



1 Hacker Way
Menlo Park, CA 94025
United States

February 22, 2023

Chair Richard Durbin
Ranking Member Lindsey Graham
US Senate Committee on the Judiciary
224 Dirksen Senate Office Building
Washington DC, 20510

Dear Chair Durbin, Ranking Member Graham, and Members of the Committee:

Thank you for the questions for the record from the Senate Committee on the Judiciary, Subcommittee on Privacy, Technology, and the Law, hearing entitled “Algorithms and Amplification: How Social Media Platforms’ Design Choices Shape Our Discourse and Our Minds” on April 27, 2021. Per your request, attached are answers to the questions posed.

Sincerely,

Meta Platforms, Inc.

Questions from Senator Coons

Question 1. As you testified, it is not in Facebook’s long-term interest to pursue short-term engagement or to push users toward increasingly extreme content. Some have suggested, however, that the metrics Facebook uses to set goals and measure product teams’ performance may create structural incentives within the company potentially at odds with this long-term interest.¹ During last Tuesday’s hearing, Tristan Harris testified that although Facebook does not measure teams’ performance based on time spent on the site, it does currently rely on metrics like numbers of sessions and seven-day active users growth. He also said that Facebook abandoned a brief experiment in setting goals for teams based on other, non-engagement-based metrics for social impact.

- a. Please enumerate the metrics that Facebook uses to set goals for its product teams—particularly the News Feed Team—for purposes of performance and pay.**

Broadly speaking, our goals and incentives process is set up with the goal of building the best app over the long term. Teams and individual employees are evaluated on a number of organization-specific factors, as well as the employee’s expectations and goals set that year. During performance reviews, managers assess and calibrate the performance of the people on their team and assign them a performance rating, which drives compensation.

The entire process is holistic and does not turn on any single goal or metric. The employees who work on developing our apps have many goals related to their app’s overall health including safety, security, engagement, and the prevalence of objectionable content. We believe that all of these goals are important to consider because they help us evaluate whether our users are having a positive experience with our apps. For instance, if changes made not only increased engagement but also problematic content, we would not consider that organization successful.

- b. Is Mr. Harris correct about the metrics that Facebook uses?**

Please see the response to your previous question.

- c. Please explain why Facebook chose the metrics that it uses.**

Our employees are evaluated on a number of goals and metrics as part of our holistic evaluation process. Many of these goals and metrics are designed to promote overall app health. For more information, please see response to your Question 1(a).

- d. Is Mr. Harris correct that Facebook experimented with non-engagement-based metrics? If so, please explain what was tried and why this was discontinued.**

Our employee evaluation process is holistic and does not turn on any single goal or metric. The employees who work on developing our apps have many goals related to their app’s overall

¹ Karen Hao, “How Facebook got addicted to spreading misinformation,” *MIT Technology Review* (Mar. 11, 2021), <https://www.technologyreview.com/2021/03/11/1020600/facebook-responsible-ai-misinformation/>.

health including safety, security, engagement, and the prevalence of objectionable content. For more information, please see the response to your Question 1(a).

- e. Does Facebook include integrity metrics in its performance and pay assessments for any of its teams, including the Integrity Team?**

Please see the response to your Question 1(a).

- f. To the extent there are tensions between topline goal metrics and integrity metrics in assessing business units or product features, how do you resolve them, and who makes those decisions?**

Our goaling and incentives process is set up with the goal of building the best app over the long term. Safety and security are an important part of our assessment of overall app health.

Question 2. The MIT Technology Review reported in March that Facebook does not have a team or initiative specifically focused on how its content ranking systems affect the spread of misinformation.² Facebook responded that this responsibility is distributed among teams with the relevant expertise.

- a. Please elaborate on this response, and specifically explain which teams share in this responsibility and how they coordinate.**

We address misinformation comprehensively—rather than treat it as a single problem that a single team can solve. We have worked to develop and expand our systems to reduce misinformation and promote reliable information.

This work is not centralized within one team, nor is this how we think about problem areas internally. The thousands of people reviewing content on our platforms 24 hours per day play an important role in these efforts, including reviewing and removing misinformation that violates our Community Standards. Our Central Integrity team, aided by our Engineering and Global Operations team, helps create technology and brings it to bear on this problem area. Our Policy team plays a critical role in deciding what exactly our rules in this area are and what we do or don't allow on our platforms. At any given moment, a large number of people are involved in identifying, labeling, or removing misinformation from our platforms, and they all depend on one another to do it at the scale at which we operate.

We also partner with more than 90 fact-checking organizations around the world who rate content in more than 60 languages.

- b. While Facebook has invested in its integrity teams, some have raised concern about the scope of integrity issues that these teams are charged with addressing and whether they have adequate staffing and resources, particularly if other teams are more oriented around growth. Please respond to this concern.**

² *Id.*

We have 40,000 people across Meta working on safety and security. As we have previously mentioned, because we embed integrity teams into various app groups, our integrity teams take safety and integrity considerations into account during app development, while also advising app teams on best practices. Many changes also go through an Integrity Review process, a cross-functional process by which proposed app changes are evaluated on integrity criteria prior to launch. This process allows us to identify and anticipate potential abuses and build in ways to mitigate them from the start. Engineers also look at a variety of metrics, including integrity metrics, as they do their work. New apps or features are reviewed against a set of integrity standards to help us provide a positive experience for our users.

Question 3. Reporting has indicated that one challenge for Facebook is the issue of “hyperactive users” or “super-inviter” who generate a lot of activity and often promote either low-value or violating content.³ This reporting has raised the concern that Facebook may be reluctant to address these users’ behavior because of the impacts that doing so would have on engagement and growth.

a. Please provide information regarding the concentration of violating content among such “hyperactive users.”

We want to help people see what they find most valuable. Our goal is not to keep people on the service for a particular length of time, and it is certainly not to give people low quality or the most provocative content. In fact, key parts of our systems are designed to do just the opposite. To that end, we reduce the distribution of many types of content as outlined in our [Content Distribution Guidelines](#). We want the content in Feed to be representative of the whole community, and not only individuals excessively sharing lots of content.

It is also against our policies for people to misrepresent themselves on Facebook, use fake accounts, artificially boost the popularity of content, or engage in behaviors designed to enable other violations under our Community Standards. We also look for and remove coordinated inauthentic behavior (CIB) campaigns that seek to manipulate public debate while relying on fake accounts across our platforms. Over the past five years, we’ve shared our findings about CIB we detect and remove from our platforms, including in our quarterly adversarial threat reports.

