

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

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

- **How many of these accounts had some form of two-factor authentication enabled on their accounts.**
- **How many of these accounts were secured with a two-factor authentication security key?**

We do not maintain public statistics on this issue.

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

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

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

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

See Response to above Question.

- **If the answer to the question above is yes, were these voter discouragement ads targeted at people of any particular race or ethnic group?**
 - **Were these voter discouragement ads predominantly targeted at people expected to vote for one party or the other?**

See Response to above Question.

- **Has any foreign government, their agent, or other foreign entity ever used Custom Audiences to target individuals in the United States?**
 - **If yes, please provide a full accounting of each case, including the party that purchased the advertising, the foreign government sponsor (if applicable), and the number of users that saw or interacted with the content (e.g. clicked or shared).**

See Response to above Question.

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

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

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

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

- **Can users opt out of being targeted with Custom Audiences?**

- **If no, why not?**

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

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

- **Does Facebook have a policy of shutting down pages and accounts that seek to suppress voting, regardless of whether they are found to be inauthentic?**
 - **If yes, to what kind of content does Facebook apply that policy (e.g., content discouraging people from voting, content providing inaccurate information on how or when to vote, etc.)?**

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

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

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

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

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

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

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

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

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

- **Is Facebook aware of any instances in which user data obtained by Kogan was subsequently used to target Facebook users, either during the 2016 Election, or at any other time?**
- **Please describe in detail all uses of user data obtained by Kogan of which Facebook is aware.**
- **What efforts have been made to ensure that user data obtained by Kogan has been completely deleted, and cannot be used in the future by any party, for any purpose?**

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

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

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

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

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

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

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

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

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

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

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

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

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

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

- Training staff, candidates, government officials, advocacy groups, and others during live events how to take common sense safety precautions, including turning on two-factor authentication.
- Creating a video explainer on two-factor authentication specifically for those involved in politics, available at: <https://politics.fb.com/learn-the-basics/#component-1-secure-your-account>.

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

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

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

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

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

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

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

- **If yes:**

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

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

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

Facebook has discussed its device integration partnerships with the FTC.

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

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

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

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

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

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

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

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

See Response to above Question.

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

See Response to above Question.

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

See Response to above Question.

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

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

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

See Response to above Question.

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

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

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

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

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

The nature of the partnership varied from partner to partner.

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

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

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

See Response to above Question.

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

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

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

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

- **Is it true that Facebook’s terms of service bar certain journalism and research focused on the platform?**
- **If you have concerns about the Knight Institute’s proposal, are there modifications to the proposal that would address your concerns while safeguarding digital journalism and researched focused on Facebook’s platform?**

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

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

[From Senator Lankford]

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

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

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

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

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

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

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

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

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

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

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

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

[From Senator Harris]

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

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

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

54. What is your definition of “organic content?”

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

55. What percent of your content is not organic?

The majority of content on Facebook is organic.

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

- We estimate that fake accounts represented approximately 3 percent to 4 percent of monthly active users (MAU) on Facebook;
- We disabled 583 million fake accounts; and
- 98.5 percent of fake accounts acted on were flagged by Facebook before users reported them.
- **In Facebook's estimation, how many more accounts are plausibly linked to Russia's disinformation campaign? If you cannot provide a specific number, please provide an estimate.**
- **In Facebook's estimation, how many more accounts are plausibly linked to other countries' disinformation campaigns? If you cannot provide a specific number, please provide an estimate.**

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

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

We are committed to finding and removing fake accounts. We continue to make improvements to our efforts to more effectively detect and deactivate fake accounts to help reduce the spread of spam, false news, and misinformation. We continually update our technical systems to identify, checkpoint, and remove inauthentic accounts, and we block millions of attempts to register fake accounts every day. These systems examine thousands of detailed account attributes and prioritize signals that are more difficult for bad actors to disguise, such as their connections to others on our platform. As with all security threats, we have been incorporating new insights into our models for detecting fake accounts, including information specific to election issues.

We do not share detailed descriptions of how our tools work in order to avoid providing a road map to bad actors who are trying to avoid detection. When we suspect that an account is inauthentic, we typically enroll the account in a checkpoint that requires the account holder to provide additional information or verification. We view disabling an account as a severe sanction, and we want to ensure that we are highly confident that the account violates our policies before we take permanent action. When we have confirmed that an account violates our policies, we remove the account.

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

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

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

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

Stopping the abuse of fake accounts and malicious bot activity is a focus for many teams, some more directly and some in more of a supportive role. For example, we are expanding our threat intelligence team, and more broadly, we have more than doubled the number of people working on safety and security and now have over 20,000. We expect to have at least 250 people specifically dedicated to safeguarding election integrity on our platforms, and that number does not include the thousands of people who will contribute to this effort in some capacity. We also continue to make improvements to our efforts to more effectively detect and deactivate fake accounts to help reduce the spread of spam, false news, and misinformation. We continually update our technical systems to identify, checkpoint, and remove inauthentic accounts, and we

block millions of attempts to register fake accounts every day. These systems examine thousands of detailed account attributes and prioritize signals that are more difficult for bad actors to disguise, such as their connections to others on our platform. As with all security threats, we have been incorporating new insights into our models for detecting fake accounts, including information specific to election issues.

We publish information and metrics about fake accounts at <https://transparency.facebook.com/community-standards-enforcement#fake-accounts> and in our quarterly SEC filings.

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

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

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

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

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

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

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

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

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

- **Has Facebook implemented this technology?**
- **Does Facebook categorize users into socioeconomic groups?**

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

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

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

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

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

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

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

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

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

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

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

Questions for the Record
Senate Select Committee on Intelligence
Hearing on Foreign Influence Operations Using Social Media September 17, 2018
Questions for the Record for Mr. Jack Dorsey, Chief Executive Officer, Twitter.

[From Vice Chairman Warner]

1. According to reports about social media usage during the Catalan Independence Referendum in Spain, “[Sputnik and RT], both financed by the Kremlin, managed to see their links shared more than those from Spanish public networks EFE and RTVE, or those of private international publications such as The Guardian and CNN.” This information operation utilized Russian bots on Twitter and came almost a year after the Russian interference in the U.S. election which used similar tactics.

- **Given that similar tactics were used and this event happened almost a year after the 2016 US election, why was Twitter not able to detect and stop this information operation by Russian-linked operatives?**

We remain vigilant about identifying and eliminating abuse on the platform perpetrated by hostile foreign actors, and we will continue to invest in resources and leverage our technological capabilities to do so. Twitter’s main focus is promoting healthy public discourse through protection of the democratic process. Twitter continues to engage in intensive efforts to identify and combat state-sponsored hostile attempts to abuse social media for manipulative and divisive purposes. We now possess a deeper understanding of both the scope and tactics used by malicious actors to manipulate our platform and sow division across Twitter more broadly. Our efforts enable Twitter to fight this threat while maintaining the integrity of peoples’ experience on the service and supporting the health of conversation on our platform. Our work on this issue is not done, nor will it ever be.

Any amount of election interference in any election is unacceptable to us. Twitter prioritizes identifying suspicious account activity, such as exceptionally high-volume Tweeting with the same hashtag or mentioning the same @handle without a reply from the account being addressed, and requires an individual using the platform to confirm control. Twitter has also increased its use of challenges intended to catch automated accounts, such as reCAPTCHAs, that require individuals to identify portions of an image or type in words displayed on screen, and password reset requests that protect potentially compromised accounts. Twitter is also in the process of implementing mandatory email or cell phone verification for all new accounts.

We conducted a comprehensive analysis of accounts that promoted election-related Tweets on the platform throughout 2016 in the form of paid ads. We reviewed nearly 6,500 accounts and our findings showed that approximately one-tenth of one-percent—only nine of the

total number of accounts—were Tweeting election-related content and linked to Russia. The two most active accounts out of those nine were affiliated with Russia Today (“RT”), which Twitter subsequently barred from advertising on Twitter. And Twitter is donating the \$1.9 million that RT spent globally on advertising to academic research and partnerships focused on election and civic engagement.

