

Questions from Senator Tillis

1. My subcommittee has spent this year holding a series of hearings on the Digital Millennium Copyright Act (DMCA) and the safe harbors that internet service providers enjoy when it comes to copyright infringement by their users. Technology company witnesses said they can't identify infringements by their users – but Twitter is very good at identifying and censoring or taking down or flagging content from conservative voices. Don't you think that if Twitter can track and take down political content that it also has the ability to take down content that infringes the IP of hard-working American creators?

Twitter responds to copyright complaints submitted under the DMCA pursuant to our [copyright policy](#). Section 512 of the DMCA outlines the statutory requirements necessary for formally reporting copyright infringement, as well as providing instructions on how an affected party can appeal a removal by submitting a compliant counter-notice. Twitter will respond to reports of alleged copyright infringement, such as allegations concerning the unauthorized use of a copyrighted image as a profile or header photo, allegations concerning the unauthorized use of a copyrighted video or image uploaded through our media hosting services, or Tweets containing links to allegedly infringing materials.

Twitter is unique among our industry peers and remains a text-first service for the creation of content. When media is shared on Twitter, it is often in service of commentary or criticism to drive real-time, public conversations and debate. For example, when our customers post snippets of third-party content (videos, photos and GIFs), they are frequently doing so in an exercise of political and newsworthy speech, and for the purpose of commentary and criticism (not to share and watch pirated content). As you can see in our [transparency reports](#), the number of takedown notices Twitter receives annually is a fraction of what other platforms receive and reflects the fact that allegedly infringing material is a small fraction of the total Tweets uploaded to Twitter.

2. Can you tell me about the number of human moderators you have at Twitter, and how many are working on copyright infringement?

Twitter uses a combination of machine learning and human review to adjudicate reports of violations of the Twitter Rules and make determinations on whether the activity violates our rules. With hundreds of millions of Tweets sent every day, we recognize that engaging in content moderation at scale requires increased use of machine learning and algorithms to surface and take action against violations of our rules. Thus, we continue to invest in efforts to automate content moderation. As a result of these investments, more than 50 percent of Tweets we take action on for abuse are now proactively surfaced using technology, rather than relying on reports to Twitter.

3. How does Twitter determine that it is appropriate to take a proactive step to counter, address, block, or remove content? Are the considerations different for political content and copyright-infringing content?

Whether Twitter takes enforcement action is dependent on whether an individual is in violation of the Twitter Rules. The Twitter Rules include a range of policies, including policies that prohibit abuse, harassment, hateful conduct, and child sexual exploitation, among many others. The Twitter Rules also include a copyright policy. Under this policy, Twitter will respond to reports of alleged

copyright infringement, such as allegations concerning the unauthorized use of a copyrighted image as a profile or header photo, allegations concerning the unauthorized use of a copyrighted video or image uploaded through our media hosting services, or Tweets containing links to allegedly infringing materials.

4. I know that Twitter has developed significant technological tools to guard against foreign interference in elections or to identify and label content that you deem politically misleading, inaccurate, or dispute. What types of technological tools do you use to identify and address content that infringes copyright?

Twitter relies on rightsholders, who are in the best position to know whether use of their content is infringing or not, to notify us of infringing material on Twitter. Twitter's response to copyright complaints may include the removal or restriction of access to allegedly infringing material. If we remove or restrict access to content in response to a copyright complaint, Twitter will make a good-faith effort to contact the affected account holder with information concerning the removal or restriction of access, including a full copy of the complaint, along with instructions for filing a counter-notice.

In an effort to be as transparent as possible regarding the removal or restriction of access to Tweeted content, we clearly mark withheld Tweets and media to indicate to viewers when content has been withheld in response to a complaint from a rightsholder. We also send a redacted copy of each copyright complaint that we process to [Lumen](#), where they are posted to a public-facing website, with all personal information removed.

One important consideration as we contemplate voluntary agreements and technological measures to address this problem is that in certain circumstances heightened specificity can hurt small businesses that are launching new apps and platforms. Small companies will not necessarily have the resources to afford expensive third-party technical solutions. Additionally, we are always concerned that voluntary agreements will further entrench the market dominant players.

5. Can you give me an update on Twitter's ongoing issues with RIAA and the steps you are taken to be more proactive in taking down the content of musicians and other artists that is being pirated on your site?

Twitter responds to all legitimate copyright complaints as laid out in Section 512 of the DMCA. Furthermore, Twitter does not allow for full-length music streaming as some of our competitors do. A full accounting of actions taken to protect copyrighted material on our platform can be found in our [Transparency Report](#). For 2019, the total copyright takedown notices we received worldwide from all rightsholders was less than 285,000; those notices resulted in 1.1 million pieces of media removed in 2019 worldwide.

We are actively engaged with the RIAA and many other rightsholder bodies regarding copyright concerns they have. We dedicate significant resources to quickly respond to takedown notices, and we work with rightsholders and representative bodies to address specific concerns if and when they arise.

6. We're now continuing the DMCA reform process with draft legislation that I want stakeholders from all perspectives to help me refine so we can build consensus. Do I have your commitment that Twitter will participate constructively and in good faith?

Yes, we look forward to engaging with you on constructive dialogue on this issue.

7. It is not uncommon for legitimate media outlets to make corrections or update stories that they publish and post on Twitter. Does your platform place a notification on posts for stories that have been corrected? If not, why not?

We have heard from the people who use Twitter that we should not determine the truthfulness of Tweets and we should provide context to help people make up their own minds in cases where the substance of a Tweet is disputed. Consistent with this feedback from our customers, we have expanded our enforcement options to allow us to label misinformation. When we label Tweets, we link to Twitter conversation that shows three things for context: (1) factual statements; (2) counterpoint opinions and perspectives; and (3) ongoing public conversation around the issue.

For now, we have focused these labeling efforts on areas where we believe there is the greatest risk of harm, including COVID-19, civic integrity, and manipulated media. However, we continue to explore additional ways in which we can provide additional context to people who use Twitter, in order to address potential harms associated with misinformation.

8. You claim you are not a publisher. At the same time you substitute your judgment for what can be published on your website for the editorial decision of traditional news outlets. The main example of this is the recent censoring of the New York Post story on Hunter Biden. How are you not engaging in editorial decision making by prohibiting a traditional news outlet from publishing a story on your site?