We seek to provide transparency about how we define and treat problematic and low-quality content. While the Community Standards make it clear what content is removed from Facebook because we don’t allow it on the platform, the [Content Distribution Guidelines](#) outline some of the types of content that receive reduced distribution in Feed. Our efforts to reduce problematic content in Feed are rooted in our commitment to responding to people’s direct feedback, incentivizing creators to invest in high-quality and accurate content, and fostering a safe community—and we continue to adjust and develop our guidelines in line with these values. A

³ See Craig Silverman, Ryan Mac & Jane Lytvynenko, “Facebook Knows It Was Used to Help Incite the Capitol Insurrection,” *BuzzFeed News* (Apr. 22, 2021), <https://www.buzzfeednews.com/article/craigsilverman/facebook-failed-stop-the-steal-insurrection>; Jeff Horwitz & Deepa Seetharaman, “Facebook Executives Shut Down Efforts to Make the Site Less Divisive” (May 26, 2020), <https://www.wsj.com/articles/facebook-knows-it-encourages-division-top-executives-nixed-solutions-11590507499>.

summary of updates we make to the Guidelines can be found on our [Changes, Corrections and Adjustments](#) page.

In addition, the [Widely Viewed Content Report](#) (“WVCR”) provides a quarterly overview of the content that reaches the most people on Facebook’s Feed in the United States. This report was created to both provide more transparency about the most-viewed organic content in Feed on Facebook in the US and to hold ourselves accountable to improve the quality of content on Facebook. Insights from the WVCR continue to inform how we evolve our apps and policies. This includes developing new policies, where needed, to address harmful or otherwise low-quality content, and making changes to ranking that have reduced problematic content from reaching the widest audience.

To address low-quality content, we have:

- Improved how we reduce highly viral engagement bait;
- Introduced spacing rules to prevent multiple posts identified as engagement bait from showing up consecutively in Feed;
- Improved our detection and enforcement systems to better identify and reduce posts with unrelated links in Feed;
- Improved systems to identify and block spammers on Facebook, as well as people engaged in inauthentic behavior who pretend to be based in one country while actually being in a different one.

As explained in the WVCR for the third quarter of 2022, this most recent quarter’s top content did not contain any violating content. There is one piece of content in the report that was removed by the LADBible Page.

We are cautiously optimistic of the progress we’ve made as we work to improve the quality of content within Facebook. We continue to rigorously work to understand the content ecosystem and evaluate the effectiveness of our policies and integrity measures—closing gaps as we find them.

b. Please describe Facebook’s efforts to address this issue, and in particular, please respond to the concerns raised in *BuzzFeed News* about “super-invited” associated with the “Stop the Steal” campaign and in the *Wall Street Journal* about the “Sparing Sharing” proposal that Facebook considered.

We worked hard to do our part in helping to protect the integrity of the 2020 election, and we’re proud of the work we’ve done to support our democracy. We followed the policies and processes that we laid out in advance to protect the democratic process in both the pre-and post-election periods, and we worked hard to apply those policies fairly and consistently. Securing the integrity of elections is an ongoing challenge for platforms, and we are committed to continuing to improve our systems.

Part of our comprehensive approach to supporting the integrity of the 2020 election included removing Pages, groups, and events that violated any of our policies, including calls for violence. Based on repeat violations of our policies, we also restricted some US-based accounts from creating live videos, events, groups, or pages. We also removed content containing the phrase

“stop the steal” from Facebook and Instagram under our Coordinating Harm policy. When it came to the “Stop the Steal” Group, we took down the Group within about 24 hours. We quickly removed the Group because it was organized around the delegitimization of the election process, and we saw worrying calls for violence from some members of the Group, in violation of our policies.

Building community is core to Facebook’s mission in bringing the world closer together, and while community has made its way into every corner of the app, Groups continue to be the central place where people go to do more together. It’s important to us that people can discover and engage safely with Facebook Groups so that they can connect with others around shared interests and life experiences, and we have taken significant steps toward that goal. For example, we remove Groups that represent QAnon, even if they contain no violent content. And we do not allow militarized social movements—such as militias or groups that support and organize violent acts amid protests—to have a presence on our platform. In addition, in 2020, we temporarily stopped recommending US civic or political Groups, and in early 2021, we announced that policy would be kept in place and expanded globally. We’ve instituted a recommendation waiting period for new Groups so that our systems can monitor the quality of the content in the Group before determining whether the Group should be recommended to people. And we limit the number of Group invites a person can send in a single day, which can help reduce the spread of harmful content from violating Groups.

We believe that Groups and members that violate our rules should have reduced privileges and reach, with restrictions getting more severe as they accrue more violations, until we remove them completely. When someone has repeated violations in Groups, we will block them from being able to post or comment for a period of time in *any* Group. They also won’t be able to invite others to any Groups, and won’t be able to create new Groups. These measures are intended to help slow down the reach of those looking to use our platform for harmful purposes and build on [existing restrictions](#) we’ve put in place.

There is always more to do to help keep Facebook Groups safe, and we will continue to build and invest to make sure people can rely on these places for connection and support.

Question 4. In Facebook’s responses to my questions for the record from November last year, Facebook confirmed that it was building a “parallel viral content review system to flag posts that may be going viral—no matter what type of content they are—as an additional safety net.” Facebook also pointed to a Crisis Assessment Dashboard to correlate spikes in hate speech or voter interference content in near real time. I appreciate Facebook’s innovations in this space.

- a. Please elaborate on the current status of these efforts, and in particular whether they are being implemented at scale.**
- b. Who decides where to set the threshold for virality in this system, and what factors are considered in making that decision?**
- c. Has Facebook found the parallel viral content review system to be effective? Please explain.**

Since 2016, we've built an advanced system combining people and technology to review the billions of pieces of content that are posted to our platform every day. State-of-the-art artificial intelligence systems detect and remove content that violates our policies, and flag for further review content that may violate our policies; people also report content to us they believe is questionable, and our own teams review content. When it comes to prioritizing content for human review, our review systems use technology to prioritize high-severity content with the potential for offline harm and viral content that is spreading quickly.

We built the Crisis Assessment Dashboard in preparation for the US 2020 elections, to surface spikes and trends of potentially harmful content. Following those elections, we further developed and expanded the system to support global elections and proactive risk monitoring. Our viral content review system prioritizes for review content that we predict, based on virality signals, may go viral.

Question 5. I appreciate Facebook's recognition of the importance of estimating the prevalence of violating content and how this is a key statistic for engineers focused on misinformation. Although Facebook reports a single global prevalence statistic for hate speech on a quarterly basis, that statistic is not broken out into more granular categories.

- a. Is Facebook able to estimate the prevalence of violating content on a more granular basis for internal purposes (e.g., as to the types of violations, the particular content sources, countries)?**

Our primary metric is the same internally as it is externally—prevalence. Prevalence tells us how often content that violates our standards is seen relative to the total amount of times any content is seen on Facebook. It is how we objectively evaluate our progress and reflects our efforts to reduce the distribution of content that violates our standards. Some pieces of content get many more views than others, so removing that content likely has a greater effect on how much violating content is actually seen by users. Meta's enforcement teams typically prioritize review of content that is likely to accumulate a high number of views, and therefore the removals Meta makes may have an outsized positive effect on reducing user encounters of violating content. We talk about prevalence in our Community Standards Enforcement Report every quarter and describe it in our Transparency Center. For example, in the third quarter of 2022, we reported that hate speech prevalence was about 0.02%, or two views for every 10,000 views on Facebook and Instagram.

