<center>**Hearing Before the United States Senate**
Committee on Commerce, Science, and Transportation
Subcommittee on Consumer Protection, Product Safety, and Data Security

Testimony of Adam Mosseri
Head of Instagram, Meta Platforms Inc.

December 8, 2021</center>

I.   Introduction

Chairman Blumenthal, Ranking Member Blackburn, and members of the Subcommittee, my name is Adam Mosseri, and I have served as the Head of Instagram since 2018. Over the last few months, this Subcommittee has held a number of hearings about the safety and well-being of young people online. This is a critically important topic, and it is something that we think about—and work on—every day at Instagram.

Our mission at Instagram is to bring people closer to the people and things they love. Our platform began a decade ago with a few million users. Today, we proudly serve well over a billion people. While our platform began as a simple photo-sharing app, we have evolved to provide new ways for people to express themselves, including Stories, Reels, and Live. Teens use our app every day to spend time with the people they care about, explore their interests, and express themselves. They are doing incredible things on our platform, and I firmly believe that Instagram can be a force for good in the lives of young people.

Much has been said recently about Instagram and its impact on young people. As a parent and as the Head of Instagram, this is an issue I care deeply about. It's an area our company has been focused on for many years, and I'm proud of our work to help keep young people safe, to support young people who are struggling, and to empower parents with tools to help their teenagers develop healthy and safe online habits.

I hope we can work together—across industry and government—to raise the standards across the internet and better serve young people. The reality is that keeping young people safe online is not just about one company. An external survey from just last month suggested that more US teens are using TikTok and YouTube than Instagram.[1] With teens using multiple platforms, it is critical that we address youth online safety as an industry challenge and develop industry-wide solutions and standards.

---

[1] Mike Prouix, Weekly Usage of TikTok Surpasses Instagram Among US Gen Z Youth, FORRESTER (Nov. 18, 2021), https://www.forrester.com/blogs/weekly-usage-of-tiktok-surpasses-instagram-among-us-gen-z-youth/.

<center>1</center>

II.   Keeping Young People Safe on Instagram

As Head of Instagram, I am especially focused on the safety of the youngest people who use our services. This work includes keeping underage users off our platform, designing age-appropriate experiences for people ages 13 to 18, and building parental controls.

Age Verification on Instagram

Instagram is built for people 13 and older. If a child is under the age of 13, they are not permitted on Instagram. When we learn someone underage has created an account, we remove them. In fact, in the third quarter of this year, we removed over 850,000 accounts on Instagram that were unable to demonstrate that they meet our minimum age requirement.

Understanding people's age on the internet is a complex and industry-wide challenge—especially considering that many young people in the US do not have a driver's license until they are 15 or 16 years old. However, we're building new technology to proactively find and remove accounts belonging to those under 13 and to identify those people who may be under the age of 18.

In addition to requiring people to share their date of birth when they register and allowing anyone to report a suspected underage account, we train our technology to identify if people are above or below 18 using multiple signals. We look at things like wishing people a happy birthday and the age written in those messages—for example, "Happy 21st Bday!" or "Happy Quinceañera." This technology isn't perfect, and we're always working to improve it, but that's why it's important that we use it alongside many other signals to understand people's ages.

There is more that we can do as an industry to ensure that there are clear standards of age verification across apps. For instance, I think it would be much more effective to solve the problem at the phone level so that a young person using a phone has an age-appropriate experience across any of the apps that they use on that device.

Keeping Instagram Safe

Understanding age is important so that we can create a more age-appropriate version of Instagram for the youngest people on our platform. We've put in place multiple protections to create safe and age-appropriate experiences for people between the ages of 13 and 18.

Wherever we can, we want to stop young people from hearing from adults they don't know or that they don't want to hear from. We believe accounts that offer people more control about who can see and respond to their content are the best way to prevent this from happening, and we recently announced that everyone who is under 16 years old in the US is defaulted into what is called a private account when they join Instagram. For young people who already have a public account on Instagram, we are sharing a notification highlighting the benefits of a private account and explaining how to change their privacy settings.