We issued the Distribution of Hacked Materials Policy in advance of the U.S. 2018 midterm elections to discourage and mitigate harms associated with hacks and unauthorized exposure of private information. Pursuant to these policies, on October 14, 2020, we took action on Tweets related to two articles published by the *New York Post* that, based on preliminary information, linked to content we determined to be in violation of our policies, including the Distribution of Hacked Materials Policy. Following our enforcement actions, we received significant feedback — both positive and negative — on the action.

After reviewing the feedback, we made changes within 24 hours to the policy to address concerns that there could be unintended consequences to journalists, whistleblowers, and others in ways that are contrary to Twitter's purpose of serving the public conversation. We also noted publicly that the only enforcement action available under the Distribution of Hacked Materials Policy was removal, which was no longer in alignment with new product capabilities, such as a label, that provide people with additional context.

Under Section 230 of the Communications Decency Act, services like Twitter are not considered publishers or speakers for hosting content by third parties. This protection allows us to act responsibly to promote healthy conversations by taking action against misinformation, abuse, harm and illegal activity that makes its way onto Twitter, and has been critical to promoting free

expression online. Eliminating Section 230 or prescribing reactionary government speech mandates will neither address concerns nor align with the First Amendment. Indeed, such actions could have the opposite effect, likely resulting in increased removal of speech, the proliferation of frivolous lawsuits, and severe limitations on our collective ability to address harmful content and protect people online. In addition, complex regulatory schemes or carveouts that favor already dominant market participants could stifle innovation and competition, ultimately giving consumers less choice online.

9. Has your company or any subsidiary entered into an agreement, either written or not, with a government to either limit access to certain information, or to not limit government officials posts or place warnings on those posts?

No, we have not entered into separate agreements with governments that would exempt them from complying with the Twitter Rules or enforce these rules differently with regards to specific government accounts.

Notwithstanding this, many countries, including the United States, have laws that may apply to Tweets and Twitter account content. In our continuing effort to make our services available to people everywhere, if we receive a valid and properly scoped request from an authorized entity, it may be necessary to withhold access to certain content in a particular country from time to time. Such withholdings will be limited to the specific jurisdiction that has issued the valid legal demand or where the content has been found to violate local laws. Upon receipt of requests to withhold content, we will promptly notify affected users unless we are prohibited from doing so (e.g., if we receive a court order under seal). We also clearly indicate within the product when content has been withheld and publish requests to withhold content on Lumen—unless, similar to our practice of notifying users, we are prohibited from doing so. The Transparency Report discloses information on government requests to remove content under local laws pursuant to our “Country Withheld Content” policy.

10. I asked if Twitter used a platform similar to the Tasks platform identified by Facebook. It was hard to hear your response but you did say that Twitter uses a similar platform.

- What is the name of this platform?
- How does it work?
- What is it primarily used for?
- Does Twitter collect metadata from this platform?
- Does any

Twitter uses a variety of methods and software to assist with businesses functions, including project management. For example, Twitter uses the software Jira to help with project management and issue tracking. Among numerous tasks, the software can help track issues logged in a range of contexts, including questions from employees to human resources or computer issues flagged to the IT Department.

11. Do you coordinate with any other company or outside group when you make decisions about content moderation?

Twitter does not coordinate with other entities when making content moderation decisions. However, we have partnerships with government agencies, nonprofits, and industry peers to facilitate information sharing to inform our policy and enforcement decisions.

For example, the National Center for Missing & Exploited Children, whose mission is to help find missing children, reduce child sexual exploitation, and prevent child victimization, is an important partner for Twitter and our industry peers. When we are made aware of content depicting or promoting child sexual exploitation, including links to images or content or third party sites where this content can be accessed, the material is removed without further notice and reported to NCMEC. While our general practice is to notify Twitter users when their content is reported to third-parties or law enforcement, we do not notify users when the reported content includes child sexual exploitation material. Furthermore, we participate in NCMEC's hash sharing database for industry and non-governmental organizations which consists of image and video hashes of known child sexual abuse material.

We also partner with nonprofits dedicated to child protection across the globe. In addition to our important relationship with NCMEC, Twitter is an active member of the Technology Coalition. This industry-led non-profit organization strives to eradicate child sexual exploitation by mentoring emerging or established companies, sharing trends and best-practices across industry, and facilitating technological solutions across the ecosystem. The Technology Coalition serves as an effective model because it gives companies the flexibility to create, test, and iterate across our diverse products and models.

12. Does Twitter receive any information from any other company or entity other than Twitter about posts and content moderation decisions?

Twitter has numerous partnerships that we rely on to better inform policy and decision making. In addition to the partnerships described above, Twitter is part of the Global Internet Forum to Counter Terrorism, which brings together industry, government, civil society, and academia to share information and collaborate to counter terrorist or extremist content online. Through the GIFCT, we have assembled a shared industry database of "hashes" or digital "fingerprints" for violent terrorist propaganda that spans more than 100,000 hashes. 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 block extremist content before it is posted.

We also began to work with a small group of companies to test a new collaborative system to share URLs. Because Twitter does not allow files other than photos or short videos to be uploaded, one of the behaviors we saw from those seeking to promote terrorism was to post links to other services where people could access files, longer videos, PDFs, and other materials. Our pilot system allows us to alert other companies when we removed an account or Tweet that linked to material that promoted terrorism hosted on their service. This information sharing ensures the hosting companies can monitor and track similar behavior, taking enforcement action pursuant with their individual policies. This is not a high-tech approach, but it is simple and effective, recognizing the resource constraints of smaller companies.

In order to safeguard the conversation regarding the 2020 U.S. election, we also have partnerships with leaders in civic tech, industry, and governments organizations, such as the National Association of Secretaries of State, National Association of State Election Directors, Department of Homeland Security, Federal Bureau of Investigation, Department of Justice, Office of the Director of National Intelligence, and elections officials across the country. We have also developed partnerships with news organizations, civil society, and others, which have been instrumental in informing policies and helping to identify potential threats regarding the integrity of the election conversation occurring on Twitter.

13. What procedure or policies does Twitter have in place to ensure content moderation is done in a objective manner?

Twitter does not use political viewpoints, perspectives, or party affiliation to make decisions, and we have taken several steps to ensure objective content moderation. For example, the Twitter Rules themselves are objective and not rooted in a particular ideology; the rules are focused on preventing harm and safeguarding the public conversation. In addition, we take a behavior-first approach to content moderation enforcement, meaning we look at how accounts behave before we look at the content they are posting. Moreover, we have invested in advancing procedural fairness, to facilitate impartial decision-making and provide avenues for individuals to appeal decisions if there has been a mistake.