While we understand the ask for more granular data, we believe that the way in which we report prevalence is the most useful metric, including because measuring data in a more granular way may often be less accurate.

- b. To what extent are Facebook's estimates of prevalence factoring into company decision making on new products and features? Please explain how so.**

We both build our apps and continually update them with safety and integrity in mind. This approach is built into the DNA of the company. We embed integrity teams into various app groups, our integrity teams take safety and integrity considerations into account during app

development, and they advise app teams on best practices. Many changes also go through an Integrity Review process, a cross-functional process by which proposed app changes are evaluated on integrity criteria including those that could impact prevalence prior to launch. This process allows us to identify and anticipate potential abuses and build in ways to mitigate them from the start.

Question 6. I appreciate your testimony that Facebook is committed to increasing transparency about your algorithmic recommendation engines. Many independent researchers have said it is essential to know not only what “signals” are considered but also how those signals are weighted; otherwise, it is not apparent which signals are actually driving the ranking. Reporting has also indicated that Facebook has developed certain higher-level “knobs and dials” such as “news ecosystem quality” that Facebook can adjust to modify the rankings of the News Feed, as it did in advance of the 2020 election.⁴

- a. Please provide more information about this news ecosystem quality metric, how it is calculated, and the effect that changing its weighting has on the News Feed.**

The measure referenced in your question was put in place on a temporary basis as part of our elections integrity efforts to help limit the spread of inaccurate claims about the US 2020 election and remains a temporary intervention we could deploy in the future, if needed. However, on a day-to-day basis as of February 2023, and as outlined in our [Content Distribution Guidelines](#), we seek to reduce certain types of content such as unoriginal news articles and fact-checked misinformation.

- b. What are the primary ways in which Facebook adjusted its weighting of factors in the News Feed algorithm in advance of events like the 2020 election and the Derek Chauvin verdict?**

In some unique circumstances, we may put a series of temporary measures in place where there are specific risks that spikes in activity on the platform could mean that the many systems we have in place to enforce our policies may not be able to keep up. For example, during the US 2020 election, we automatically removed potentially violating content at lower confidence levels than we normally would have even before they were reviewed by our team. We took these steps to respond to specific signals we were seeing on the platform, such as spikes in reported content—and turned some of them off responsibly and gradually as those signals returned to their previous levels.

Similarly, as the trial of Derek Chauvin drew to a close, we worked to prepare for the verdict by preventing online content from being linked to offline harm. In addition to removing content that violated our Community Standards, we sought to limit the spread of content that our systems predicted was likely to violate our Community Standards in the areas of hate speech, graphic violence, and violence and incitement, as we have done in emergency situations in the past. And

⁴ See Kevin Roose, Mike Isaac & Sheera Frenkel, “Facebook Struggles to Balance Civility and Growth,” *New York Times* (Nov. 24, 2020), <https://www.nytimes.com/2020/11/24/technology/facebook-election-misinformation.html>.

we took extra steps to limit misinformation. This included using several tools to help ensure that potential misinformation was flagged to our network of third-party fact-checking partners.

- c. Will Facebook provide greater transparency into how it weighs different inputs for its News Feed rankings, whether through reporting, granting access to independent researchers, or conducting an audit?**

We have made significant progress in providing greater transparency into [how the Facebook Feed ranking process works](#), what gets distributed, and why. For example, the quarterly Widely Viewed Content Report aims to provide more transparency and context about what people are seeing on Facebook by sharing the most-viewed domains, links, Pages and posts for a given quarter on Feed in the United States. We provide insights into the various content types that appear on Feed to help people better understand our distribution systems and how they influence the content people see on Facebook. We work to offer a variety of tools to help people understand what they're seeing on Facebook and why, like the "Why Am I Seeing This?" feature, which gives users information about why they're seeing a post or ad. We also publish posts that highlight major updates to Feed and explain the thinking behind them.

Recently, we announced new ways to customize what you see in your Facebook Feed so you can discover what's most relevant to you. Now, you can select "Show more" or "Show less" on posts from the people and communities you're connected to and posts that Facebook recommends to you. Selecting "Show more" will temporarily increase the [ranking score](#) for that post and posts like it. If you select "Show less," you'll temporarily decrease its ranking score.

For more technical information, our Engineering team has shared [details](#) of how we designed a machine learning powered Feed ranking system.

Question 7. You testified that Facebook has prioritized "meaningful social interactions" over engagement or time spent on the platform.

- a. How does Facebook measure "meaningful social interaction"?**
- b. Has Facebook analyzed the correlation of this metric with engagement or growth?**
- c. What is the practical impact of building products to prioritize meaningful social interactions?**

We want to help people see what they find most valuable, and, to that end, we have algorithms to help rank the content people follow and create their personalized feed. The goal of ranking is not to keep people on the service for a particular length of time. We reduce the distribution of many types of content—meaning that content appears lower in Feed—because they are sensational, misleading, gratuitously solicit engagement or are found to be false by our independent fact-checking partners.

In 2018, we publicly announced that we were making ranking changes that we thought would help people see more content that was meaningful to them—but that would nonetheless lead to people spending less total time on Facebook.

The prediction proved correct; at the time, the change led to a decrease of 50 million hours' worth of overall time spent on Facebook per day, and while we saw a loss of billions of dollars in the company's market cap, we viewed this as a success because it improved people's experience, and we think building good experiences is good for the business in the long term.

Question 8. During last Tuesday's hearing, I asked whether Facebook generally requires employees to sign non-disclosure agreements (NDAs). You said you did not know the answer but would follow up. Can you provide an answer here?

All employees are required to sign an agreement relating to intellectual property ownership, and as part of that agreement, employees also agree to maintain the confidentiality of trade secrets and other proprietary business information. Additionally, we typically include non-disclosure and other confidentiality provisions in our severance and settlement agreements, as is industry standard.

a. If Facebook does require employees to sign NDAs, please briefly describe the scope of those agreements and the types of information that employees and former employees are barred from disclosing.

Please see the above response. Our agreement relating to intellectual property ownership governs confidential information obtained in the course of employment at Meta, which could include proprietary technical, business, customer, employee information, and other non-public information about the company. Separate confidentiality agreements may also be used in the context of specific confidential team projects. No provisions in our confidentiality, severance, or settlement agreements prevent an employee from reporting violations of law, or talking to regulators or Congress about violations of law.

b. Do the NDAs cover more than trade secrets, private user data, and/or technical information that employees learn while employed by Facebook?

Please see the response to your Question 8(a).

c. Do the NDAs or any other contractual agreements bar employees from criticizing or disparaging Facebook after they leave the company?

Our severance and settlement agreements generally contain non-disparagement clauses, as is standard form and practice.

d. If an employee disagrees with a policy decision made by Facebook during the employee's period of employment, would the employee's NDA prohibit him or her from discussing that disagreement publicly after leaving the company?