In any election, we take action on accounts that break our rules. This included accounts Tweeting about the Catalan referendum. Some accounts outside of Spain did engage in the conversation, including sharing international media links. As part of our standard review patterns and related enforcement actions, we took action on a number of spammy accounts, including some posting at a high volume. We cannot make definitive attributions of these accounts. As is always the case, we are committed to protecting the integrity of the public conversation and that is never more important than during elections. We will always err on the side of transparency. We recently made the full database of potentially state-backed interference on Twitter available with the goal of empowering researchers to conduct independent, investigatory analysis, including the Catalan referendum.

2. This is a screenshot of an advertisement run on Twitter during the 2016 election season, suggesting it was possible to vote via text message:

- **Have you conducted analysis into who created these ads, and which country or countries those individuals are located in, and if so, can you share your findings?**
- **Were these voter discouragement ads targeted at people of any particular race or ethnic group?**
- **Since 2016, have you changed your policies and operations in any way to disallow similar ads that create confusion as to where or how to vote?**

Twitter unequivocally condemns the use of our platform for any election interference activity. Tweets aimed at suppressing voter turnout generally are surfaced through reports from people using our service. This content is reviewed, then promptly removed as illegal interference with voting rights: the content is either restricted as inaccessible pending deletion by the individual (i.e., other individuals on the platform are unable to see the content) or the responsible accounts are permanently suspended. In addition, in order to proactively surface additional Tweets with a given text-to-vote meme, Twitter utilizes technology for identifying instances where the same image appears across multiple Tweets. Content identified through this process is then subject to manual review.

Depending on the number of violations for any given account disseminating voter suppression Tweets, Twitter will either restrict access to the Tweet or suspend the account. During the period leading up to the 2016 election, for example, Twitter labeled and restricted

access to the vote-to-text Tweets pursuant to the Twitter User Agreement, which contains the Twitter Terms of Service, Twitter Privacy Policy, and Twitter Rules. According to the unlawful use provision of the Twitter Rules, individuals are prohibited from using Twitter's "service for any unlawful purpose or in furtherance of illegal activities" and "[i]nternational users agree to comply with all local laws regarding online conduct and acceptable content."

Because the Tweet in question appeared to mislead people into believing that they could vote online or vote by text, Twitter viewed the Tweets as an unlawful interference with the voting process. Twitter labeled as "restricted pending deletion" a total of 918 such Tweets from 529 Twitter accounts, which rendered the Tweets inaccessible and disabled the accounts' ability to use the platform until those Tweets were deleted. In connection with this activity, Twitter also suspended 106 of those accounts, a majority of which were found to be in violation of the Twitter Rules prohibiting spam, including posting duplicate content over multiple accounts or multiple duplicate updates on one account. In a few instances, however, Twitter suspended accounts of people who shared the voting-related content and had previous, but otherwise unrelated, violations of the Twitter Rules against abusive behavior.

The specific Tweet identified in the question was an organic Tweet, not a Promoted Tweet, and they could not be targeted to any particular individual on the platform. In this specific instance, Twitter identified and suspended this Tweet on November 6, 2016. We permanently suspended the account -- which we believe was located within the United States -- on November 7, 2016 as it violated our rules against repeatedly posting content with the intent to deceive.

Twitter identified, but did not take action against, an additional 286 Tweets of the relevant content from 239 Twitter accounts. With respect to those Tweets, Twitter determined that they propagated the content in order to refute the message and alert other people that the information is false and misleading. And partly as a result of our enforcement decisions, those refuting Tweets generated significantly greater engagement across the platform compared to the Tweets spreading the misinformation—eight times as many impressions, engagement by ten times as many people on the platform, and twice as many replies.

Since 2016, we have taken a number of important steps to safeguard the integrity of elections. As part of these efforts, we have developed well-established relationships with law enforcement agencies active in this arena, including the Federal Bureau of Investigation Foreign Influence Task Force and the Department of Homeland Security's Election Security Task Force. We look forward to continued cooperation with them on these issues, as only they have access to information critical to our joint efforts to stop bad faith actors.

Additionally, to further promote information sharing and to tap into the experience and expertise of active stakeholders, we recently updated a Partner Support Portal. Our goal is to expedite our response to reports from people active in the election arena. This includes election support organizations, U.S.-based research organizations, and universities and academics who study the spread of misinformation in the media. Reports from accounts within this select group are expedited and can be actioned promptly.

3. Twitter has a tool called Tailored Audiences, which is similar to Facebook’s Custom Audiences and allows advertisers to upload lists of specific users to target them with ads. Your policies state that advertisers who use Tailored Audiences must obtain consent from users about the data they have acquired.

- **How do you ensure that all of your advertisers actually follow those policies when there are bad actors like the Internet Research Agency who are willing to break the law and illicit data freely available to them?**

Tailored audiences is a product feature that advertisers use to target existing people on the platform and customers. For example, Advertisers can reach existing customers by uploading a list of email addresses. Advertisers can also target those that have recently visited the advertisers’ websites or reach those that have taken specific action in an application, such as installation or registration.

Twitter informs individuals on the platform about Tailored Audiences in several ways. For example, Twitter describes this activity in its Privacy Policy, an “Ads info” footer on twitter.com, and the “Why am I seeing this ad?” section of the drop down menu on Twitter ads themselves. Each of these locations describe interest-based advertising on Twitter and explain how to use the associated privacy controls. In addition, the “Your Twitter Data” tool allows individuals on Twitter to download a list of advertisers that have included them in a Tailored Audience.

If people on Twitter do not want Twitter to show them Tailored Audience ads on and off of Twitter, there are several ways they can turn off this feature: using their Twitter settings, they can visit the Personalization and data settings and adjust the Personalize ads setting; if they are on the web, they can visit the Digital Advertising Alliance’s consumer choice tool at optout.aboutads.info to opt out of seeing interest-based advertising from Twitter in their current browser; if they do not want Twitter to show them interest-based ads in Twitter for iOS on their current mobile device, they can enable the “Limit Ad Tracking” setting in their iOS phone’s settings; and if they do not want Twitter to show them interest-based ads in Twitter for Android

on their current mobile device, they can enable “Opt out of Ads Personalization” in an Android phone’s settings.

In addition to explaining Tailored Audiences to people on the platform, offering them several ways to disable the feature, and enabling them to view the advertisers who have included them in Tailored Audiences, as described above, the Tailored Audience legal terms require that advertisers have secured all necessary rights, consents, waivers, and licenses for use of data.

Advertisers are also required to provide all people from whom the data is collected with legally-sufficient notice that fully discloses the collection, use, and sharing of the data that is provided to Twitter for purposes of serving ads targeted to people’s interest, and legally sufficient instructions on how they can opt out of interest-based advertising on Twitter.

4. Mr. Dorsey, you have indicated your company’s strong support for the Honest Ads Act. Thank you for your support and your efforts to largely abide by the terms of that legislation.

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

Twitter supports the goals of the Honest Ads Act. Through our own initiative, we have announced voluntary, industry-leading steps to improve transparency and accountability in our ads platform that strongly aligns with the goals and standards in the Act. In fact, in some cases, our new transparency requirements go further than the draft legislation—for example, by requiring transparency for all advertisers regardless of topic, and by committing to the inclusion of advertisements for candidates on state and local levels.

We do have suggestions for potential improvements of the bill. First, we want to be sure that the proposed requirements, including in-ad disclosure language, are sufficiently flexible to account for character-constrained platforms like Twitter. Second, we hope that legislation on this topic would clarify that, while the duty to collect and display disclosure information lies with the platforms, the duty to provide accurate information lies with the advertisers.