14. In response to Senator Ernst you mentioned a tool for tracking all content moderation decisions.

- **What is the name of this tool?**
- **What information does it track?**
- **Who has access to this tool?**
- **Does it track who makes a content moderation decision?**
- **Do you collect metadata on this platform? Are you willing to share trends based on this metadata with Congress?**
- **Are Republicans subject to content moderation more than Democrats based on the information you have from this platform?**
- **You indicated there are safeguards in this system. Please outline and describe each step in the content moderation process, including the process for appeals.**

Twitter does not use political viewpoints, perspectives, or party affiliation to make content moderation decisions. We apply the Twitter Rules impartially. Twitter uses a variety of methods and software to assist with businesses functions, including project management. For example, Twitter uses the software Jira to help with project management and issue tracking, including in some cases where there may be potential violations of Twitter Rules. We work to be transparent about our enforcement, and have provided information about enforcement trends in the [Twitter Transparency Center](#).

With regards to appeals, if an account was suspended or locked in error, an individual can appeal. First, the individual must log in to the account that is suspended and file an appeal. The individual must describe the nature of the appeal and provide an explanation of why the account is not in violation of the Twitter Rules. Twitter employees will typically engage with the account holder via email to resolve the appeal. Mistakes in enforcement — made either by a human or algorithm —

are inevitable, and why we strive to make appeals easier. We recognize that enhancing procedural fairness, including through a straightforward appeals process is a critical part of building consumer trust and we look forward to working with the Committee on this issue.

15. Why do you only place warnings on some stories published by traditional news sources and not go one step farther and provide more context on all posts from traditional news sources when the text in the post or the title of the post may not fully tell an accurate story?

Our policies focus on labeling misinformation in the contexts where there is the greatest risk of harm, including COVID-19, civic integrity, and manipulated media. However, we continue to explore additional ways in which we can provide additional context to people who use Twitter, in order to address potential harms associated with misinformation.

Questions from Senator Blackburn

1. During the November 3, 2020 election (before and after this date), did Twitter maintain any informal or formal lists of U.S. public officials who were specifically targeted for special monitoring of their Twitter posts?

During the election period, we used a combination of human and automated mechanisms to enforce our policies. For example, we reviewed Tweets reported as potential violations by the public, civil society partners, or government agencies. We also used automated systems to detect suspicious behaviors or identify potential violations of our rules. Twitter's enforcement teams prioritized the review of Tweets from the accounts of each of the presidential candidates and their campaigns and reviewed each to ensure compliance with our terms of service, beginning two weeks prior to election day.

2. During the November 3, 2020 election (before and after this date), did Twitter's enforcement teams prioritize the review of Tweets from the official and campaign accounts of any U.S. Senators for compliance with Twitter's terms of service?

As noted in the response to Question 1, we used a variety of mechanisms to enforce our policies during the election period. All people who use Twitter are governed by The Twitter Rules and all incorporated policies, Privacy Policy, and Terms of Service, which collectively make up the "Twitter User Agreement." Our enforcement teams receive reports of Tweets and other media to review against our Terms of Service through a variety of channels, including through in-app reporting from the people who use our service to reports from government partners, political stakeholders, the press, and civil society.

3. In advance of the January 5, 2021 Georgia Senate run-off elections, are Twitter's enforcement teams prioritizing the review of Tweets from the official and campaign accounts of any U.S. Senators for compliance with Twitter's terms of service?

Twitter's work to safeguard the conversation regarding the 2020 U.S. election is ongoing; we continue to work to safeguard the conversation around the 2021 Georgia Senate run-off elections.

4. In advance of the January 5, 2021 Georgia Senate run-off elections, are Twitter's enforcement teams prioritizing the review of Tweets from the official and campaign accounts of any candidates for U.S. Senate for compliance with Twitter's terms of service?

Twitter's work to safeguard the conversation regarding the 2020 U.S. election is ongoing; we continue to work to safeguard the conversation around the 2021 Georgia Senate run-off elections.

5. Does Twitter's enforcement team prioritize the review of Tweets from any reporters from One America News for compliance with Twitter's terms of service?

All people who use Twitter are governed by The Twitter Rules and all incorporated policies, Privacy Policy, and Terms of Service, which collectively make up the "Twitter User Agreement." Our enforcement teams receive reports of Tweets and other media to review against our Terms of

Service through a variety of channels, including through in-app reporting from the people who use our service to reports from government partners, political stakeholders, the press, and civil society.

6. Does Twitter’s enforcement team prioritize the review of Tweets from any reporters from Fox News for compliance with Twitter’s terms of service?

As noted in the response to Question 1, we used a variety of mechanisms to enforce our policies during the election period. All people who use Twitter are governed by The Twitter Rules and all incorporated policies, Privacy Policy, and Terms of Service, which collectively make up the “Twitter User Agreement.” Our enforcement teams receive reports of Tweets and other media to review against our Terms of Service through a variety of channels, including through in-app reporting from the people who use our service to reports from government partners, political stakeholders, the press, and civil society.

7. Twitter placed a blue elections flag label on a November 5, 2020 clip from Fox News’ Hannity Show quoting Rep. Jim Jordan. The flag stated, “This claim about election fraud is disputed.” Upon clicking the notice, the user is directed to a page which reads, “Voter fraud of any kind is exceedingly rare in the US, election experts confirm.” In the flagged clip, Congressman Jim Jordan discusses Justice Samuel Alito’s decision ordering Pennsylvania counties to comply with a state directive to separate late ballots received after Election Day. Please explain why Twitter’s election fraud label was imposed on a Tweet that discussed a Supreme Court decision ordering Pennsylvania officials to follow election laws.

We have heard from the people who use Twitter that we should not determine the truthfulness of Tweets and we should provide context to help people make up their own minds in cases where the substance of a Tweet is disputed. Consistent with this feedback from our customers, we have expanded our enforcement options to allow us to label misinformation. When we label Tweets, we link to Twitter conversation that shows three things for context: (1) factual statements; (2) counterpoint opinions and perspectives; and (3) ongoing public conversation around the issue. In the case referenced, we labeled the Tweet to provide individuals with additional context regarding the claims being made.

8. Does Twitter maintain a list of users that have a history of posting “misleading” or “false” speech as determined by Twitter’s criteria?