Please see the responses to your Questions 8(a) and 8(c).

- e. **Do Facebook’s NDAs contain binding arbitration clauses—meaning employees must submit to arbitration in the event of any legal dispute under the terms of the agreement?**

As a general matter, employees sign a mutual arbitration agreement as a condition of employment, although there are carveouts for certain types of harassment and discrimination claims, among others.

- f. **Has Facebook ever threatened legal action against an employee (current or former) based on potential violations of that employee’s NDA?**

Based on a reasonable inquiry and review, we are aware of no instance of the company initiating a legal action against a current or former employee based on violations of a NDA.

Question 9. According to public reporting, some Facebook employees alleged that Facebook made special exceptions to its “repeat offender” policies to avoid taking measures against certain conservative pages, including Breitbart, because Facebook feared accusations of anti-conservative bias.⁵

- g. **According to Facebook’s description, “Pages and websites that repeatedly share false information will see their distribution reduced and their ability to advertise removed.”⁶ Can you provide any more details about this repeat offender policy—such as how many times a page or user must share false information before Facebook reduces the distribution of content from that page or user?**

Facebook is a platform for people to share information and ideas from a variety of viewpoints and across the political spectrum. We strive to make Facebook a place where people can access reliable information and discuss important issues, and we craft and enforce policies that support these goals.

In general, we do not think a private company like Meta should be deciding what is true or false, which is why we work with more than 90 fact-checking organizations who independently review and rate potential misinformation on our platform. Fact-checkers review content, check facts, and rate the accuracy of content. These organizations are all certified through the non-partisan International Fact-Checking Network and they decide what content to review and what rating to apply to that content. Please see [here](#) for more information on Facebook’s fact-checking policies.

If content is rated False, Altered, or Partly False by a fact-checker, we significantly reduce its distribution in Feed and reject it if it’s being run as an ad. And critically, people who see it, try to

⁵ Olivia Solon, “Sensitive to Claims of Bias, Facebook Relaxed Misinformation Rules for Conservative Pages,” NBC News (Aug. 7, 2020), <https://www.nbcnews.com/tech/tech-news/sensitive-claims-bias-facebook-relaxed-misinformation-rules-conservative-pages-n1236182>; Craig Silverman & Ryan Mac, “Facebook Fired an Employee Who Collected Evidence of Right-Wing Pages Getting Preferential Treatment,” BuzzFeed News (Aug. 6, 2020), <https://www.buzzfeednews.com/article/craigsilverman/facebook-zuckerberg-what-if-trump-disputes-election-results>.

⁶ Facebook, “How is Facebook addressing false information through independent fact-checkers?”, <https://www.facebook.com/help/1952307158131536>.

share it, or have already shared it will see warnings alerting them that it's been rated false by a fact-checker.

In order to more effectively fight misinformation, we may also take action against Pages, Groups, accounts, and domains that repeatedly share or publish content that is rated False or Altered, near-identical to what fact-checkers have debunked as False or Altered, and content we enforce against under our policies on COVID and vaccine misinformation. These actions will last for a subsequent 90 days.

If Pages, Groups, accounts or websites repeatedly share such content they will typically see their distribution reduced. They could also face other penalties that include: removal from recommendations surfaces; removal of the ability to monetize and advertise; removal of the ability to register as a news Page; or pop-ups that inform users the Page, Group, or account repeatedly shares misinformation before users follow or share content. Additionally, if the Page is a part of the News Tab, they may lose eligibility.

We believe that Page owners, like publishers, have a responsibility for the content they share with their audiences. When Pages do share something that has been marked false by a fact-checker, we notify them that their shared post will be demoted—regardless of whether they are the original publisher of that content.

It is important to keep in mind that bad actors intent on breaking our rules will continue to update their tactics—and even the terminology they use—to avoid detection. We therefore do not provide specific details on our policy enforcement to avoid bad actors using that information as a playbook to game the system, and we work to stay on top of emerging trends.

h. NBC News stated in its reporting that “The company operates on a ‘strike’ basis, meaning a page can post inaccurate information and receive a one-strike warning before the platform takes action. Two strikes in 90 days places an account into ‘repeat offender’ status.”⁷ Is that true?

In addition to removing content that violates the Facebook Community Standards when we become aware of it, we use a strike system to count violations, educate people on our standards, and hold people accountable when they post violating content. Whether we apply a strike depends on the severity of the content, the context in which it was shared, and when it was posted.

As of December 2022, for most violations, the first strike will result in a warning with no further restrictions. If Meta removes additional posts that go against the Facebook Community Standards in the future, we may apply additional strikes to the account, and the account may lose the ability to create content for set periods of time. Generally, the more strikes an account accrues, the longer the penalties get. More information about how we [count strikes](#) and [restrict accounts](#) is available in our public Transparency Center, and we have an ongoing review of our enforcement policies and protocols to see how we can continue to improve them.

⁷ Solon, *supra*.

As discussed above in response to your previous question, if a person shares content on Facebook that is rated False, Altered, or Partly False by one of our independent third-party fact-checking partners, we apply penalties to the content and, if done repeatedly, to their account, Page, Group, or website. These penalties are separate from strikes for violating our Community Standards.

- i. If employees and independent fact-checkers responsible for enforcing Facebook’s content policies determine that a page or user is a “repeat offender,” is that decision subject to review by anyone outside of Facebook’s content policy or integrity teams?**

Decisions about the consequences for repeatedly sharing misinformation are governed by our policies. While third-party fact checkers review and rate potential misinformation on our platform, we take action based on those ratings in accordance with our policies. We are constantly reviewing and re-evaluating those policies to help ensure that the content review process is objective.

- j. Is public perception ever taken into consideration when determining whether or how to enforce Facebook’s repeat offender policies? If so, how?**

As noted above, decisions about the consequences for repeatedly sharing misinformation are governed by our policies. We seek to write policies that clearly distinguish between violating and non-violating content and help ensure that the content review process is objective.

For more information about our enforcement penalties for sharing fact-checked content, please refer to our [Transparency Center](#).

- k. Are potential political repercussions ever taken into consideration when determining whether or how to enforce Facebook’s repeat offender policies? If so, how?**

Please see the response to your previous question.

- l. Are potential business consequences—such as potential dissatisfaction from advertisers—ever taken into consideration when determining whether or how to enforce Facebook’s repeat offender policies? If so, how?**

Please see the response to your Question 9(j).

Questions from Senator Hirono

Question 1. During the hearing, I asked you how Facebook ensures it does not violate civil rights laws by discriminating on the basis of characteristics such as age, gender, and race when it targets ads for housing, employment, and financial services. Without equivocation, you testified that Facebook does not allow the use of “some of the more sensitive targeting criteria” when targeting ads for “certain types of advertisements, such as financial services advertisements or housing advertisements.” Two days after the hearing, *The Markup*⁸ published an article in which it identified 91 unique ads promoting financial services products that explicitly discriminated on the basis of age. And recent litigation has shown that Facebook routinely serves non-credit financial services ads that are targeted based on age and gender.