5. What is Twitter’s current policy on the posting or promotion of hacked emails on your platform?

Twitter rules prohibit the distribution of hacked material that contains private information or trade secrets, or could put people in harm’s way. According to our Rules, Twitter does not

permit the use of our services to directly distribute content obtained through hacking that contains personally identifiable information, may put people in imminent harm or danger, or contains trade secrets. Direct distribution of hacked materials includes posting hacked content on Twitter (for instance, in the text of a Tweet, or in an image), or directly linking to hacked content hosted on other websites.

We also expanded the criteria for when we will take action on accounts which claim responsibility for a hack, which includes threats and public incentives to hack specific people and accounts. We also may suspend accounts in which Twitter is able to reliably attribute a hack to the account distributing that content. Commentary about a hack or hacked materials, such as news articles discussing a hack, are generally not considered a violation of this policy.

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

- **How is Twitter complying with GDPR?**
- **Are there protections that will flow to U.S. users as a result?**
- **What lessons should we be learning from the European experiment with data protection?**
- **Should we consider policy solutions like first-party consent?**
- **Why shouldn't companies be required to obtain explicit and informed consent before collecting or processing user data like in Europe?**

Twitter has undertaken a variety of internal and public facing updates to comply with the obligations imposed by the coming into force of the General Data Protection Regulation. This includes, for example, appointment of a Global Data Protection Officer, providing mechanisms to allow people to download their data from Twitter, mechanisms to allow people to contact Twitter's Office of Data Protection, and ensuring internal systems and processes exist to support Twitter's compliance with the GDPR.

The GDPR was developed over many years and thus, the underlying goals are commendable. Notably the GDPR's objectives of protecting consumers by providing for data protection that includes core tenants from the Federal Information Processing Standards developed in the 1970s in the U.S. Most of the internal and external facing updates Twitter has undertaken for compliance with the GDPR apply to all people who use Twitter's services irrespective of where they reside.

There are areas that should be examined carefully before considering adoption of the same regulations in the U.S. For example, the language around automated decision making may prove restrictive for business and innovation. Similarly, the GDPR's language around the commonly described right to be forgotten does not comport with the First Amendment.

Twitter believes that informed consent should be obtained for data processing. Twitter believes people should know and have control over the types of data that are received about them by data processors, how it is used, and when it is shared. However, Twitter does not believe that a blanket opt-in consent requirement should be imposed. This can lead to operational and technical difficulties. For example, to provide a person with a landing page to a service in their language, their IP address is processed to determine their approximate location to infer language. Required opt-in consent for such processing would make such useful features difficult to provide and result in friction for consumers. Thus, the type of consent mechanism used should be informed by the type of service, the type of data at issue, when in the use of the service the consent is being solicited, and the information and controls available to the consumer.

7. Do you think Twitter might benefit from more independent insight into anonymized activity?

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

Information sharing and collaboration are critical to Twitter's success in preventing hostile foreign actors from disrupting meaningful political conversations on the platform. The threat we face requires extensive partnership and collaboration with our government partners and industry peers. We each possess information the other does not have, and our combined information is more powerful in combating these threats together.

We recognize the value of inputs we receive from our industry peers about hostile foreign actors. We have shared and remain committed to sharing information across platforms to better understand and address the threat of hostile foreign interference with the electoral process. On August 24, 2018, Twitter hosted our industry peers to discuss data sharing about hostile foreign actors regarding 2018 election security. We continue to meet in regular cadence with our industry peers about election integrity efforts.

We also have well-established relationships with law enforcement agencies active in this arena, including the Federal Bureau of Investigation Foreign Influence Task Force and the Department of Homeland Security's Election Security Task Force. We look forward to continued cooperation with them on these issues, as only they have access to information critical to our joint efforts to stop bad faith actors.

Additionally, we committed to the United States Congress and the public to provide regular updates and information regarding our investigation into foreign interference in political conversations on Twitter. Since that time, we have shared examples of these types of content posted on Twitter by the Internet Research Agency (IRA) and provided the public with a direct notice if they interacted with these accounts. In August, we also disclosed details of another attempted influence campaign we identified as potentially located within Iran.

In line with our strong principles of transparency and with the goal of improving understanding of foreign influence and information campaigns, on October 17, 2018, Twitter released the full, comprehensive archives of the Tweets and media that are connected with these two previously disclosed and potentially state-backed operations on the service. We are making this data available with the goal of encouraging open research and investigation of these behaviors from researchers and academics around the world.

These large datasets included 3,841 accounts affiliated with the IRA, originating in Russia, and 770 other accounts, potentially originating in Iran. They include more than 10 million Tweets and more than 2 million images, GIFs, videos, and Periscope broadcasts, including the earliest on-Twitter activity from accounts connected with these campaigns, dating back to 2009.

8. The fact that Twitter failed to anticipate misuse is extremely troubling.

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

Twitter is committed to protecting the integrity of elections. We have made recent improvement to three critical areas of our election integrity efforts: (1) Updates to the Twitter Rules (2) Detection and Enforcement; and (3) Product Improvements.

We have updated the Twitter Rules to provide clearer guidance around several key issues, including fake account, attributed activity, and distribution of hacked materials. We have heard feedback that people think our rules about spam and fake accounts only cover common spam tactics like selling fake goods. As platform manipulation tactics continue to evolve, we are updating and expanding our rules to better reflect how we identify fake accounts, and what types of inauthentic activity violate our guidelines. We now may remove fake accounts engaged in a variety of emergent, malicious behaviors. Some of the factors that we will take into account

when determining whether an account is fake include the use of stock or stolen avatar photos, use of stolen or copied profile bios, and use of intentionally misleading profile information, including profile location

Additionally, as per the Twitter Rules, if we are able to reliably attribute an account on Twitter to an entity known to violate the Twitter Rules, we will take action on additional accounts associated with that entity. We are expanding our enforcement approach to include accounts that deliberately mimic or are intended to replace accounts we have previously suspended for violating our rules. Further, our rules prohibit the distribution of hacked material that contains private information or trade secrets, or could put people in harm's way. We are also expanding the criteria for when we will take action on accounts which claim responsibility for a hack, which includes threats and public incentives to hack specific people and accounts. Commentary about a hack or hacked materials, such as news articles discussing a hack, are generally not considered a violation of this policy.

We have seen positive results from our investments in conversational health and information integrity. We continue to enforce our rules against intentionally misleading election-related content. In August, we removed approximately 50 accounts misrepresenting themselves as members of various state Republican parties. We have also taken action on Tweets sharing media regarding elections and political issues with misleading or incorrect party affiliation information. We continue to partner closely with the RNC, DNC, and state election institutions to improve how we handle these issues. In August, we removed 770 accounts engaging in coordinated behavior which appeared to originate in Iran. Our investigation into this activity continues, and we will share further updates on our findings with law enforcement, our industry peers, and the public.

Our automated detections continue to identify and challenge millions of potentially spammy and automated accounts per week. In the first half of September, we challenged an average of 9.4 million accounts each week. As a result of our proactive detections and enforcements, we have continued to see a decline in the average number of spam-related reports we receive from individuals each day — from an average of approximately 17,000 per day in May, to approximately 16,000 per day in September. We are continuing to roll out improvements to our proactive enforcements against common policy violations, including building new proprietary systems to identify and remove ban evaders at speed and scale.

Finally, we continue to make improvements to the Twitter product to help people stay informed and to see the best content first. We heard feedback that people want an easy way to see the most recent Tweets in their home timeline. We recently updated the timeline personalization setting to allow people to select a strictly reverse-chronological experience,

without recommended content and recaps. This ensures you have more control of how you experience what's happening on our service. We are continuing to roll out new features to show people context about accounts on Twitter. In May, we launched an election labels beta for candidates in the 2018 U.S. midterm elections. We are also going to send candidates a message prompt to ensure they have two-factor authentication enabled on their account so it is safe and secure.