When determining whether to take enforcement action, we may consider a number of factors, including whether the account holder has a history of violating our policies. If an account holder repeatedly violates our Rules then the penalty for violating our policies will increase. This includes requiring violators to remove the Tweet(s) and taking additional actions like verifying account ownership and/or temporarily limiting their ability to Tweet for a set period of time.

9. Does Twitter maintain a list of users with links to Chinese state-sponsored actors that have a history of posting “misleading” or “false” speech as determined by Twitter’s criteria?

In August 2019, Twitter disclosed 936 accounts originating from within the People’s Republic of China (PRC). These accounts were deliberately and specifically attempting to sow political discord in Hong Kong, including undermining the legitimacy and political positions of the protest

movement on the ground. Based on our intensive investigations, we have reliable evidence to support that this is a coordinated state-backed operation. Specifically, we identified large clusters of accounts behaving in a coordinated manner to amplify messages related to the Hong Kong protests.

As Twitter is blocked in PRC, many of these accounts accessed Twitter using VPNs. However, some accounts accessed Twitter from specific unblocked IP addresses originating in mainland China. The accounts represent the most active portions of this campaign; a larger, spammy network of approximately 200,000 accounts — many created following our initial suspensions — were proactively suspended before they were substantially active on the service. These accounts were disclosed in Twitter’s comprehensive public archive of Tweets and media associated with suspected state-backed information operations on Twitter.

10. Based on previous cybersecurity breaches, has Twitter suspected in the past that Chinese state-sponsored actors have targeted Twitter accounts affiliated with Chinese dissidents for hacking or other internet attacks?

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 State’s Global Engagement Center. We look forward to continued cooperation with federal partners on the threats posed by state-backed information operations, because in certain circumstances only they have access to information critical to our joint efforts to stop bad faith actors.

11. Does Twitter’s enforcement team prioritize the review of Tweets from any accounts affiliated with China’s state propaganda outlets, which include the People’s Daily, China Daily, China News Service, the Global Times, and CCTV, for compliance with Twitter’s terms of service?

An important part of our work is providing people with context so that they can make informed decisions about the content they see on Twitter. Accordingly, in August 2020, we made the decision to add labels to the Twitter accounts of key government officials, with a focus on senior officials who are the voices of the state abroad, and accounts belonging to state-affiliated media entities, their editors-in-chief, and senior staff. We believe providing these labels are an important step, so that when people see an account discussing geopolitical issues from another country, they have context about its affiliation and who it represents. We are applying labels to the accounts that represent the five permanent members of the UN Security Council: China, France, Russian Federation, the United Kingdom, and the United States before expanding to a wider range of countries. Additionally, we do not permit news media entities controlled by state authorities to purchase advertisements. This policy extends to individuals reporting on behalf of or who are directly affiliated with such entities.

12. Of the \$3.59 billion in revenue Twitter generated in 2019, how much ad revenue did Twitter generate from relationships with Chinese businesses or government agencies?

The Twitter [annual report](#) contains additional information about the company’s revenue streams.

13. How much ad revenue did Twitter generate from Huawei for marketing the launch of its tablet, the Huawei Ascend Mate 7? For reference, this marketing campaign is found on

Twitter's "Success Stories" page:

https://marketing.twitter.com/en_apac/success-stories/huaweis-new-tablet-goes-global-with-promoted-tweets.

The Twitter [annual report](#) contains additional information about the company's revenue streams. As a global platform, Twitter provides advertising services to companies around the world. All ads must comply with Twitter policies, which among other things, prohibit ads from state media or that promote unacceptable business practices.

14. How much ad revenue did Twitter generate from Huawei for marketing the launch of its phone, the Huawei Mate 30 Pro 5G? For reference, this marketing campaign is found on Twitter's "Success Stories" page:

https://marketing.twitter.com/en_gb/success-stories/huawei-mate-30-launch.

The Twitter [annual report](#) contains additional information about the company's revenue streams. All ads must comply with Twitter policies, which among other things, prohibit ads from state media. As a global platform, Twitter provides advertising services to companies around the world. All ads must comply with Twitter policies, which among other things, prohibit ads from state media or that promote unacceptable business practices.

15. How much ad revenue did Twitter generate from Huawei for marketing the @HuaweiMobile brand in advance of the GSMA Mobile World Congress, a global tech conference? For reference, this marketing campaign is found on Twitter's "Success Stories" page:

https://marketing.twitter.com/en_apac/success-stories/how-huawei-became-the-most-talked-about-brand-at-a-global-event.

The Twitter [annual report](#) contains additional information about the company's revenue streams. As a global platform, Twitter provides advertising services to companies around the world. All ads must comply with Twitter policies, which among other things, prohibit ads from state media or that promote unacceptable business practices.

Questions from Senator Cruz

1. The following questions examine how Twitter views its activities in moderating and directing its platform:

a. When Twitter hosts, unaltered, the material of a third party – as in the form of a basic Tweet – is Twitter acting as publisher?

Under subsection (c)(1), Section 230 of the Communications Decency Act provides that neither providers nor the people who use our service are to “be treated as the publisher or speaker of any information provided by another information content provider.” Whether a plaintiff’s claim treats an entity as a publisher for purposes of Section 230 is a determination that a court would make based on the specific details of a case.

Section 230’s protection has been critical to preserving free expression online, as well as promoting innovation. Looking forward, we encourage Congress to work with industry and civil society to build upon Section 230’s foundation, whether it be through additions to Section 230, industry-wide self-regulation best practices, or a new legislative framework. We believe that the best way to address concerns with content moderation is to require the publication of moderation processes and practices, a straightforward process to appeal decisions, and best efforts around algorithmic choice, while protecting the privacy of individuals. We look forward to working with the committee to achieve these goals.

b. When Twitter blocks a Tweet, is it acting as a publisher?

Under subsection (c)(1), Section 230 of the Communications Decency Act provides that neither providers nor the people who use our service are to “be treated as the publisher or speaker of any information provided by another information content provider.” Whether a plaintiff’s claim treats an entity as a publisher for purposes of Section 230 is a determination that a court would make based on the specific details of a case. Twitter does not block Tweets. In cases where an account violates the Twitter Rules, it may be subject to a range of enforcement actions, which can include account suspension or removal of a Tweet.

c. When Twitter intentionally limits the reach of a Tweet, is it acting as a publisher?

