

**HEARING BEFORE
THE UNITED STATES SENATE COMMITTEE ON
HOMELAND SECURITY AND GOVERNMENTAL AFFAIRS**

September 14, 2022

Testimony of Chris Cox
Chief Product Officer, Meta

I. Introduction

Chairman Peters, Ranking Member Portman, and distinguished members of the Committee, thank you for the opportunity to appear before you today. My name is Chris Cox. I am Meta's Chief Product Officer, and I oversee our apps and privacy teams.

I first joined Meta (then Facebook) in 2005 as one of the company's first 15 software engineers. I care deeply about the work we do to help people connect with the things and people they care the most about. It is important to us that we help people feel safe on our family of apps, and 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. We regularly publish transparency reports so people can see how we're doing over time and how we compare to other internet platforms. We're proud of our work in this space—but we're always working to improve and to share our lessons learned with other companies in the US and around the world.

II. Combating Hate and Dangerous Organizations

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 groups.

We work tirelessly to enforce these policies. We've designated more than 1,000 militarized social movements and 270 white supremacist organizations, and we removed 2.3 million pieces of content from Facebook tied to organized hate globally in the second quarter of 2022, nearly 97% of which we found before someone reported it.

We have a team of cross-functional experts focused on these issues at Meta, including more than 350 highly trained professionals who work exclusively or primarily to prevent terrorist and violent extremist content from appearing on our platform and to quickly identify and remove it if it does. These professionals possess expertise ranging from law enforcement and national security experience to counterterrorism intelligence and academic studies in radicalization.

We invested around \$5 billion on safety and security last year alone—more than any other tech company, even adjusted for scale. We are proud that we have over 40,000 people working on safety and security issues, and our efforts are making a difference. For example, we’ve more than halved the amount of hate speech people see on Facebook over the last eighteen months. Hate speech now 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. We found more than 95% of the hate speech we removed before anyone reported it to us—up from just 23% a few years ago.

When it comes to advertisements that run on Facebook and Instagram, they must comply with the same Community Standards and Community Guidelines that apply to other content people post, and we place additional requirements on advertisers to help further protect people from poor experiences. For example, our Ads Policies prohibit predatory advertising practices and using ad targeting to discriminate against, harass, provoke, or disparage users. We have an ad review system that is designed to review all ads before they go live. If a violation is found at any point in the review process, the ad will be rejected.

Our efforts to combat extremism and hate don’t 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 (GNET), the academic research network of the Global Internet Forum to Counter Terrorism (GIFCT).

We are proud to have co-launched GIFCT five years ago to help fight terrorism and extremism online alongside our peers at other technology companies, including Google, YouTube, Microsoft, and Twitter. GIFCT was the culmination of years of informal partnerships among our companies on these issues, and the group has created a shared industry database for violent terrorist imagery that we have removed from our services. Sharing information allows all of us to identify more quickly and more accurately potential extremist content on our respective platforms. Most importantly, it also allows smaller companies the ability to take advantage of our technology and tactics, even with fewer people and resources. We believe that terrorism and extremism are shared problems that require shared solutions, and we encourage all tech companies to continue to partner with us in our efforts to keep such violence and hate off of online communities.

III. Using Algorithms to Keep People Safe and Improve their Experience

Like most platforms, Facebook and Instagram use many different algorithms for various app features, including to enable our search function and to help enforce our policies. For example, we use algorithms to identify and help remove content that violates our policies, including hate speech, incitement, and terrorism. This often happens before anyone reports content to us, sometimes even at the point of creation.

In addition to safety and enforcement, we also have algorithms to help rank the content people follow and create their personalized Feed. Oftentimes when people refer to Facebook’s “algorithm,” they are talking about this sort of content ranking algorithm. This personalization is important because people often have thousands of posts in their Facebook Feed each day from their friends, Pages they choose to follow, and Groups they choose to join. Most people don’t have time to look at all of this content every day, so the ranking process helps sort it and put the things they will find most valuable toward the top of their Feed.

I want to stress that the goal of ranking 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—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; the change led to a decrease of 50 million hours’ worth of time spent on Facebook per day, and we saw a loss of billions of dollars in the company’s market cap. We view this as a success because it improved the experience of our users, and we think building good experiences is good for the business in the long term.

Transparency and choice are also important to giving people the best possible experience. We have made significant progress over the past several years in providing greater transparency into how the Feed ranking process works, what gets distributed and why. Notably, the “Why Am I Seeing This” feature lets people understand why a particular post shows up where it does—and to change their settings easily if something’s not right. And we have published Content Distribution Guidelines, which explain why content might receive reduced distribution. Recently, we also made it easier for people to switch between, for example, seeing the newest content first or seeing only content from people they’ve chosen to be a “favorite.”

IV. Incorporating Safety into Our Products by Design

Safety and integrity are key to the product experience, and we build and update our products with these values in mind. We embed teams focusing specifically on safety and integrity directly into product development teams across the company, allowing us to address potential issues during product development. And individual product teams use our integrity tools to build in preventative safeguards at the start.

Before launch, potential new products or product changes are reviewed against a set of integrity standards to help us provide a positive experience for users. The process also allows us to identify and anticipate potential abuses and build in mitigations from the start. This process helps us build-in effective privacy, security, and safety protections before a product launches.

After products launch, we continue to monitor their impact, including by looking at integrity

metrics, to ensure products are best serving our community. Some of these same integrity metrics are released quarterly to the public, as part of our Community Standards Enforcement Report, which provides data on how Meta enforces its policies.

Abuse of our products isn't static—and neither is the way we approach our integrity work. We're continuing to evolve how we approach integrity, embedding integrity teams with product teams across Facebook, Instagram, Messenger, WhatsApp, and the teams that are going to build the metaverse in the years to come. And we are committed to continuing to improve so that we can help keep people safe.

V. Conclusion

At the end of the day, our job is to build the best product for people, and that's a product that is reliable, fast, secure, relevant, and safe—a product that connects people to content relevant to their interests and connects them to their family and their friends. That's the product people want and the product we wake up every day trying to build.

We appreciate your attention to these important issues and look forward to continuing to work with your offices to find ways we can continue to improve our products, processes, and partnerships.

Thank you, and I look forward to your questions.