We are also offering electoral institutions increased support via an elections-specific support portal, which is designed to ensure we receive and review critical feedback about emerging issues as quickly as possible. We will continue to expand this program ahead of the elections and will provide information about the feedback we receive in the near future. As part of our civic engagement efforts, we are building conversation around the hashtag #BeAVoter with a custom emoji, sending U.S.-based individuals a prompt in their home timeline with information on how to register to vote, and drawing attention to these conversations and resources through the top US trend. This trend is being promoted by @TwitterGov, which will create even more access to voter registration information, including election reminders and an absentee ballot FAQ.

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

- **Do you believe that the Russian-linked operatives continue to utilize Twitter for information operations to undermine our democracy?**
- **Have you seen non-IRA, Russian-linked activity on your platform conducting similar types of information operations?**
- **What percentage of Russian-linked activity do you think the IRA represents?**
- **Have you seen evidence of additional Russian-linked troll farms?**
- **Have you identified any troll farms backed by countries other than Russia?**
- **Do you anticipate additional account take-downs in the weeks ahead?**
- **Will you commit to notifying the public should you identify other foreign influence operations?**
- **Will you alert users when they've been exposed to these types of operations?**

It is clear that information operations and coordinated inauthentic behavior will not cease. These types of tactics have been around for far longer than Twitter has existed — they will adapt and change as the geopolitical terrain evolves worldwide and as new technologies emerge. For our part, we are committed to understanding how bad-faith actors use our services. We will continue to proactively combat nefarious attempts to undermine the integrity of Twitter, while

partnering with civil society, government, our industry peers, and researchers to improve our collective understanding of coordinated attempts to interfere in the public conversation.

Our dedicated site integrity team, in partnership with a diverse range of committed organizations and personnel across the company, continue to invest heavily in this area. We are constantly seeking to improve our own ability to detect, understand, and neutralize these campaigns as quickly and robustly as technically possible. Twitter has learned from 2016 and more recently from other nation's elections how best to protect the integrity of our elections. Better tools, stronger policy, and new partnerships are already in place. We intend to understand the efficacy of these measures to continue to get better.

[From Senator Feinstein]

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

- **In the recent case of the Iranian-associated influence campaign, did an external company have to alert you to the activity; and if so, why?**
- **What specific steps are you taking to enhance your ability to find and mitigate influence operations?**

On August 21, 2018, working with our industry peers, Twitter suspended 770 accounts from Twitter for engaging in coordinated manipulation. Based on our analysis, it appears that many of these accounts originated from Iran. As with all investigations, we are committed to engaging with other companies and relevant law enforcement entities. Our goal is to assist investigations into these activities and where possible, we will provide the public with transparency and context on our efforts.

Fewer than 100 of the 770 suspended accounts claimed to be located in the U.S. and many of these were sharing divisive social commentary. On average, these 100 accounts Tweeted 867 times, were followed by 1,268 accounts, and were less than one year old. In line with our strong principles of transparency and with the goal of improving understanding of foreign influence and information campaigns, on October 21, 2018, we released the full, comprehensive archives of the Tweets and media that are connected with these two previously disclosed and potentially state-backed operations on our service. We are making this data available with the goal of encouraging open research and investigation of these behaviors from researchers and academics around the world.

Independent analysis of this activity by researchers is a key step toward promoting shared understanding of these threats. To support this effort, we have provided early access to a small group of researchers with specific expertise in these issues. Working with law enforcement and the authorities will always be our first priority, but we strongly believe that this level of transparency can enhance the health of the public conversation on the internet. This is our singular mission.

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

- **Do you believe there is a need for better information sharing between the social media companies?**
- **What is prohibiting your company from sharing more with your peers, government actors, and the public with respect to foreign information operations?**

Information sharing and collaboration are critical to Twitter’s success in preventing hostile foreign actors from disrupting meaningful political conversations on the platform. We recognize the value of inputs we receive from our industry peers about hostile foreign actors. We have shared and remain committed to sharing information across platforms to better understand and address the threat of hostile foreign interference with the electoral process. On August 24, 2018, Twitter hosted our industry peers to discuss data sharing about hostile foreign actors regarding 2018 election security. These conversations continue to occur with regular cadence in the lead-up to the 2018 midterm election.

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

- **Have your companies agreed on a common standard for what constitutes prohibited extremist or terrorist content? If not, why not?**
- **Would a shared standard and the deployment of similar software used to detect spam and copyrighted material, facilitate the automated blocking of such content across all four platforms?**
- **In the interest of transparency, would you make this database open to the public or researchers to know which images are prohibited?**

We agree that collaboration with our industry peers and civil society is critically important to addressing common threats and that it has been successful in meeting shared challenges. In June 2017, for example, we launched the Global Internet Forum to Counter Terrorism (the “GIFCT”), a partnership among Twitter, YouTube, Facebook, and Microsoft.

The GIFCT facilitates, among other things, information sharing; technical cooperation; and research collaboration, including with academic institutions. In September 2017, the members of the GIFTC announced a multimillion dollar commitment to support research on terrorist abuse of the Internet and how governments, tech companies, and civil society can respond effectively. We are looking to establish a network of experts that can develop these platform-agnostic research questions and analysis that consider a range of geopolitical contexts.

The GIFCT has created a shared industry database of “hashes”—unique digital “fingerprints”—for violent terrorist imagery or terrorist recruitment videos or images that have

been removed from our individual services. The database allows a company that discovers terrorist content on one of its sites to create a digital fingerprint and share it with the other companies in the forum, who can then use those hashes to identify such content on their services or platforms, review against their respective policies and individual rules, and remove matching content as appropriate, or even block extremist content before it is posted in the first place.

The database now contains more than 88,000 hashes. Instagram, Justpaste.it, LinkedIn, Oath, and Snap have also joined this initiative, and we are working to add several additional companies in 2018. Twitter also participates in the Technology Coalition, which shares images to counter child abuse. The database works to surface content for human review against each platform's respective terms of service. This is essential to take into account the context, for example academic or news media use.

Because each platform is unique, there are many elements of our coordinated work that do not translate easily across platforms. Although we share with other companies our approach to addressing shared threats, including certain signals that we use to identify malicious content, solutions applicable to the Twitter platform are not always applicable to other companies. We describe our tools as "in-house and proprietary" to distinguish them from tools that are developed by and licensed from third-party vendors.

13. Where Twitter has identified content as advancing a foreign influence campaign, will you commit to providing public access to a library of all ads that target users based on demographics (what content, purchased by whom, targeting whom)? If not, why not?

Twitter is committed to providing greater transparency to our account holders and the public, particularly as it relates to election integrity. In the future, we commit to releasing all the accounts and related content associated with potential information operations as appropriate. Following our investigation into the propaganda effort by the Internet Research Agency (IRA), we notified approximately 1.4 million individuals on our platform who interacted with this malicious content.

Twitter sent notices to people on the platform with an active email address who our records indicate are based in the U.S. and fall into at least one of the following categories:

- People who directly engaged during the election period with the 3,814 IRA-linked accounts we identified, either by Retweeting, quoting, replying to, mentioning, or liking those accounts or content created by those accounts;
- People who were actively following one of the identified IRA-linked accounts at the time those accounts were suspended; and

- People who opt out of receiving most email updates from Twitter and would not have received our initial notice based on their email settings.

[From Senator Wyden]

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

- **If yes, please provide a detailed accounting of each case, including the suspected foreign entity, and the number of users that saw or interacted with the content (e.g. clicked or shared).**
- **What steps has Twitter taken to inform its users, the public, and the United States Government of each case that you have listed in response to the previous question?**

In early 2018, we committed to the United States Congress and the public to provide regular updates and information regarding our investigation into foreign interference in political conversations on Twitter. Since that time, we have shared examples of these types of content posted on Twitter by the Internet Research Agency (IRA) and we notified approximately 1.4 million individuals on our platform who interacted with this malicious content. In August, we also disclosed details of another attempted influence campaign we identified as potentially located within Iran.