Under subsection (c)(1), Section 230 of the Communications Decency Act provides that neither providers nor the people who use our service are to “be treated as the publisher or speaker of any information provided by another information content provider.” Whether a plaintiff’s claim treats an entity as a publisher for purposes of Section 230 is a determination that a court would make based on the specific details of a case. In cases where an individual violates the Twitter Rules, they may be subject to various enforcement actions, including actions designed to provide additional context to individuals about the Tweet. For example, in some circumstances where we do not remove content which violates the civic integrity policy, we may take other actions, including applying a label and/or warning message to the content where it appears in the Twitter product; showing a warning to people before they share or like the content; reducing the visibility of the

content on Twitter and/or prevent it from being recommended; and/or providing a link to additional explanations or clarifications, such as in a Twitter Moment or relevant Twitter policies.

d. When Twitter covers the face of a Tweet with a warning or other label written by Twitter, requiring the user to click through in order to access the content of that Tweet, is Twitter acting a publisher?

Under subsection (c)(1), Section 230 of the Communications Decency Act provides that neither providers nor the people who use our service are to “be treated as the publisher or speaker of any information provided by another information content provider.” Whether a plaintiff’s claim treats an entity as a publisher for purposes of Section 230 is a determination that a court would make based on the specific details of a case.

In cases where an individual violates the Twitter Rules, they may be subject to various enforcement actions. In certain cases involving a violation of the Twitter Rules by a world leader, we will err on the side of leaving the content up if there is a clear public interest in doing so. In such cases, we may place the violative content behind a warning notice that provides context about the violation and allows people to click through should they wish to see the content.

e. When Twitter labels a Tweet, instructing the user that the information in the Tweet is subject to controversy or is in dispute, is Twitter acting as a publisher?

Under subsection (c)(1), Section 230 of the Communications Decency Act provides that neither providers nor the people who use our service are to “be treated as the publisher or speaker of any information provided by another information content provider.” Whether a plaintiff’s claim treats an entity as a publisher for purposes of Section 230 is a determination that a court would make based on the specific details of a case.

In cases where an individual violates the Twitter Rules, they may be subject to various enforcement actions. We have heard from the people who use Twitter that we should not determine the truthfulness of Tweets and we should provide context to help people make up their own minds in cases where the substance of a Tweet is disputed. When we label Tweets, we link to Twitter conversation that shows three things for context: (1) factual statements; (2) counterpoint opinions and perspectives; and (3) ongoing public conversation around the issue. We will only add descriptive text that is reflective of the existing public conversation to let people determine their own viewpoints.

f. When Twitter labels a Tweet, informing the user that Twitter has determined that the information in the Tweet is untrue, is Twitter acting as a publisher?

Under subsection (c)(1), Section 230 of the Communications Decency Act provides that neither providers nor the people who use our service are to “be treated as the publisher or speaker of any information provided by another information content provider.” Whether a plaintiff’s claim treats an entity as a publisher for purposes of Section 230 is a determination that a court would make based on the specific details of a case.

We have heard from the people who use Twitter that we should not determine the truthfulness of Tweets and we should provide context to help people make up their own minds in cases where the substance of a Tweet is disputed. When we label Tweets, we link to Twitter conversation that shows three things for context: (1) factual statements; (2) counterpoint opinions and perspectives; and (3) ongoing public conversation around the issue. We will only add descriptive text that is reflective of the existing public conversation to let people determine their own viewpoints.

g. When Twitter makes a judgment regarding the content of a tweet, and accordingly reduces or limits its dissemination, is Twitter acting as a publisher?

Under subsection (c)(1), Section 230 of the Communications Decency Act provides that neither providers nor the people who use our service are to “be treated as the publisher or speaker of any information provided by another information content provider.” Whether a plaintiff’s claim treats an entity as a publisher for purposes of Section 230 is a determination that a court would make based on the specific details of a case. In cases where an individual violates the Twitter Rules, they may be subject to various enforcement actions.

h. When Twitter conditions usage of its platform upon not sharing information, is Twitter acting as publisher?

Under subsection (c)(1), Section 230 of the Communications Decency Act provides that neither providers nor the people who use our service are to “be treated as the publisher or speaker of any information provided by another information content provider.” Whether a plaintiff’s claim treats an entity as a publisher for purposes of Section 230 is a determination that a court would make based on the specific details of a case. When individuals use Twitter, they agree to comply with the Twitter Rules, which are designed to prevent harm. For example, the Twitter Rules prohibit the [promotion](#) of terrorism or violent extremism.

i. When Twitter categorizes and organizes Tweets, deciding on their order or presence in a “Trending” list, or employs similar tools to suggest content, is Twitter acting as a publisher?

Under subsection (c)(1), Section 230 of the Communications Decency Act provides that neither providers nor the people who use our service are to “be treated as the publisher or speaker of any information provided by another information content provider.” Whether a plaintiff’s claim treats an entity as a publisher for purposes of Section 230 is a determination that a court would make based on the specific details of a case.

The goal of Trends is to provide individuals who use Twitter topics that are popular now, so they can discover emerging topics of discussion. Trends are, by default, tailored to an individual based on who they follow, their interest, and other factors. However, individuals can choose to see Trends that are not tailored for them by selecting a specific location, which identify popular topics among people in a specific geographic location.

2. Twitter recently limited the distribution and dissemination of reporting from the *New York Post* regarding Joe Biden, his family’s business dealings, and his campaign.

a. Has Twitter ever acted to limit the distribution of:

i. News stories regarding President Trump’s personal financial information, including but not limited to confidential tax records?

The Twitter Rules prohibit posting other people's private information (such as home phone number and address) without their express authorization and permission. They also prohibit threatening to expose private information or incentivizing others to do so. News stories referencing the fact that a disclosure has been made, which does not directly share private information, would not fall under this prohibition. In addition, our Distribution of Hacked Materials policy prohibits distribution of hacked materials, but does not prohibit reporting on such materials.

ii. The disclosures of Edward Snowden, including but not limited to the illegally obtained material for which he is currently under federal investigation?

Our Distribution of Hacked Materials was not in place at the time of the Snowden disclosures. On October 23, 2020, we revised the Distribution of Hacked Materials to state that we will no longer remove hacked content unless it is directly shared by hackers or groups directly associated with a hack.

iii. Stories regarding First Lady Trump, and personal conversations recorded without her knowledge or consent?