Private accounts let people control who sees or responds to their content. If a young person has a private account, people have to follow them to see their posts, Stories, and Reels, unless they choose to allow others to re-share their content. We're also—by default—eliminating the ability for young people to be tagged or mentioned by others or to have their content included in Reels Remixes or Guides. Additionally, people can't comment on their content in those places, and they won't see the young person's content at all in places like Explore or through hashtags.

Encouraging young people to have private accounts is important when it comes to stopping unwanted contact from adults. But we've gone even further to make young people's accounts difficult to find for certain adults. We developed technology that allows us to find accounts that have shown potentially suspicious behavior—for example, an adult account that might already have been blocked by another young person—and to stop those accounts from interacting with young people's accounts. Using this technology, we don't show young people's accounts in Explore, Reels, or 'Accounts Suggested For You' to these adults. If they find young people's accounts by searching for their usernames, they are not able to follow them. They also are not able to see comments from young people on other people's posts nor are they able to leave comments on young people's posts.

Additionally, we've launched a number of tools to restrict direct messaging between teens and adults and to prompt teens to be more cautious about interactions in direct messaging. To protect teens from unwanted contact from adults, we introduced a new feature that prevents adults from sending messages to people under 18 who don't follow them. For instance, when an adult tries to message a teen who doesn't follow them, they receive a notification that says that sending a Direct Message isn't an option.

In addition to preventing conversations between adults and teens who don't follow one another, we started using prompts—or safety notices—to encourage teens to be cautious in conversations with adults they're already connected to. These safety notices alert young people when an adult who has been exhibiting potentially suspicious behavior is interacting with them. For example, if an adult is sending a large amount of friend or message requests to people under 18, we use this tool to alert the recipients and give them an option to end the conversation, or block, report, or restrict the adult.

Our work to create age-appropriate experiences for teenagers on Facebook and Instagram also includes age gating certain content, prohibiting certain types of ads from being served to minors, and limiting options for serving any ads to these users.

We've always had rules about the kinds of content we suggest to people in places like the Explore tab. These rules apply to everyone, but we're going to go a step further for young people. We're developing a new experience that will raise the bar even higher for what we recommend for them in Search, Explore, hashtags, and suggested accounts. This new experience will make it harder for young people to find potentially sensitive content on Instagram.

We're also optimistic about using nudges to point people towards different topics. External experts have suggested that, if people are dwelling on one topic for a while, it could be helpful to nudge them towards other topics.[2][3] That's why we're building a new experience that will nudge people towards other topics if they've been spending time on one topic for a while.

When it comes to advertising, we've long restricted certain kinds of ads from being served to minors, and we recently limited advertisers' options for serving ads to people under 18. Now, advertisers can only serve ads to people under 18 based on age, gender, and location but not interests or activity. This means that previously available targeting options, like those based on interests or on their activity on other apps and websites, are no longer available to advertisers.

Supporting Teens Who May Be Struggling

In addition to making sure young people are safe on Instagram, we believe it's important to support young people who are struggling with mental health and well-being.

Sometimes young people come to Instagram dealing with hard things in their lives. I believe Instagram can help many of them in those moments. This is something that our research has suggested as well. One of the internal studies that has been the subject of much discussion showed that teen boys and girls who reported struggling with loneliness, anxiety, sadness, and eating disorders were more likely to say that Instagram made those difficult times *better* rather than worse.

We care deeply about the teens on Instagram, which is in part why we research complex issues like bullying and social comparison and make changes. We have a long track record of using research and close collaboration with our Safety Advisory Board, Youth Advisors, and additional experts and organizations to inform changes to our apps and provide resources for the people who use them.

We don't allow people to post graphic suicide and self-harm content, content that depicts methods or materials involved in suicide and self-harm (even if it's not graphic), or fictional content that promotes or encourages suicide or self-harm. In the third quarter of 2021, we removed 96 percent of this content before it was reported to us.

Since 2019, we've taken steps to protect more vulnerable members of our community from being exposed to suicide and self-harm related content that is permissible under our policies, for example, if someone posts about their recovery journey. We remove known suicide- and self-harm-related posts from places where people discover new content, including our Explore page, and we will not recommend accounts we have identified as featuring suicide or self-injury content.