Do you stand by your testimony that Facebook does not allow the use of “some of the more sensitive targeting criteria” when targeting ads for “certain types of advertisements, such as financial services advertisements or housing advertisements?” If not, please explain what in your testimony was inaccurate. If you contend your testimony was accurate, what is your explanation for the advertisements identified in the article in *The Markup* and pending litigation such as *Liapes v. Facebook, Inc.*?

Discrimination and discriminatory advertising have no place on Facebook, and we remove such content when we become aware of it. Our policies prohibit advertisers from discriminating against people based on personal attributes such as race, ethnicity, color, national origin, religion, age, sex, sexual orientation, gender identity, family status, disability, and medical or genetic conditions. It is also accurate that Meta does not allow the use of some categories of targeting criteria for certain types of advertisements, as explained in additional detail below.

We strongly believe that the most useful advertising experiences for consumers are personalized. Personalized ads enable people to discover products and services from small businesses that may not have the ability to market them on broadcast television or other forms of media. They also enable nonprofits, social causes, and organizations to reach the people most likely to support and benefit from them, such as connecting people to fundraisers for charitable causes they care about.

At the same time, we want to better match people’s evolving expectations of how advertisers may reach them on Facebook, and address feedback from civil rights experts, policymakers, and other stakeholders on the importance of preventing advertisers from abusing the targeting options we make available.

Facebook’s ads generally allow marketers to select audiences for their ads based on a variety of factors including age, gender, location, interests, and behavior. These audience selection (or targeting) tools are not available for ads for housing, employment, and credit opportunities (credit ads include certain types of financial services advertisements, but not all financial services ads).

⁸ <https://themarkup.org/citizen-browser/2021/04/29/credit-card-ads-were-targeted-by-age-violating-facebooks-anti-discrimination-policy>

We also announced specific changes to detailed targeting options that apply more broadly than just housing, employment, or credit ads. In early 2022, we removed detailed targeting options that relate to topics people may perceive as sensitive, such as options referencing causes, organizations, or public figures that relate to health, race or ethnicity, political affiliation, religion, or sexual orientation. Examples include:

- Health causes (e.g., “Lung cancer awareness,” “World Diabetes Day,” “Chemotherapy”)
- Sexual orientation (e.g., “same-sex marriage” and “LGBT culture”)
- Religious practices and groups (e.g., “Catholic Church” and “Jewish holidays”)
- Political beliefs, social issues, causes, organizations, and figures

Question 2. Facebook’s Help Center describes what factors it uses to target ads to users, including “[o]ther information about you from your Facebook account” such as “your age [and] your gender.” Facebook also allows ads to be targeted based on things like zip code that can be used as a proxy for race.

Does Facebook allow advertisers to use factors such as age, gender, and zip code to target ads for housing, employment, credit, and other financial services (such as insurance, bank accounts, and debit cards)? If it does, how does Facebook ensure it does not violate civil rights laws when targeting these ads to users?

As discussed above, Facebook’s ads generally allow marketers to select audiences for their ads based on a variety of factors including age, gender, location, interests, and behaviors. However, we have [made significant investments](#) to better prevent discrimination on our ads platform. We restrict how housing, employment, and credit advertisers [can create their target audiences](#), and we have disallowed their use of gender or age targeting, and require that location targeting have a minimum 15-mile radius.

Our policies prohibit advertisers from discriminating against people based on personal attributes such as race, ethnicity, color, national origin, religion, age, sex, sexual orientation, gender identity, family status, disability, and medical or genetic conditions. And our policies require all advertisers to comply with the laws in their jurisdiction, including not engaging in discriminatory practices.

Question 3. In 2019, Facebook entered into a settlement with the National Fair Housing Alliance and the Communications Workers of America in which it agreed to implement changes to its ad platform to prevent advertisers from targeting housing, employment, and credit ads based on age, gender, zip code (a proxy for race), and other factors.

- a.** What changes did Facebook make to its ad platform to ensure factors such as age, gender, and zip code could not be used to target ads for housing, employment, and credit?

Our policies have long prohibited discrimination, and we have made significant changes to help prevent advertisers from misusing our tools to discriminate in their ad targeting. As you note, as part of settlement agreements with civil rights organizations like the National Fair Housing Alliance, and based on input from civil rights experts, we have changed the way advertisers may

select the audience for housing, employment, and credit (“HEC”) ads. Specifically, we have removed age, gender, and zip code targeting options for HEC ads, and we have severely restricted the number of interest category targeting options available. We have also expanded our enforcement of these restrictions across the tools businesses use to buy ads. Additionally, we have added HEC sections to the [Ad Library](#), making it easy to search for and view all currently active US ads about housing, employment, and credit opportunities, regardless of the advertiser’s intended audience.

For more information about our ad delivery system, please see the response to your Question 4(a) below.

Facebook does not give advertisers the ability to target ads based on race. More broadly, we are following through on requests to look at the experiences people from historically and systemically marginalized communities have on our services and how our technologies may impact them. We have been meeting with the civil rights community, privacy experts, academics, regulators, and other organizations to determine the best way to assess how people of different races experience our services. For more information about technical methods and overall principles we are using to approach this work responsibly, please see our Newsroom post, available [here](#).

b. What additional steps has Facebook taken to ensure compliance with the settlement agreement?

Please see the response to your Question 3(a) above.

c. Do the ads identified by *The Markup* comply with the settlement agreement? If you contend they do, please explain why. If they do not, how did they get through Facebook’s system?

Our policies require credit card advertisers to identify themselves and use restricted targeting options. We have worked to review and remove ads from these businesses that ran in violation of this policy. Ads on Facebook are subject to our ad review system, which we use to check ads against our Advertising Standards. These policies are posted publicly online. We leverage automated review and use teams of reviewers to improve and train our automated systems, and in some cases, review specific ads.

Reviewing ads from millions of advertisers globally against our Advertising Standards is essential, but not without challenges. Feedback from our community of advertisers, global experts and people who use our platform everyday are an integral part of our ongoing efforts to protect the people who use our platform. We will continue to work alongside them as we regularly evaluate our policies, improve our enforcement, and stay accountable to our progress over time.

d. Does Facebook allow advertisers to use age, gender, and zip code to target non-credit financial services ads (e.g., insurance, bank accounts, debit cards)? If it does, why does Facebook consider it appropriate or desirable to allow discrimination for these types of non-credit financial services?

All ads on Facebook, including non-credit financial services ads, must comply with our Advertising Standards, which can be found at <https://transparency.fb.com/policies/ad-standards/>.

Our Advertising Standards prohibit all advertisers from using targeting options to discriminate against, harass, provoke, or disparage users or to engage in predatory advertising practices in any ads. Our policies prohibit advertisers from discriminating against people based on personal attributes such as race, ethnicity, color, national origin, religion, age, sex, sexual orientation, gender identity, family status, disability, and medical or genetic conditions. Further, our policies require all advertisers to comply with the laws in their jurisdiction, which includes not engaging in discriminatory practices.