The Twitter Rules prohibit posting other people's private information (such as home phone number and address) without their express authorization and permission. They also prohibit threatening to expose private information or incentivizing others to do so. News stories referencing the fact that a disclosure has been made, which does not directly share private information, would not fall under this prohibition.

b. If the answer to the previous questions is “no,” please explain, with specificity, the disparate treatment between these news items. If the answer is “yes,” please state with specificity what actions were taken to limit distribution.

Twitter is constantly refining its rules and enforcement practices to better safeguard the public conversation and respond to the feedback from the public. For example, following the application of our Distribution of Hacked Materials Policy against specific content shared by @NYPost, we received significant feedback — both positive and negative — on the action.. After reviewing the feedback, we made changes within 24 hours to the policy to address concerns that there could be unintended consequences to journalists, whistleblowers and others in ways that are contrary to Twitter’s purpose of serving the public conversation. We also noted publicly that the only enforcement action available under the Distribution of Hacked Materials Policy was removal, which was no longer in alignment with new product capabilities, such as the application of a label, that provides people with additional context.

3. Despite the unprecedented steps Twitter has taken to define for the average user which information is “false,” and which information is to be relied upon, you were unable to answer

questions on the subject of voter fraud in your testimony before the Committee. Please answer the following questions regarding the public discussion of voter fraud:

a. Does voter fraud exist?

Twitter does not tell people what information is true or false as it relates to civic processes. Our intention is to connect the dots of conflicting statements and show the information in dispute so people can judge for themselves.

b. Are any of the executives at Twitter, including but not limited to those who are making content moderation policy decisions, experts in voter fraud?

We seek to employ individuals with a range of backgrounds, expertise, and skill sets.

c. What sources of authority does Twitter rely upon in making determinations regarding the truthfulness of claims of voter fraud?

Twitter's enforcement of its civic integrity policy does not tell people what information is true or false. Our intention is to connect the dots of conflicting statements and show the information in dispute so people can judge for themselves. We have heard from the people who use Twitter that we should not determine the truthfulness of Tweets and we should provide context to help people make up their own minds in cases where the substance of a Tweet is disputed. When we label Tweets, we link to Twitter conversation that shows three things for context: (1) factual statements; (2) counterpoint opinions and perspectives; and (3) ongoing public conversation around the issue. We will only add descriptive text that is reflective of the existing public conversation to let people determine their own viewpoints.

d. Would the following statement violate Twitter's election information policies: "Absentee ballots remain the largest source of potential voter fraud"?

When making enforcement decisions, Twitter considers a range of behaviors, including the content of the Tweet itself, but also platform behavior. Thus, it is difficult to provide an answer to such hypothetical scenarios.

e. Would the following statement violate Twitter's election information policies: "Voter fraud is particularly possible where third-party organizations, candidates, and political party activists are involved in handling absentee ballots"?

When making enforcement decisions, Twitter considers a range of factors. This can include the content of the Tweet itself, but also platform behavior. Thus, it is difficult to provide an answer to such hypothetical scenarios.

f. In making determinations regarding the truth or verifiability of voter fraud claims, does Twitter employ corporate values, beliefs, priorities, or opinions when deciding what content is removed, republished, moderated, labelled, or otherwise promoted or demoted? What are those values?

The Twitter Rules are objectively enforced, without regards to political viewpoints, party affiliation, political ideology, or personal beliefs. As outlined in our testimony, in order to build trust, we fully support efforts to enhance transparency, procedural fairness, privacy, and algorithmic choice.

g. Does Twitter make these determinations on a viewpoint-neutral basis?

In developing and enforcing our rules for the service we seek to be impartial. As part of this, Twitter supports and has invested in advancing procedural fairness to ensure decisions are made impartially and individuals have mechanisms to appeal in cases where they believe a mistake has been made.

4. The day before you testified before the Committee, you and I spoke on the telephone and you told me that Twitter was committed to transparency with regard to its content moderation policies and enforcement. Accordingly, the following questions relate to Twitter's enforcement of its content moderation policies. For this question and its subparts, please construe "content moderation policies" broadly, including decisions regarding the position or order in which content is displayed, the position or order in which users or content appear in searches, whether users or content are promoted or demoted, and all other modifications of content, such as flagging, qualifying, labelling, and denoting.

a. In your enforcement of content moderation policies:

- i. How many times has Twitter blocked, flagged, censored, limited the reach of, or otherwise affected the tweets, posts, or content of Republican candidates for office between January 1, 2016 and November 24, 2020? Please include enforcements that were later reversed or recognized to be in error.**

Twitter does not track the political affiliation of the people who use our service. Our biannual [Twitter Transparency Center](#) highlights trends in enforcement of our Rules, legal requests, intellectual property-related requests, and email privacy best practices.

- ii. How many times has Twitter blocked, flagged, censored, limited the reach of, or otherwise affected the tweets, posts, or content of Democratic candidates for office between January 1, 2016 and November 24, 2020? Please include enforcements that were later reversed or recognized to be in error.**

Twitter does not track the political affiliation of the people who use our service. Our biannual [Twitter Transparency Center](#) highlights trends in enforcement of our Rules, legal requests, intellectual property-related requests, and email privacy best practices.

- iii. How many times has Twitter blocked, flagged, censored, limited the reach of, or otherwise affected the tweets, posts, or content of Republican elected officials between January 1, 2016 and November 24,**

2020? Please include enforcements that were later reversed or recognized to be in error.

Twitter does not track the political affiliation of the people who use our service. Our biannual [Twitter Transparency Center](#) highlights trends in enforcement of our Rules, legal requests, intellectual property-related requests, and email privacy best practices.

- iv. How many times has Twitter blocked, flagged, censored, limited the reach of, or otherwise affected the tweets, posts, or content of Democratic elected officials between January 1, 2016 and November 24, 2020? Please include enforcements that were later reversed or recognized to be in error.**

Twitter does not track the political affiliation of the people who use our service. Our biannual [Twitter Transparency Center](#) highlights trends in enforcement of our Rules, legal requests, intellectual property-related requests, and email privacy best practices.

- b. For each instance marked and counted in the previous question, please provide the name of the account affected, the content of the material affected, and the specific reason for the enforcement action.**

Our biannual [Twitter Transparency Center](#) highlights trends in enforcement of our Rules, legal requests, intellectual property-related requests, and email privacy best practices.

- c. In your enforcement of election information content moderation policies, are decisions made and executed by humans, or by algorithm with preset code?**

Twitter uses a combination of automation and human review to adjudicate reports of violations and make determinations on whether activity violates our rules.