---

[2] Aditya Purohit et al., Designing for Digital Detox: Making Social Media Less Addictive with Digital Nudges, ASSOC. FOR COMPUTING MACHINERY (Apr. 2020), https://dl.acm.org/doi/10.1145/3334480.3382810.

[3] Christoph Schneider et al., Digital Nudging: Guiding Online User Choices through Interface Design. Communications of the ACM (July 2018), https://cacm.org/magazines/2018/7/229029-digital-nudging/fulltext.

We also remove certain hashtags and accounts from appearing in search. When someone starts typing a known hashtag or account related to suicide and self-harm into search, we restrict these results. We also add sensitivity screens to blur more content that isn't graphic but could have a negative impact on someone searching.

We have a resource center[4] developed with help from mental health partners, and, when a post is identified as being about suicide (either because a friend reported it or our technology detected it), a person at Meta reviews the post. If it's about suicide, we provide resources to the poster such as a one-click link to the Crisis Text Line. Additionally, whomever reported the post also receives resources and information about how to help the person in distress.

Similarly, we don't allow content that promotes or encourages eating disorders on our platforms. We use technology and reports from our community to find and remove this content as quickly as we can, and we're always working to improve. We follow expert advice from academics and mental health organizations, like the National Eating Disorder Association ("NEDA"), to strike the difficult balance between allowing people to share their mental health experiences while protecting them from potentially harmful content.

We've made a number of changes to support those struggling with eating disorders. When someone searches for or posts content related to eating disorders or body image issues, they'll see a pop-up with tips and an easy way to connect to organizations like NEDA in the US.

We also introduced a dedicated reporting option for eating disorder content. People have always been able to report content related to eating disorders, but, until recently, this was combined with the option to report suicide and self-harm-related content, because they are part of one policy— but now people will see a separate dedicated option for eating disorder content.

We also worked with the JED Foundation to create expert- and research-backed educational resources for teens on how to navigate experiences like negative social comparison.[5]

Lastly, we don't allow people to bully or harass other people on Instagram and have rules in place that prohibit this type of content. We've also built tools that help prevent bullying from happening in the first place and empower people to manage their accounts so they never have to see it.

We launched Restrict in 2019, which allows people to protect themselves from bullying without the fear of retaliation.[6] We also created comment warnings when people try to post potentially offensive comments. So far, we've found that, about 50 percent of the time, people edited or deleted their comments based on these warnings.

---

[4] Suicide Prevention, https://www.facebook.com/safety/wellbeing/suicideprevention.
[5] More information on this work is available here: https://pressuretobeperfect.jedfoundation.org/.
[6] Introducing the "Restrict" Feature to Protect Against Bullying, Instagram Blog (Oct. 2, 2019), https://about.instagram.com/blog/announcements/stand-up-against-bullying-with-restrict.

We recently announced a new tool called 'Limits' that lets people automatically hide comments and direct message requests from people who don't follow them, or who only recently followed them. We developed this feature because we heard that creators and public figures sometimes experience sudden spikes of comments and message requests from people they don't know. In many cases, this is an outpouring of support, but sometimes it can also mean an influx of unwanted comments or messages. Now, if you're going through that—or think you may be about to—you can turn on Limits and protect yourself.

We also recently launched Hidden Words, which automatically filters message requests containing offensive words, phrases, and emojis into a separate inbox so people never have to see them. Because messages are private conversations, we don't proactively look for hate speech or bullying the same way we do elsewhere on Instagram, so Hidden Words allows people to control what they see and receive in messages and protect themselves from abuse. In addition, all accounts on Instagram have the option to switch off messages from people they don't follow. This means people never have to receive a message from anyone they don't know.

These are just a few examples of the tools we developed to protect people from bullying and harassment. We have numerous other tools including comment controls, blocking, and managing who can comment on your posts and who can tag and mention you.