Question 4. Even when advertisers do not expressly target ads for housing, employment, and financial services on the basis of age, gender, and race, Facebook’s ad delivery algorithm may still serve housing, employment, and financial services ads in a discriminatory way. For example, during the hearing, I referenced an article in *MIT Technology Review*⁹ that found “Facebook’s ad-delivery system shows different job ads to women and men even though the jobs require the same qualifications.”

- a. Does Facebook’s ad delivery algorithm rely on protected statuses such as age, gender, and race or well-known proxies for protected statuses to determine which users will receive ads on Facebook, including ads for employment, housing, and financial services?**

Meta embeds civil rights and responsible AI principles into our app development process to help advance our algorithmic fairness efforts while protecting privacy.

Ads are ranked for delivery using different sets of machine learning models that are constantly learning and changing. There are hundreds of models that use billions of data points. The machine learning models are only one part of the ads delivery process. Before ranked ads are delivered to a particular user, they compete against each other in an auction and against other content, such as posts from friends and family, for the limited space available in user Feeds and other placements. As a result, no single input is determinative of the delivery of a particular ad to a particular user, let alone of how an ad is ranked.

In June 2022, we announced a settlement with the US Department of Housing and Urban Development (“HUD”) and the US Department of Justice (“DOJ”) regarding the way we deliver housing ads to people residing in the US. As a part of our settlement with HUD and DOJ, we announced our plan to create the Variance Reduction System (“VRS”) to help advance the equitable distribution of ads on Meta technologies. We launched the VRS in the United States for housing ads and shared an update in January on our collaboration with the DOJ. Over the coming year, we will extend its use to US employment and credit ads. The VRS builds on our longstanding efforts to help protect against discrimination, including our previous work to restrict certain targeting options for campaigns that advertise housing, employment, or credit ads. For example, we don’t allow advertisers that are either based in or trying to reach people in the

⁹ <https://www.technologyreview.com/2021/04/09/1022217/facebook-ad-algorithm-sex-discrimination/>

US, Canada, and certain European countries from targeting their housing, employment, or credit ads based on age, gender, or ZIP code.

Additionally, as a part of this settlement, we sunsetted Special Ad Audiences, a tool that lets advertisers expand their audiences for ad campaigns related to housing. We also chose to sunset this for employment and credit ads. In 2019, in addition to eliminating certain targeting options for housing, employment, and credit ads, we introduced Special Ad Audiences as an alternative to Lookalike Audiences. But the field of fairness in machine learning is a dynamic and evolving one, and Special Ad Audiences was an early way to address concerns.

We believe this type of work is unprecedented in the advertising industry and represents a significant technological advancement for how machine learning is used to deliver personalized ads. We are excited to pioneer this effort, but given the complexity and technical challenges involved, it will take some time to test and implement. We will also engage with external stakeholders throughout the development process so we can hear and integrate their feedback.

We will continue to make this work a priority as we collaborate with stakeholders to support important industry-wide discussions around how to make progress toward more fair and equitable digital advertising. Meta is committed to being transparent about its ad delivery process; we publish information about it, including in our Business Help Center. For more information, please visit: <https://www.facebook.com/business/news/good-questions-real-answers-how-does-facebook-use-machine-learning-to-deliver-ads>.

b. What steps does Facebook take to ensure its ad delivery algorithm does not result in a disparate impact based on protected statuses such as age, gender, and race?

Please see the response to your Question 4(a) above.

c. How often, if ever, does Facebook audit housing, employment, and financial services ads served by its ad platform to ensure Facebook's ad delivery algorithm does not result in a disparate impact based on protected statuses such as age, gender, and race? If Facebook has performed such audits, please provide the results, including any audits or reports on audits that Facebook has provided to any federal or state agencies, including the Department of Justice or the Department of Housing and Urban Development. If Facebook has not performed such audits, will it commit to doing so and sharing the results with this Subcommittee?

Please see the response to your Question 4(a) above for information on our recent collaboration with HUD and DOJ regarding the way we deliver housing, employment, and credit ads to people residing in the US.

d. Has Facebook considered any methods for reducing or eliminating the extent to which its ad delivery algorithm results in a disparate impact based on protected statuses such as age, gender, and race, including in any discussions with federal, state, or local agencies? If so, what are those methods.

Please see the responses to your Questions 4(a) and 4(c) above.

Question 5. Last fall, Jack Dorsey testified before the full Senate Judiciary Committee that “people should have choices about the key algorithms that affect their experience online.” Following the hearing, I asked Mark Zuckerberg in a Question for the Record whether he agreed with Mr. Dorsey’s statement and if he would commit Facebook to allowing users to choose from a selection of curation algorithms. He failed to answer the question and instead described how Facebook frequently makes changes to the algorithms that drive Facebook’s News Feed ranking “in an effort to improve people’s experience on Facebook.”

- a. **Will Facebook commit to allowing users and/or advertisers to choose the algorithm that ranks and filters the content they see on the platform, including a choice of algorithms developed by third parties? If not, why not?**

We want to provide people with tools to help them manage their experiences on our platforms. For example, on Facebook, we’ve made it easier to sort and browse Feed, giving people more control over what they see. We also launched Favorites, a tool where people can control and prioritize posts from the friends and Pages they care about most in their Feed. Specifically, people can select up to 30 friends and Pages to include in Favorites, and posts from these selections will appear higher in ranked Feed and can also be viewed as a separate filter.

We also give people ways to tell us that they would rather not see ads based on their interests or on their activities on other websites and apps, such as through controls within our ad settings. We believe it is important to provide people with controls to help them determine the ads they see. For example, we allow people to update their ad preferences so they can take more control over what information we use when deciding what ads to show them. People can opt out of seeing ads based on data from partners, and can also hide or report any ad from any advertiser with a few taps. We also have a tool called “Why Am I Seeing This?,” which shows people information about targeting options the advertiser chose to reach them. It lets them tap on posts and ads in Feed, get context on why they’re appearing, and take action to further personalize what they see. Finally, we have a tool called “Off-Facebook Activity,” which lets people see a summary of the information we received from business partners and clear it from their account if they want to.

In addition, we have several mechanisms for advertisers and Page admins to control the audience eligible to view the content they produce. When an advertiser decides to create an ad, we provide age and location targeting options during the ad creation process. The advertiser must comply with our Advertising Standards and any applicable local laws, and they can do so, for example, by specifying that their ads be shown only to users that meet a minimum age or are located in a specific country. Page admins can also use age restrictions to limit the audience of their Page.

- b. **Will Facebook commit to giving advertisers the ability to deliver their ads with an algorithm that does not rely on protected statuses or proxies for them to target their ads, including ads for employment, housing and financial services?**

Please see the response to your Question 5(a) above.

Question 6. Facebook’s independent Oversight Board recently began accepting appeals of content left up on the platform.