- i. Has anyone been tasked with keeping track of which content or users are affected by these policies?**

We have made efforts to be as transparent as possible about our enforcement decisions. Because Twitter is a public platform, enforcement decisions are often apparent to the public. In some contexts, we have also taken steps to provide additional datasets to researchers. For example, in [October 2018](#), we published the first comprehensive archive of Tweets and media associated with known state-backed information operations on Twitter. This one of a kind resource, used by researchers, journalists and experts around the world, now spans operations across 15 countries, including more than nine terabytes of media and 200 million Tweets.

- ii. If not, has anyone been tasked with keeping track of other information of this type, in other contexts?**

We have made efforts to be as transparent as possible about our enforcement decisions, providing information about enforcement through the Twitter Transparency Center. Apart from this, we fully support efforts to increase transparency around content moderation. We believe that companies like

Twitter should publish their moderation process, and should be transparent about how cases are reported and reviewed, how decisions are made, and what tools are used to enforce.

iii. If so, please explain the difference in the maintenance of records.

d. In drafting Twitter’s election information content moderation policies and enforcing those same policies with regard to the 2020 elections, did Twitter collaborate with, confer with, or defer to any outside individuals or organizations? If so, please list the individuals and organizations and state the nature of their relationship with Twitter.

As part of our civic integrity efforts, we have developed partnerships that allowed us to share information, gather input from experts, and better gain context on how misinformation was being spread and impacting the public conversation. These partnerships included leaders in civic tech, our peers, federal, state, and local governments organizations (e.g., National Association of Secretaries of State, National Association of State Election Directors, Department of Homeland Security, Federal Bureau of Investigation, Department of Justice, Office of the Director of National Intelligence, and elections officials across the country), news organizations, and civil society, among others.

i. Does Twitter take account of the political, philosophical, or ideological orientation or reputation of those sources with which it cooperates in executing its election information content moderation policies?

Twitter does not use political viewpoints, perspectives, or party affiliation to make any decisions.

e. The following questions relate to the individuals with supervisory authority who are responsible for the formulation and implementation of content moderation policies.

i. Among those individuals with supervisory authority who make substantive decisions regarding content moderation policy, how many:

1. Self-identify or are registered as Democrats?

Twitter does not ask employees to disclose their political affiliation.

2. Self-identify or are registered as Republicans?

Twitter does not ask employees to disclose their political affiliation.

3. Would identify themselves as “liberal?”

Twitter does not ask employees to disclose their political affiliation.

4. Would identify themselves as “conservative?”

Twitter does not ask employees to disclose their political affiliation.

5. Have donated to:

a. The Democratic Party?

Twitter does not ask employees to publicly disclose political contributions, apart from disclosures required under existing laws.

b. A candidate running for office as a Democrat?

Twitter does not ask employees to publicly disclose political contributions, apart from disclosures required under existing laws.

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

Twitter does not ask employees to publicly disclose political contributions, apart from disclosures required under existing laws.

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

Twitter does not ask employees to publicly disclose political contributions, apart from disclosures required under existing laws.

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

Twitter does not ask employees to publicly disclose political contributions, apart from disclosures required under existing laws.

f. The Republican Party?

Twitter does not ask employees to publicly disclose political contributions, apart from disclosures required under existing laws.

g. A candidate running for office as a Republican?

Twitter does not ask employees to publicly disclose political contributions, apart from disclosures required under existing laws.

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

Twitter does not ask employees to publicly disclose political contributions, apart from disclosures required under existing laws.

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

Twitter does not ask employees to publicly disclose political contributions, apart from disclosures required under existing laws.

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

Twitter does not ask employees to publicly disclose political contributions, apart from disclosures required under existing laws.

6. Worked on or volunteered for a Democratic campaign?

Decisions at Twitter are made without regard to political viewpoints, party affiliation, or political ideology.

7. Worked on or volunteered for a Republican campaign?

Decisions at Twitter are made without regard to political viewpoints, party affiliation, or political ideology.

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

Decisions at Twitter are made without regard to political viewpoints, party affiliation, or political ideology.

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

Decisions at Twitter are made without regard to political viewpoints, party affiliation, or political ideology.

10. Worked on or interned for a Democratic administration or candidate?

Decisions at Twitter are made without regard to political viewpoints, party affiliation, or political ideology.

11. Worked on or interned for a Republican administration or candidate?

Decisions at Twitter are made without regard to political viewpoints, party affiliation, or political ideology.

ii. Following a Constitution Subcommittee in April 2019 in which a Twitter executive testified, I submitted questions similar to those in 4.d.1 above. Twitter responded by refusing to answer the questions asked and instead stating that it “does not use political ideology as a factor in its hiring decisions.”

1. Does Twitter believe that it can adopt and enforce viewpoint neutral content moderation rules if the individuals developing and enforcing those rules overwhelmingly identify as Democrats and progressives?

Yes. Decisions at Twitter are made without regard to political viewpoints, party affiliation, or political ideology. To help ensure that decisions are made objectively, we have invested in developing fair processes.

2. When hiring individuals to make substantive decisions regarding content moderation policies or to enforce those policies, does Twitter consider any personal characteristics or demographic information, including but not limited to race, ethnicity, sex, gender, or sexual orientation, to ensure that Twitter’s content moderation policies and enforcement actions reflect diverse views?

We believe that being a diverse and inclusive company is key to serving the public conversation. By 2025 we have set a goal to have at least 25% of our overall U.S. workforce be underrepresented minorities. We consistently share our progress with the public every quarter. Our most recent report reflects U.S. ethnicity for all roles and levels at 41.4% White, 28.4% Asian, 6.3% Black, 5.2% Latinx, 3.7 Multiracial, and less than 1% Indigenous.

a. If the answer is “yes”, please identify the characteristics or demographic information Twitter considers and explain how that characteristic promotes view-point neutrality but ensuring diverse political ideologies does not.

As a company, Twitter is focused on advancing the principle of procedural fairness in our decision-making across the board. Procedural fairness at Twitter means we ensure that all decisions are made without using political viewpoints, party affiliation, or political ideology, whether related to automatically ranking content on our service or how we develop or enforce the Twitter Rules. Our Twitter Rules are not based on ideology or a particular set of beliefs. We believe strongly in being impartial, and we strive to enforce our Twitter Rules fairly.