Giving Teens Tools to Control their Experience

We want to give people on our platform—especially teenagers—tools to help them manage their experiences in the ways that they want and need, including the time they spend. We have built time management tools including Daily Limit, which lets people know when they've reached the total amount of time they want to spend on Instagram each day; 'You're All Caught Up,' which notifies people when they've caught up with new content on their feed; and controls to mute notifications.

This week, we launched 'Take A Break' to go even further and empower people to make informed decisions about how they're spending their time on Instagram. We'll show reminders suggesting that people close Instagram if they've been scrolling for a certain amount of time, and we'll show them expert-backed tips to help them reflect and reset. We want to make sure young people are aware of this feature, so we'll show them notifications suggesting they turn the reminders on.

Also this week, we began testing a new activity center, a central place for people to see and manage their information on Instagram. For the first time, people will be able to bulk delete content they've posted like photos and videos as well as their previous likes and comments. While available to everyone, this tool will help young people more fully understand what information they've shared on Instagram and what is visible to others and give them an easy way to manage their digital footprint.

Prioritizing and Expanding Parental Controls

We want parents to have the information to help their teens have a safe and positive experience on Instagram. That's why in March we're launching Instagram's first set of controls for parents and guardians, allowing them to see what their teens are up to on Instagram and manage things like the time they spend in our app. These new features, which parents and teens can opt into, will give parents tools to meaningfully shape their teen's experience.

In the US, we've also collaborated with The Child Mind Institute and ConnectSafely to publish a new Parents Guide that includes the latest safety tools and privacy settings as well as a list of tips and conversation starters to help parents navigate discussions with their teens about their online presence.[7]

III.   Using Research to Improve Instagram

A lot of focus in recent weeks has been about internal research. As our Head of Research Pratiti Raychoudhury has written, the public reporting about our internal research was mischaracterized, so I want to take a moment to address it. Among other things, the research in question actually demonstrated that many teens said that using Instagram helped them when they were struggling with the kinds of hard moments that teenagers have always faced.

In addition to putting specific findings in context, it is also critical to make the nature of this research clear. This research, some of which relied on input from only 40 teens, was designed to inform internal conversations about teens' most negative perceptions of Instagram. It did not measure causal relationships between Instagram and real-world issues.

Our goal with all of the research that we do is to improve the services that we offer. That means our insights often shed light on problems so that we can evaluate possible solutions and work to improve. We believe this work is critical to delivering a better Instagram.

Moving forward, we will continue to collaborate and engage in data-sharing with researchers on issues related to young people. We have been working with external academics and research partners in this space for many years, and we plan to do even more early next year. This is something that we have done in our program with independent academics around the US 2020 elections. We will take the methodology from the US 2020 program and apply it to well-being research over the coming year. This will involve collaborative co-design of studies and peer-reviewed publication of findings.

In addition, we are continuing our investment in external research to better understand how to keep young people safe and to ensure their well-being is protected in the metaverse. For example, we committed to providing $5 million over three years to the Digital Wellness Lab at Boston Children's Hospital for independent research on these important topics.

---

[7] Instagram Teen Safety for Parents, https://about.instagram.com/community/parents#guide.

IV. Supporting Industry Regulation to Protect Young People

The reality is that keeping young people safe online is not just about one company. We've been calling for updated regulations for nearly three years. From where I sit, there is no area more important than youth safety.

Specifically, we believe there should be an industry body that will determine best practices when it comes to at least three questions: how to verify age, how to design age-appropriate experiences, and how to build parental controls. This body should receive input from civil society, parents, and regulators to create standards that are high and protections that are universal. And I believe that companies like ours should have to adhere to these standards to earn some of our Section 230 protections.

In addition, the body could take steps to require each member to publish regular reports on the progress they are making against each standard and to develop a free and accessible information hub for parents and educators.

This proposal is a work in progress, but we hope that it will contribute to the ongoing discussion about how appropriate regulation can help us address these critical issues. In the meantime, we will continue to push forward on safety and well-being for young people online.

V. Conclusion

We want young people to enjoy using Instagram while making sure we don't compromise on their privacy and safety. As we work toward that goal, we'll continue listening to them, their parents, lawmakers, and experts to build an Instagram that works for everyone and is trusted by parents.