In line with our strong principles of transparency and with the goal of improving understanding of foreign influence and information campaigns, on October 21, 2018, we released the full, comprehensive archives of the Tweets and media that are connected with these two previously disclosed and potentially state-backed operations on our service. We made this data available with the goal of encouraging open research and investigation of these behaviors from researchers and academics around the world.

These large datasets comprise 3,841 accounts affiliated with the IRA, originating in Russia, and 770 other accounts, potentially originating in Iran. They include more than 10 million Tweets and more than 2 million images, GIFs, videos, and Periscope broadcasts, including the earliest on-Twitter activity from accounts connected with these campaigns, dating back to 2009.

Additionally, we have well-established relationships with law enforcement agencies active in this arena, including the Federal Bureau of Investigation Foreign Influence Task Force and the Department of Homeland Security’s Election Security Task Force. We look forward to continued cooperation with them on these issues, as only they have access to information critical to our joint efforts to stop bad faith actors.

15. In September 2017, Twitter confirmed that it had found approximately 200 accounts linked to the same Russian groups that had purchased ads on Facebook. In August 2018, Twitter confirmed that it had suspended 770 accounts for “coordinated manipulation.”

- **In addition to the cases listed above, has any foreign government, their agent, or an entity acting on the behalf of a foreign government, created Twitter accounts or written tweets, that masquerade as American for the purpose of influencing political debate or policymaking within the United States, not limited to elections?**
- **Has any other foreign entity, even if it is not known to be acting on behalf of a foreign government, created Twitter accounts, or written tweets, that masquerade as American for the purpose of influencing political debate or policymaking within the United States, not limited to elections?**
- **If the answer to either of the previous two questions is yes, please provide a detailed accounting of each case, including the foreign government (if applicable), the issue, and the number of users that saw or interacted with the content (e.g. clicked or shared).**
- **What steps has Twitter taken to inform users, the public, and the United States Government of any case listed in response to the previous question?**

Twitter is unaware of any instances, beyond the Russian-linked and Iran-affiliated accounts we have already disclosed publicly, of foreign government or entity acting on the behalf of a foreign government that have created Twitter accounts or written tweets, that masquerade as American for the purpose of influencing political debate or policymaking within the United States.

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

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

Twitter is committed to providing greater transparency to our account holders and the public, particularly as it relates to election integrity. Following our investigation into the propaganda effort by the Internet Research Agency (IRA), we notified approximately 1.4 million individuals on our platform who interacted with this malicious content.

Twitter sent notices to people on the platform with an active email address who our records indicate are based in the U.S. and fall into at least one of the following categories:

- People who directly engaged during the election period with the 3,814 IRA-linked accounts we identified, either by Retweeting, quoting, replying to, mentioning, or liking those accounts or content created by those accounts;
- People who were actively following one of the identified IRA-linked accounts at the time those accounts were suspended; and
- People who opt out of receiving most email updates from Twitter and would not have received our initial notice based on their email settings.

Twitter does not have data on the the number of accounts with a two-factor authentication key that interacted with with the IRA, although in this instance the security of the accounts that interacted with the IRA were not compromised.

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

- **Prior to being notified by Twitter, how many of these accounts had some form of two-factor authentication enabled on their accounts?**
- **Prior to being notified by Twitter, how many of these accounts were secured with a two-factor authentication security key?**

We will provide additional information to the Committee concerning the targeting of elected officials or their staff in the United States in a more secure setting.

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

- **How many of these accounts had some form of two-factor authentication enabled on their accounts.**
- **How many of these accounts were secured with a two-factor authentication security key?**

Twitter recommends to the individuals on its platform certain best security practices in order to help keep their accounts secure. These include the use of a strong password that is not reused on other websites, the use login verification, and requiring email and phone number to request a reset password link or code. Twitter also suggests that individuals on the platform be cautious of suspicious links and always make sure an individualis on twitter.com before he or she

enters login information. We caution people to never give their username and password out to third parties, especially those promising to grow followers, make money, or verify an account. A relatively small number of people using Twitter within the United States have two-factor authentication enabled. Since May 2018, Twitter estimates that approximately 3 million accounts may have potentially been impacted by data breaches, although there is no indication these have been associated with foreign government activity.

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

- **How many of these accounts had some form of two-factor authentication enabled on their accounts?**
- **How many of these accounts were secured with a two-factor authentication security key?**

Please see the response to question 18.

20. Twitter has a tool called Tailored Audiences, which is similar to Facebook's Custom Audiences and allows advertisers to upload lists of specific users to target them with ads. Twitter's policies state that advertisers who use Tailored Audiences must obtain consent from users about the data they have acquired.

- **How does Twitter ensure that all of its advertisers actually follow those policies?**
- **Is Twitter aware of any advertisements targeted with Tailored Audiences that appear to be designed to discourage any United States citizen from voting?**
 - **If yes, please provide a full accounting of each case, including the advertisement, what Twitter knows about the party that purchased the advertising, and the number of users that saw or interacted with the content (e.g. clicked).**
 - **If the answer to the question above is yes, were these voter discouragement ads targeted at people of any particular race or ethnic group?**
 - **Were these voter discouragement ads predominantly targeted at people expected to vote for one party or the other?**

Tailored audiences is a product feature that advertisers use to target existing people on the platform and customers. For example, Advertisers can reach existing customers by uploading a list of email addresses. Advertisers can also target those that have recently visited the advertisers' websites or reach those that have taken specific action in an application, such as installation or registration.

Twitter informs individuals about Tailored Audiences in several ways. For example, Twitter describes this activity in its Privacy Policy, an “Ads info” footer on twitter.com, and the “Why am I seeing this ad?” section of the drop down menu on Twitter ads themselves. Each of these locations describe interest-based advertising on Twitter and explain how to use the associated privacy controls. In addition, the “Your Twitter Data” tool allows people on the platform to download a list of advertisers that have included them in a Tailored Audience.

If people do not want Twitter to show them Tailored Audience ads on and off of Twitter, there are several ways they can turn off this feature: using their Twitter settings, they can visit the Personalization and data settings and adjust the Personalize ads setting; if they are on the web, they can visit the Digital Advertising Alliance’s consumer choice tool at optout.aboutads.info to opt out of seeing interest-based advertising from Twitter in their current browser; if they do not want Twitter to show them interest-based ads in Twitter for iOS on their current mobile device, they can enable the “Limit Ad Tracking” setting in their iOS phone’s settings; and if they do not want Twitter to show them interest-based ads in Twitter for Android on their current mobile device, they can enable “Opt out of Ads Personalization” in an Android phone’s settings.

In addition to explaining Tailored Audiences to people on the platform, offering them several ways to disable the feature, and enabling them to view the advertisers who have included them in Tailored Audiences, as described above, the Tailored Audience legal terms require that advertisers have secured all necessary rights, consents, waivers, and licenses for use of data. Advertisers are also required to provide all individuals from whom the data is collected with legally-sufficient notice that fully discloses the collection, use, and sharing of the data that is provided to Twitter for purposes of serving ads targeted to people’s interest, and legally sufficient instructions on how they can opt out of interest-based advertising on Twitter.

Twitter is not aware of any advertisements targeted with Tailored Audiences that appear to be designed to discourage any United States citizen from voting.

21. Has any foreign government, their agent, or other foreign entity ever used Tailored Audiences to target individuals in the United States?

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

In 2017, the U.S. intelligence community named Russia Today (RT) and Sputnik as implementing state-sponsored Russian efforts to interfere with and disrupt the 2016 U.S.

Presidential election. We made the policy decision to off-board advertising from all accounts owned by RT and Sputnik based on the retrospective work we had conducted around the 2016 U.S. election and the U.S. intelligence conclusion that both RT and Sputnik attempted to interfere with the election on behalf of the Russian government.