Does Facebook commit to expanding the Oversight Board’s jurisdiction to include reviews of Facebook’s algorithm for ranking and filtering content? If so, when will the Oversight Board be given this authority? If not, why not?

Ultimately, the goal of the Oversight Board is to play an important role in Facebook and Instagram’s overall governance, and the Board’s charter and bylaws are set up for the scope of the Board to grow and evolve over time, including expanding to review cases involving Groups, Pages, profiles, Events, and advertisements. For example, the scope of the Board recently expanded to include the ability to issue binding judgment on cases involving applying a warning screen to photo and video content. Now, if the Board determines that content should be restored or remain on our platforms, it can also issue a binding judgment about whether that content qualifies for the application or removal of a warning screen. We see the use of screens like these as an important tool for nuanced content moderation and a valuable opportunity to seek expanded Board guidance. When it comes to algorithmic design, we currently have additional partnerships and tools in place to provide checks and transparency on this aspect of our technology. We work to help establish best practices for AI systems across the industry and the broader policy community through external collaborations, such as our founding membership in the Partnership on AI and our participation in the expert group that helped formulate the Organisation for Economic Co-operation and Development (“OECD”) AI Principles. To help define what it means to implement these principles in practice, we worked with Business at OECD and other member companies (AXA, Amazon Web Services, IBM, Microsoft, NEC, and PricewaterhouseCoopers) on in-depth business case studies highlighting how industry is operationalizing the OECD AI principles. The resulting [report](#), published in July 2022, includes best practices and key challenges in using tools to implement the AI principles. In addition, we have introduced new tools and processes to further improve the transparency of our machine-learning systems. For example, the “Why Am I Seeing This?” feature explains some of the reasons our automated systems show a user particular content.

Question 7. Facebook has claimed that it has Section 230 immunity when Facebook publishes ads that exclude people based on age, gender, and race and when it relies on protected statuses in its own ad delivery algorithm to decide which users receive ads, including ads for employment, housing, and financial services.

Why does Facebook believe it should receive Section 230 immunity for structuring its own platform in a discriminatory way and for taking its own actions to discriminate through its ad delivery algorithm even in cases where advertisers do not want Facebook to engage in such discrimination on their behalf?

As explained above, we prohibit advertisers from discriminating against people based on personal attributes such as race, ethnicity, color, national origin, religion, age, sex, sexual orientation, gender identity, family status, disability, and medical or genetic conditions. Our Advertising Standards prohibit advertisers from using targeting options to discriminate against, harass, provoke, or disparage users, or to engage in predatory advertising practices. We have HEC sections to the Ad Library, making it easy to search for and view all currently active US ads about housing, employment, and credit opportunities, regardless of the advertiser’s intended audience.

For additional information, including on the targeting restrictions we have added to our HEC ad setup process and the technical mitigations we are implementing to help ensure the audience that ends up seeing a housing ad more closely reflects the eligible targeted audience for that ad, please see the responses to your Questions 1 and 3(a) above.

Questions from Senator Sasse

Question 1. How does your firm define the ideal user experience? Does your answer differ between adult and children?

As discussed in more detail in the responses below, we work to put in place the right policies and features to help people who use our apps have an age-appropriate and positive experience. As part of these efforts, we developed a process to help us apply the UN’s Convention on the Rights of the Child directly to the features and experiences we build at Meta. We complemented our own internal surveys and analysis with input from global data protection regulators to create [Meta’s Best Interests of the Child Framework](#), which distills the “best interests of the child” standard into six key considerations that teams can consult throughout the development process.

Question 2. From your firm’s perspective, what is the ideal amount of time for a user to spend on the platform over a specified period of time? How much is too much?

There are many challenges with conducting research in this space, and we are not aware of a consensus among studies or experts about how much screen time is “too much.” Many experts and research studies suggest it’s not necessarily about how much time you spend on social media but more about what you’re doing and the experiences you’re having that’s important. Meta has implemented new features designed to give people more control over the time they spend online. For example, we recently launched [Quiet Mode](#) on Instagram, which, once enabled, turns off notifications and sends an auto-reply when someone sends a direct message. Anyone can use Quiet mode, but we’ll prompt teens to do so when they spend a specific amount of time on Instagram late at night. Indeed, and specifically with regards to Instagram, Meta has invested considerable effort in developing additional tools designed to give young users more control over their time including Take a Break, Daily Time Limit, You’re All Caught Up, Notification Controls, and more.

Question 3. If an individual pulled up her account on your platforms, could somebody at your company walk her through the exact calculation behind why the first post ended up where it is? And if not, what are the limitations, and how close could you get to a full explanation?

Like most platforms, Facebook and Instagram use many different algorithms for various app features, including to help rank the posts people see in Feed. Our Engineering team has shared details of how we designed a machine learning powered Feed ranking system [here](#). We summarize the different types of signals we use to rank content in Feed in our [Transparency Center](#). We have made significant progress in providing greater transparency into [how the Facebook Feed ranking process works](#), what gets distributed, and why.

Question 4. If a family member asked you for advice on whether and under what conditions the family member should allow her child to use your firm’s social media platform, what advice would you offer? How would it change based on age?

We are committed to working with parents and families, as well as experts in child development, online safety, and children’s health and media, to build appropriate tools and features for families. That includes building tools to help people manage their time on our platforms. It also means providing information, resources, and tools for parents and teens to work together to develop healthy and safe online habits. And it means continued learning in this area.

We’ve developed more than 30 tools to support teens and families, including developing supervision tools that allow parents and guardians to ask teens to enable them to be more involved in their teens’ experiences. In March 2022, we rolled out supervision tools on Instagram that, when turned on, allow parents and guardians to: (1) view how much time their teens spend on Instagram and set time limits; (2) be notified when their teen shares they’ve reported someone; and (3) view and receive updates on what accounts their teens follow and the accounts that follow their teens. In June 2022, we introduced additional features that allow parents and guardians to send invitations to their teens to initiate supervision tools, set specific times when they would like to limit their teen’s use of Instagram, and see more information when their teen reports an account or post, including who was reported, and the type of report. We make video tutorials on how to use these supervision tools available for parents and guardians in the Family Center’s education hub.

Question 5. Do you think that there is anything fundamentally wrong with the business mode of your firm along the lines of the criticisms outlined by Dr. Donovan and Mr. Harris?

We understand the responsibility that comes with operating a global tech platform, and we take it very seriously. Our goal is to help people see what they find most valuable. It is not to keep people on the service for a particular length of time, and it’s certainly not to give people the most provocative or enraging content. In fact, key parts of those systems are designed to do just the opposite. We reduce the distribution of many types of content as outlined in our [Content Distribution Guidelines](#).

The reality is that it is not in Meta’s interest—financially or reputationally—to push people towards increasingly extreme content. The company’s long-term growth will be best served if people continue to use and value its services for years to come. If we prioritized trying to keep a person online for a few extra minutes, but in doing so lead to a bad experience and made that person less likely to return in the future, it would be self-defeating.