In addition, we strive to give people an easy, clear way to appeal decisions we make that they think are not right. Mistakes in enforcement — made either by a human or algorithm — are inevitable, and why we strive to make appeals easier. We believe that all companies should be required to provide a straightforward process to appeal decisions. This makes certain people can let us know when we do not get it right, so that we can fix any mistakes and make our processes better in the future.

f. A top Twitter executive—Vijaya Gadde—has been praised as the leader of the “resistance” and the “lead architect of the policy approach that led Twitter to clamp down on . . . President Donald Trump.”¹

i. Does this description of the top policy executive at Twitter reflect the orientation and priorities of the election information content moderation team as a whole?

Decisions at Twitter are made without regard to political viewpoints, party affiliation, or political ideology. We focus our content moderation efforts to reduce harm and facilitate the public conversation.

ii. Is Ms. Gadde ultimately responsible for what is determined to be “true” and “untrue” regarding claims made by users on Twitter?

Twitter does not tell people what information is true or false as it relates to information concerning COVID-19, civic integrity, or manipulated media. Our intention is to connect the dots of conflicting statements and show the information in dispute so people can judge for themselves. We have heard from the people who use Twitter that we should not determine the truthfulness of Tweets and we should provide context to help people make up their own minds in cases where the substance of a Tweet is disputed. When we label Tweets, we link to Twitter conversation that shows three things for context: (1) factual statements; (2) counterpoint opinions and perspectives; and (3) ongoing public conversation around the issue. We will only add descriptive text that is reflective of the existing public conversation to let people determine their own viewpoints.

iii. If she is not ultimately responsible, who is?

No one individual makes content decisions at Twitter. As Twitter’s Chief Executive Officer Jack Dorsey Tweeted in May 2020: “Fact check: there is someone ultimately accountable for our actions as a company, and that’s me.” We fully support efforts to enhance transparency and procedural fairness, so that individuals have trust in our company and the platforms they use. We also support efforts to increase competition and innovation, so that individuals have choices if a platform fails to earn their trust.

¹<https://www.politico.com/news/magazine/2020/10/28/twitter-vijaya-gadde-free-speech-policies-technology-social-media-429221>

Questions from Senator Grassley

1. How do your users hold you and your companies accountable for your content moderation practices? Congress has the benefit of bringing you before this Committee to answer questions. But when one of my constituents thinks his or her speech was wrongfully moderated or fact-checked, why should they have any faith that their objection will even be heard?

As a company, Twitter is focused on advancing the principle of procedural fairness in our decision-making across the board. We strive to give people an easy, clear way to appeal decisions we make that they think are not right. Mistakes in enforcement — made either by a human or algorithm — are inevitable, and why we strive to make appeals easier. We believe that all companies should be required to provide a straightforward process to appeal decisions. This makes certain people can let us know when we do not get it right, so that we can fix any mistakes and make our processes better in the future. We fully support efforts to increase procedural fairness and look forward to working with your office on this issue.

2. Do you agree that users should be entitled to due process when content they post is taken down or moderated? And do you agree that there could be more transparency by your companies in explaining why certain speech or content is moderated?

We agree that increased transparency and procedural fairness are key to address concerns with content moderation. We believe that all companies should be required to provide a straightforward process to appeal decisions. In addition, we believe that all companies should publish their content moderation process, and be transparent about how cases are reported and reviewed, how decisions are made, and what tools are used to enforce rules.

Questions from Senator Hawley

1. Does Twitter have a policy prohibiting its employees from coordinating content moderation decisions with outside companies such as Facebook or Google, where such moderation is not strictly required by law?

Twitter does not coordinate its content moderation decisions with outside entities. However, Twitter has numerous partnerships that we rely on to better inform decision making and facilitate information sharing. For example, we share information in three critical areas: combatting child sexual exploitation, prohibiting terrorism and violent extremism, and safeguarding the conversation about the U.S. election.

The National Center for Missing & Exploited Children is a nonprofit whose mission is to help find missing children, reduce child sexual exploitation, and prevent child victimization. NCMEC is an important partner for Twitter and our industry peers. When we are made aware of content depicting or promoting child sexual exploitation, including links to images or content or third party sites where this content can be accessed, the material is removed without further notice and reported to NCMEC. We also partner with nonprofits dedicated to child protection across the globe. In addition to our important relationship with NCMEC, Twitter is an active member of the Technology Coalition. This industry-led non-profit organization strives to eradicate child sexual exploitation by mentoring emerging or established companies, sharing trends and best-practices across industry, and facilitating technological solutions across the ecosystem.

In addition, Twitter is part of the Global Internet Forum to Counter Terrorism, which brings together industry, government, civil society, and academia to share information and collaborate to counter terrorist or extremist content online. Through the GIFCT, we have assembled a shared industry database of "hashes" or digital "fingerprints" for violent terrorist propaganda that spans more than 100,000 hashes. 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 block extremist content before it is posted.

We have also begun to work with a small group of companies to test a new collaborative system to share URLs. Because Twitter does not allow files other than photos or short videos to be uploaded, one of the behaviors we saw from those seeking to promote terrorism was to post links to other services where people could access files, longer videos, PDFs, and other materials. Our pilot system allows us to alert other companies when we removed an account or Tweet that linked to material that promoted terrorism hosted on their service. This information sharing ensures the hosting companies can monitor and track similar behavior, taking enforcement action pursuant with their individual policies. This is not a high-tech approach, but it is simple and effective, recognizing the resource constraints of smaller companies.

Furthermore, in order to safeguard the conversation regarding the 2020 U.S. election, we have critical partnerships with leaders in civic tech, industry, and government organizations, such as the National Association of Secretaries of State, National Association of State Election Directors, Department of Homeland Security, Federal Bureau of Investigation, Department of Justice, Office of the Director of National Intelligence, and elections officials across the country. We also have

partnerships with news organizations, civil society, and others, which have been instrumental in informing policies and helping to identify potential threats regarding the integrity of the election conversation occurring on Twitter.

2. Please provide a list of all hashtags, URLs, and real-world individuals that Twitter has banned across its platform.

Our biannual [Twitter Transparency Center](#) highlights trends in enforcement of our Rules, legal requests, intellectual property-related requests, and email privacy best practices. In addition, in 2018, we published the first comprehensive archive of Tweets and media associated with known state-backed information operations on Twitter. This one of a kind resource, used by researchers, journalists and experts around the world, now spans operations across 15 countries, including more than nine terabytes of media and 200 million Tweets.