In 2014, @RT_com used Tailored Audiences to deliver advertisements totaling \$8,487 and @RTUKnews used Tailored Audiences to deliver advertisements totaling \$165 in 2015.

Based on our internal investigation of their behavior as well as their inclusion in the January 2017 intelligence community report, Twitter decided to take the \$1.9 million we were projected to have earned from RT global advertising since they became an advertiser in 2011, which includes the \$274,100 in 2016 U.S.-based advertising, and donated those funds to support external research into the use of malicious automation and misinformation, with an initial focus on elections and automation.

22. Does Twitter have a policy of shutting down accounts that seek to suppress voting?

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

Twitter unequivocally condemns the use of our platform for any election interference activity. Tweets aimed at suppressing voter turnout generally are surfaced through reports from people using our service. This content is reviewed, then promptly removed as illegal interference with voting rights: the content is either restricted as inaccessible pending deletion by the individuals (i.e., other individuals on the platform are unable to see the content) or the responsible accounts are permanently suspended. In addition, in order to proactively surface additional Tweets with a given text-to-vote meme, Twitter utilizes technology for identifying instances where the same image appears across multiple Tweets. Content identified through this process is then subject to manual review.

Depending on the number of violations for any given account disseminating voter suppression Tweets, Twitter will either restrict access to the Tweet or suspend the account. During the period leading up to the 2016 election, for example, Twitter labeled and restricted access to the vote-to-text Tweets pursuant to the Twitter User Agreement, which contains the Twitter Terms of Service, Twitter Privacy Policy, and Twitter Rules. According to the unlawful use provision of the Twitter Rules, people are prohibited from using Twitter's "service for any unlawful purpose or in furtherance of illegal activities" and "[i]nternational users agree to comply with all local laws regarding online conduct and acceptable content."

Additionally, we have updated the Twitter Rules to provide clearer guidance around several key issues, including fake account, attributed activity, and distribution of hacked materials. We have heard feedback that people think our rules about spam and fake accounts only cover common spam tactics like selling fake goods. As platform manipulation tactics continue to evolve, we are updating and expanding our rules to better reflect how we identify fake accounts, and what types of inauthentic activity violate our guidelines. We now may remove fake accounts engaged in a variety of emergent, malicious behaviors. Some of the factors that we will take into account when determining whether an account is fake include the use of stock or stolen avatar photos, use of stolen or copied profile bios, and use of intentionally misleading profile information, including profile location

As per the Twitter Rules, if we are able to reliably attribute an account on Twitter to an entity known to violate the Twitter Rules, we will take action on additional accounts associated with that entity. We are expanding our enforcement approach to include accounts that deliberately mimic or are intended to replace accounts we have previously suspended for violating our rules. Further, our rules prohibit the distribution of hacked material that contains private information or trade secrets, or could put people in harm's way. We are also expanding the criteria for when we will take action on accounts which claim responsibility for a hack, which includes threats and public incentives to hack specific people and accounts. Commentary about a hack or hacked materials, such as news articles discussing a hack, are generally not considered a violation of this policy.

23. Has the Internet Research Agency ever used Tailored Audiences to target users, in the United States or elsewhere, with advertisements?

Twitter is not aware of any efforts by the Internet Research Agency to use Tailored Audiences to target individuals, in the United States or elsewhere, with advertisements.

24. Does Twitter believe that any of the content created by the Russian Internet Research Agency was designed to discourage anyone from voting?

Twitter did not see that content created by the Russian Internet Research Agency that constituted voter suppression.

25. Can users opt out of being targeted with Tailored Audiences?

- **If no, why not?**

Yes. If individuals do not want Twitter to show them Tailored Audience ads on and off of Twitter, there are several ways they can turn off this feature: using their Twitter settings, they

can visit the Personalization and data settings and adjust the Personalize ads setting; if they are on the web, they can visit the Digital Advertising Alliance’s consumer choice tool at optout.aboutads.info to opt out of seeing interest-based advertising from Twitter in their current browser; if they do not want Twitter to show them interest-based ads in Twitter for iOS on their current mobile device, they can enable the “Limit Ad Tracking” setting in your iOS phone’s settings; and if they do not want Twitter to show them interest-based ads in Twitter for Android on their current mobile device, they can enable “Opt out of Ads Personalization” in your Android phone’s settings.

26. For several years, Twitter has allowed its customers to protect their accounts from hacking through the use of two-factor authentication. Since June, Twitter has also supported the use of a physical security token as an enhanced form of two-factor authentication. However, two-factor authentication remains an opt-in feature for Twitter users.

- **Does Twitter require that its employees use two-factor authentication for their work accounts?**
- **If yes, does Twitter require, like Google, that employees use a security key?**

We will provide additional information to the Committee concerning our security protocols for employees in a more secure setting.

27. Do you have two-factor authentication enabled for your personal Twitter and personal email accounts?

- **If yes, are you using a security key?**

Yes, two-factor authentication is enabled and we will provide additional information to the Committee concerning our security protocols in a more secure setting.

28. What percentage of Twitter’s U.S. customers have enabled any form of two-factor authentication?

Twitter recommends to the individuals on its platform certain best security practices in order to help keep their accounts secure. These include the use of a strong password that is not reused on other websites, the use login verification, and requiring email and phone number to request a reset password link or code. Twitter also suggests that individuals on the platform be cautious of suspicious links and always make sure an individual is on twitter.com before he or she enters login information. We caution people to never give their username and password out to third parties, especially those promising to grow followers, make money, or verify an account.

We will provide additional information to the Committee concerning the use of two-factor authentication in the U.S. in a more secure setting.

29. What percentage of Twitter's U.S. customers have enabled enhanced two-factor authentication using a security key?

We will provide additional information to the Committee concerning the use of two-factor authentication using a security key in the U.S. in a more secure setting.

30. Since May, Twitter now specially identifies the accounts of individuals running for public office.

- **What percentage of these Twitter accounts currently have any form of two-factor authentication enabled?**
- **What percentage are using a security key?**
- **What specific outreach, if any, has Twitter engaged in to encourage elected officials and individuals running for public office to enable two-factor authentication on their official and personal Twitter accounts?**

Twitter recommends to the individuals on its platform certain best security practices in order to help keep their accounts secure, including those individuals who are running for public office. Twitter has developed a new U.S. election label to identify political candidates. The label includes information about the office the candidate is running for, the state the office is located in, and the district number, if applicable. Accounts of candidates who have qualified for the general election and who are running for governor or for the U.S. Senate or House of Representatives will display an icon of a government building. These new features are designed to instill confidence that the content people are viewing is reliable and accurately reflects candidates' and elected officials' positions and opinions.

In our correspondence with candidates participating in the election label program, Twitter encouraged them to review all of our security best practices. We stated: "As a friendly reminder, we highly recommend that you review our account safety and security best practices" and included a link to the relevant content. We have also distributed detailed information that includes security best practices to the political parties and other election stakeholders to encourage them learn about all of our integrity efforts.

We will provide additional information to the Committee concerning the use of two-factor authentication of badged candidates in the U.S. and their use of security keys in a more secure setting.

31. Twitter will place a blue verification badge on accounts “of public interest” which have been verified as authentic by Twitter.

- **Does Twitter currently require that verified accounts enable two-factor authentication?**

No, two-factor authentication is not required.

[From Senator Heinrich]

32. In July 2018, Twitter acknowledged it has a problem with fake and automated accounts, or bots, and announced that in the final three months of 2017, the company had suspended 58 million accounts, another 70 million in May and June, and continuing at a rate of a million per day.

- **How has Twitter improved detection of automated accounts? Has this been a technical challenge?**

Twitter continues to develop the detection tools and systems needed to combat malicious automation on our platform. Twitter has refined its detection systems. Twitter prioritizes identifying suspicious account activity, such as exceptionally high-volume Tweeting with the same hashtag or mentioning the same @handle without a reply from the account being addressed, and requires an individual using the platform to confirm control. Twitter has also increased its use of challenges intended to catch automated accounts, such as reCAPTCHAs, that require people to identify portions of an image or type in words displayed on screen, and password reset requests that protect potentially compromised accounts. Twitter is also in the process of implementing mandatory email or cell phone verification for all new accounts. We will continue to undertake important steps to improve detection of automated accounts in the coming months and years.