Question 6. If somebody like Dr. Donovan who spends her time researching online extremists has her feed filled with extremist content, should we consider that to be a job well done in optimizing her user experience or something that should worry us?

At Meta, we stand firmly against the exploitation of social media by those committed to inciting violence and hate. That’s why, under the Facebook Community Standards, we prohibit hate speech, terrorism, and other harmful content. We employ tens of thousands of people and use industry-leading technology, including advanced artificial intelligence to enforce these rules—finding the majority of the content we remove before people even see it. We regularly publish transparency reports so people can see how we’re doing over time and how we compare to other

internet platforms. We are proud of our work in this space—but we are always working to improve and to share our lessons learned with other companies in the US and around the world.

Our Violence and Incitement policy prohibits content calling for or advocating violence, and we ban individuals and organizations that proclaim a violent mission under our Dangerous Individuals and Organizations policy. In August 2020, we expanded this policy further to address militarized social movements, such as certain militias, and violence-inducing conspiracy networks, like QAnon. We remove content that represents, praises, or supports those dangerous groups. We work tirelessly to enforce these policies. We have designated more than 270 white supremacist organizations, and we took action on 1.2 million pieces of content from Facebook tied to organized hate globally in the third quarter of 2022, over 94% of which we found before someone reported it.

We have cross-functional experts focused on these issues at Meta, including highly trained professionals who work exclusively or primarily to prevent terrorist and violent extremist content from appearing on our platform, to quickly identify and remove it if it does, and to track it as it evolves. These professionals possess expertise ranging from law enforcement and national security experience to counterterrorism intelligence and academic studies in radicalization.

We are proud that we have more than 40,000 people working on safety and security issues, and our efforts are making a difference. For example, as of the third quarter of 2022, hate speech represents only about 0.02% of content views, or around 2 views per every 10,000. Our advanced artificial intelligence systems have also improved at keeping people safe on our platform by proactively removing content that violates our standards.

Our efforts to combat extremism and hate do not end with our policies and enforcement efforts. We also have a number of programs to direct people to content or organizations to help them disengage from dangerous or violent movements. When people search for terms related to QAnon on Facebook and Instagram, we redirect them to credible resources from the Global Network on Extremism and Technology, the academic research network of the Global Internet Forum to Counter Terrorism (“GIFCT”).

We also recognize the importance of cross-industry collaboration to counter terrorism effectively. Meta is a founding member of GIFCT and we work closely with other GIFCT members to stop terrorist content from spreading online and to help smaller companies with fewer resources.

Question 7. If a user’s behavior indicates that he or she consistently chooses to focus on content that makes them angry or anxious or depressed, what is the responsibility of the platform in terms of curating the content that it places in front of that user?

Our goal is to encourage expression and help create a positive environment on Facebook and Instagram. The reality is that it’s not in Meta’s interest—financially or reputationally—to push people toward increasingly extreme content. The company’s long-term growth will be best served if people continue to use and value its services for years to come. If we prioritized trying

to keep a person online for a few extra minutes, but in doing so led to a bad experience and made that person less likely to return in the future, it would be self-defeating.

Question 8. Do you measure whether content that stimulates a reaction of anger elicits more engagement from users? If so, what are the results, and how are those results used to shape the platform?

Our systems are not designed to reward provocative content. In fact, key parts of those systems are designed to do just the opposite, and we are committed to enhancing user choice and transparency. We believe a better understanding of the relationship between people and the algorithms is in everyone's interest, and that anyone who uses our platforms should have more control over the content that they see. That's why we rolled out a chronological Facebook Feed years ago, which turns off algorithmic ranking for anyone who wants that instead, as well as tools such as Favorites and "Why Am I Seeing This." Recently, we announced additional new ways to customize what you see in your Facebook Feed so you can discover what's most relevant to you. Now, you can select "Show more" or "Show less" on posts from the people and communities you're connected to and posts that Facebook recommends to you. We have also expanded on our work to use direct feedback from people who use Facebook to understand the content people find most valuable and will continue to incorporate this feedback into our Feed ranking process.

The company's long-term growth will be best served if people continue to use and value its services for years to come. If we prioritized trying to keep people online for a few extra minutes, but in doing so made them unhappy or angry and less likely to return in the future, it would be self-defeating.

Question 9. In the hearing, I noted that political debate about regulating social media platforms often fails to respect the distinctions between the First Amendment limitations of government actors and the greater freedom enjoyed by private actors. While your firm is not a government entity, we should not overlook the history of cooperation between social media platforms and federal, state, and local governments in some content decisions. Many of these efforts, such as limiting access to terrorist propaganda, have long earned bipartisan support. Nevertheless, learning more about the role of government actors in your content decisions would be a significant and helpful step in informing the ongoing debate. Please give an overview of your partnerships with federal, state, and local governments in making all aspects of content decisions (e.g., developing standards, removing content that violates these standards, deprioritizing or demonetizing content, flagging content as potentially inaccurate or misleading or providing links to authoritative resources, promoting certain kinds of content).

Facebook is a platform for ideas across the political spectrum. We moderate content on Facebook according to our policies to help keep people on the platform safe, reduce objectionable content, and help ensure people participate on the platform responsibly. We seek to be clear and transparent about what our standards are and to apply them consistently.

As a general matter, we seek to be transparent about the various ways in which we engage with

government agencies and law enforcement. And we regularly release transparency reports that provide information on our responses to government requests for user data (<https://transparency.fb.com/data/government-data-requests/>), content restrictions based on local law (<https://transparency.fb.com/data/content-restrictions/>), and our enforcement of our Community Standards (<https://transparency.fb.com/data/community-standards-enforcement/>).

Question 10. At the federal level, to what extent is this liaising with the federal government done by political appointees instead of career officials? Does contact occur at the agency level or the White House level? Please characterize the nature of this contact.

Please see the response to your previous question.

Question 11. How often do you receive communications from legislators urging more aggressive content moderation?

Please see the response to your Question 9.

Question 12. Has your firm ever been the subject of threats of negative consequences by federal officials for refusing to take certain content moderation steps?

Please see the response to your Question 9.

Question 13. When the changes announced as part of the “Meaningful Social Interactions” initiative were rolled out, how much did traffic to Facebook decrease?

In 2018, we publicly announced that we were making ranking changes that we thought would help people see more content that was meaningful to them—but that would nonetheless lead to people spending less total time on Facebook. The prediction proved correct; at the time, the change led to a decrease of 50 million hours’ worth of overall time spent on Facebook per day, and while we saw a loss of billions of dollars in the company’s market cap, we viewed this as a success because it improved people’s experience, and we think building good experiences is good for the business in the long term.

Question 14. What is Facebook’s current thinking in regards to the potential for creating an ad-free, subscription-based version of the platform?

To build secure apps with extensive infrastructure that connect people wherever they live or whatever their income level, we need to make sure everyone can afford them. Advertising lets us keep our apps free, which makes them more accessible to everyone.