33. According to the New York Times, Twitter’s “purge” of fake and automated accounts resulted in over 300,000 followers lost for President Trump’s Twitter account, about .58% of his total.

- **Following the purge, does Twitter estimate that any fake or automated accounts still follow President Trump on Twitter?**
- **How accurately can Twitter or other Twitter audit sites estimate the number of real and fake Twitter followers for any particular account?**

Twitter continues to develop the detection tools and systems needed to combat malicious automation on our platform. Twitter has refined its detection systems. Twitter prioritizes identifying suspicious account activity, such as exceptionally high-volume Tweeting with the same hashtag or mentioning the same @handle without a reply from the account being addressed, and requires an individual using the platform to confirm control. Twitter has also increased its use of challenges intended to catch automated accounts, such as reCAPTCHAs, that require people to identify portions of an image or type in words displayed on screen, and password reset requests that protect potentially compromised accounts. Twitter is also in the process of implementing mandatory email or cell phone verification for all new accounts.

Our efforts have been effective. Due to technology and process improvements, we are now removing 214% more accounts year-over-year for violating our platform manipulation policies. For example, over the course of the last several months, our systems identified and challenged between 8.5 million and 10 million accounts each week suspected of misusing automation or producing spam. Spam can be generally described as unsolicited, repeated actions that negatively impact other people. This includes many forms of automated account interactions and behaviors as well as attempts to mislead or deceive people. This constitutes more than three times the 3.2 million we were catching in September 2017. We thwart 530,000 suspicious logins a day, approximately double the amount of logins that we detected a year ago.

These technological improvements have brought about a corresponding reduction in the number of spam reports from people on Twitter, evidence to us that our systems' ability to automatically detect more malicious accounts and potential bad faith actors than they did in the past. We received approximately 25,000 such reports per day in March of this year; that number decreased to 17,000 in August.

We also removed locked accounts from people's follower counts, including President Trump, to ensure these figures are more reliable. Accounts are locked when our systems detect unusual activity and force a password change or other challenge. If the challenge has not been met or the password has not been changed within a month, the account is locked, barring it from sending Tweets, Retweets or liking posts from others.

34. In reply to a question about whether Twitter users should be notified whether they are communicating with a human or a machine, you testified that "we can identify these automations, we can label them, and I think that is useful context and it's an idea that we have been considering over the past few months. It's really a question of the implementation, but we are interested in it and we are going to do something along those lines." You also noted that Twitter is looking at "expanding that transparency report around suspensions of any account."

- **Do you have any more information on what a transparency report on numbers of and suspension of automated account activity might look like or when Twitter might begin issuing such reports?**
- **What are the challenges in distilling such information?**

Twitter is committed to the open exchange of information. First published on July 2, 2012, our biannual Twitter Transparency Report highlights trends in legal requests, intellectual property-related requests, and email privacy best practices. The report also provides insight into whether or not we take action on these requests. The Transparency Report includes information requests from worldwide government and non-government legal requests we have received for

account information. Removal requests are also included in the Transparency Report and include worldwide legal demands from governments and other authorized reporters, as well as reports based on local laws from trusted reporters and non-governmental organizations, to remove or withhold content.

The Transparency Report also discloses information on third-party requests that compel Twitter to remove content for legal reasons (“legal requests”) under our Country Withheld Content (“CWC”) policy. Governments (including law enforcement agencies), organizations chartered to combat discrimination, and lawyers representing individuals are among the many complainants that submit legal requests included below. For example, we may receive a court order requiring the removal of defamatory statements in a particular jurisdiction, or law enforcement may ask us to remove prohibited content such as Nazi symbols in Germany.

In December 2017, Twitter updated its in-product messaging about withheld content to better explain why content has been withheld. Subsequently, we began to differentiate between legal demands (e.g., court orders) and reports based on local law(s) (e.g., reports alleging the illegality of particular content in a certain country). To further increase transparency, this change is also reflected in the report below.

The Transparency Report also includes information on government requests to remove content that may violate Twitter’s Terms of Service (TOS) under the following Twitter Rules categories: abusive behavior, copyright, promotion of terrorism, and trademark. It does not include legal demands, regardless of whether they resulted in a TOS violation, which will continue to be published in our removal request section report. As we take an objective approach to processing global TOS reports, the fact that the reporters in these cases happened to be government officials had no bearing on whether any action was taken under our Rules.

The Transparency Report also includes the total number of Digital Millennium Copyright Act (DMCA) takedown notices and counter notices received for Twitter and Periscope content, along with data about the top five copyright reporters across both platforms. The Vine app was transitioned in January of 2017. Trademark notices include reports of alleged Trademark Policy violations received for Twitter and Periscope.

The forthcoming Transparency Report will also include information on automated manipulation.

35. Twitter estimates that fewer than 8.5 percent of its users use automation tools, yet it has recently announced the suspension of millions of accounts, which calls that estimate into question.

- **What is Twitter's latest estimate of numbers of its accounts that are automated?**
- **What is Twitter's estimate of numbers of automated accounts that are used maliciously, as opposed to for positive purposes?**
- **Why is it difficult to provide these kinds of estimates?**

Twitter continues to develop the detection tools and systems needed to combat malicious automation on our platform. Twitter has refined its detection systems. Twitter prioritizes identifying suspicious account activity, such as exceptionally high-volume Tweeting with the same hashtag or mentioning the same @handle without a reply from the account being addressed, and requires an individual using the platform to confirm control. Twitter has also increased its use of challenges intended to catch automated accounts, such as reCAPTCHAs, that require people to identify portions of an image or type in words displayed on screen, and password reset requests that protect potentially compromised accounts. Twitter is also in the process of implementing mandatory email or cell phone verification for all new accounts.

Our efforts have been effective. Due to technology and process improvements, we are now removing 214% more accounts year-over-year for violating our platform manipulation policies. For example, over the course of the last several months, our systems identified and challenged between 8.5 million and 10 million accounts each week suspected of misusing automation or producing spam. Spam can be generally described as unsolicited, repeated actions that negatively impact other people. This includes many forms of automated account interactions and behaviors as well as attempts to mislead or deceive people. This constitutes more than three times the 3.2 million we were catching in September 2017. We thwart 530,000 suspicious logins a day, approximately double the amount of logins that we detected a year ago.

These technological improvements have brought about a corresponding reduction in the number of spam reports from people on Twitter, evidence to us that our systems' ability to automatically detect more malicious accounts and potential bad faith actors than they did in the past. We received approximately 25,000 such reports per day in March of this year; that number decreased to 17,000 in August.

We also removed locked accounts from people's follower counts, to ensure these figures are more reliable. Accounts are locked when our systems detect unusual activity and force a password change or other challenge. If the challenge has not been met or the password has not been changed within a month, the account is locked, barring it from sending Tweets, Retweets or liking posts from others.

36. Twitter has disputed estimates by outside researchers that up to 15 percent of its accounts are bots rather than real people.

- **Is Twitter collaborating with academics and the research community in order to better quantify the extent of its bot problem?**

Information sharing and collaboration are critical to Twitter's success in preventing hostile foreign actors from disrupting meaningful political conversations on the platform. The threat we face requires extensive partnership and collaboration with our government partners and industry peers. We each possess information the other does not have, and our combined information is more powerful in combating these threats together.

We recognize the value of inputs we receive from our industry peers about hostile foreign actors. We have shared and remain committed to sharing information across platforms to better understand and address the threat of hostile foreign interference with the electoral process. On August 24, 2018, Twitter hosted our industry peers to discuss data sharing about hostile foreign actors regarding 2018 election security. We continue to meet in regular cadence with our industry peers about election integrity efforts.

We also have well-established relationships with law enforcement agencies active in this arena, including the Federal Bureau of Investigation Foreign Influence Task Force and the Department of Homeland Security's Election Security Task Force. We look forward to continued cooperation with them on these issues, as only they have access to information critical to our joint efforts to stop bad faith actors.

Additionally, we committed to the United States Congress and the public to provide regular updates and information regarding our investigation into foreign interference in political conversations on Twitter. Since that time, we have shared examples of these types of content posted on Twitter by the Internet Research Agency (IRA) and provided the public with a direct notice if they interacted with these accounts. In August, we also disclosed details of another attempted influence campaign we identified as potentially located within Iran.

In line with our strong principles of transparency and with the goal of improving understanding of foreign influence and information campaigns, on October 17, 2018, Twitter released the full, comprehensive archives of the Tweets and media that are connected with these two previously disclosed and potentially state-backed operations on the service. We are making this data available with the goal of encouraging open research and investigation of these behaviors from researchers and academics around the world.

These large datasets consist of 3,841 accounts affiliated with the IRA, originating in Russia, and 770 other accounts, potentially originating in Iran. They include more than 10 million Tweets and more than 2 million images, GIFs, videos, and Periscope broadcasts,

including the earliest on-Twitter activity from accounts connected with these campaigns, dating back to 2009.

[From Senator Lankford]

37. How do you verify that third-parties with access to Twitter’s data do not violate the company’s terms of use?

We recognize that access to that data could be manipulated, so we have taken steps to prevent the use of our application programming interfaces (“APIs”) for products and services that are abusive or that disrupt the health of conversations. Those to whom we grant access to our APIs are prohibited from using the data to manipulate conversations or otherwise abuse the data. Between April and June 2018 alone we removed more than 143,000 applications that we determined to be in violation of our developer policies. Most violated our policies against producing spam via APIs. And we continue to invest in and improve our detection tools to stop misuse of public Twitter data.

In July 2018, we introduced a new measure designed to increase developers’ accountability for applications that create and engage with Twitter content and accounts. Twitter now reviews and conducts compliance checks of all developers’ stated use of the data that they wish to access. We have also added new protections aimed to prevent the registration of low quality and spam-generating applications. We believe that these additional steps will help protect the integrity of our platform.

38. What are Twitter’s platform threat detection capabilities?

- **What are the limitations of Twitter’s ability to detect threats to the platform?**
- **How much of Twitter’s platform threat detection is outsourced?**

Twitter has created an internal cross-functional analytical team whose mission is to monitor site and platform integrity. Drawing on expertise across the company, the analytical team can respond immediately to escalations of inauthentic, malicious automated or human-coordinated activity on the platform. The team’s work enables us to better understand the nature of the malicious activity and mitigate it more quickly.

To supplement its own analyses, Twitter’s analytical team also receives and responds to reports from across the company and from external third parties. The results from all of the team’s analyses are shared with key stakeholders at Twitter and provide the basis for policy changes and product initiatives and removal of accounts.

The primary focus of the cross-functional analytical team is election readiness. Leading up to and during the 2018 election period, the team will examine, respond to, and escalate instances of suspected inauthentic, election-related coordinated activity in political conversation and conduct in-depth analyses of relevant Twitter data.

[From Senator Harris]

39. In his November 2017 testimony to the Committee, Twitter’s general counsel stated that “false or spam accounts represent less than 5% of our [Monthly Active Users] (MAU). On July 7, 2018, the Washington Post reported that Twitter suspended over 70 million accounts deemed fake or suspicious in May and June. Additionally, on July 17th, the Associated Press reported Twitter suspended 58 million accounts in the final three months of 2017.

- How many accounts has Twitter suspended, in total, since November of 2016?
- What percentage of Twitter’s total registered users does that represent?
- What percentage of your active users does that represent?
- How many of the suspended accounts claimed a location in the United States but had technical access that suggested a foreign location?
- How many of the suspended accounts connected to Twitter from an IP address in a foreign country?
- How many of the suspended accounts used a Virtual Private Network (VPN)?
- How many of the suspended accounts were automated?
- How many of the suspended automated accounts used Twitter’s application program interface (API)?
- How many malicious automated accounts used “headless” browsers, i.e., browsers without a visual user interface, or other methods of device impersonation?
 - What steps have you taken to detect such activity?

Twitter continues to develop the detection tools and systems needed to combat malicious automation on our platform. Twitter has refined its detection systems. Twitter prioritizes identifying suspicious account activity, such as exceptionally high-volume Tweeting with the same hashtag or mentioning the same @handle without a reply from the account being addressed, and requires an individual using the platform to confirm control. Twitter has also increased its use of challenges intended to catch automated accounts, such as reCAPTCHAs, that require people to identify portions of an image or type in words displayed on screen, and password reset requests that protect potentially compromised accounts. Twitter is also in the process of implementing mandatory email or cell phone verification for all new accounts.

Our efforts have been effective. Due to technology and process improvements, we are now removing 214% more accounts year-over-year for violating our platform manipulation policies. For example, over the course of the last several months, our systems identified and challenged between 8.5 million and 10 million accounts each week suspected of misusing automation or producing spam. Spam can be generally described as unsolicited, repeated actions that negatively impact other people. This includes many forms of automated account interactions and behaviors as well as attempts to mislead or deceive people. This constitutes more than three

times the 3.2 million we were catching in September 2017. We thwart 530,000 suspicious logins a day, approximately double the amount of logins that we detected a year ago.

These technological improvements have brought about a corresponding reduction in the number of spam reports from people on Twitter, evidence to us that our systems' ability to automatically detect more malicious accounts and potential bad faith actors than they did in the past. We received approximately 25,000 such reports per day in March of this year; that number decreased to 17,000 in August.

We also removed locked accounts from people's follower counts, to ensure these figures are more reliable. Accounts are locked when our systems detect unusual activity and force a password change or other challenge. If the challenge has not been met or the password has not been changed within a month, the account is locked, barring it from sending Tweets, Retweets or liking posts from others.

40. Do you stand by previous estimates of false or spam accounts, including for previous quarters in your SEC filings?

Yes.

41. Do you intend to revise or update testimony previously provided to the Committee concerning Twitter's estimates of the proportion of MAU comprising false or spam accounts?

No.

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

- **Are deepfakes a violation of Twitter's terms of use?**
- **What is Twitter doing to identify deepfakes on its platform and to alert users when they may be seeing deepfakes?**
- **How many deepfakes has Twitter identified on its platform to date?**

Twitter is aware of deepfakes in the context of intimate media on the platform. Deepfakes in the context of intimate media are clear violations of our terms of services and our intimate media policy. Twitter suspends any account we identify as the original poster of intimate media that has been produced or distributed without the subject's consent. We also suspend any account dedicated to posting this type of content.

43. Can Twitter commit to:

- **Assessing how foreign disinformation campaigns can use deepfakes;**
- **Developing a strategy to combat it; and,**
- **Reporting its findings and efforts to the Committee by the end of the year?**

The public conversation occurring on Twitter is never more important than during elections, the cornerstone of our democracy. Our service shows the world what is happening, democratizes access to information and—at its best—provides people insights into a diversity of perspectives on critical issues; all in real-time. We work with commitment and passion to do right by the people who use Twitter and the broader public. Any attempts to undermine the integrity of our service is antithetical to our fundamental rights and undermines the core tenets of freedom of expression, the value upon which our company is based. This issue affects all of us and is one that we care deeply about as individuals, both inside and outside the company.

We appreciate the continued partnership with the Committee, and we share your concern about malicious foreign efforts to manipulate and divide people in the United States and throughout the world, including through the use of foreign disinformation campaigns that rely upon the use of deepfakes. We will continue to share our ongoing work to safeguard elections with the members of this Committee.