

FACEBOOK, INC.
COMMUNITY STANDARDS ENFORCEMENT REPORT
Q2 2021
August 18, 2021

FACEBOOK
August 18, 2021
11:00 a.m. ET

Operator: This is Conference # 289430

Operator: Hello and welcome to today's Community Standards Enforcement Report Global Press Call. There will be prepared remarks and a Q&A to follow. To ask a question after the prepared remarks conclude, please press the 1 followed by the 4 on your telephone. As a reminder, this conference is being recorded, Wednesday, August 18, 2021.

Now, I'd like to turn the call over to Sabrina, who will kick this off. Please go ahead.

Sabrina Nadia Siddiqui: Hi, everyone. Thank you for joining us. You should have received embargoed materials ahead of this call. We are on the record and we are embargoed until 10 a.m. Today, you will hear opening remarks from Vice President of Integrity, Guy Rosen; Vice President of Content Policy, Monika Bickert; Director of Product Management, Anna Stepanov; and CTO, Mike Schroepfer. We will then open up the call for questions.

In addition to releasing our Community Standards Enforcement Report this quarter, we are also sharing additional integrity and transparency items. This includes our first Widely Viewed Content Report and updates on our COVID-19 response. With that, I'll kick it over to Guy.

Guy Rosen: Hey, everyone. Thanks for joining us today. Today, we're publishing the Community Standards Enforcement Report for the second quarter of 2021. This report provides metrics on how we enforce our policies from April through June.

Today, we're also releasing the first in a series of reports that we call the Widely-Viewed Content Report, which will give an overview of the most widely-viewed content in Facebook news feed starting with the top 20 most viewed domains, links, pages, and posts in the U.S. This report will be public and available in our Transparency Center, and we're going to update it each

FACEBOOK, INC.
COMMUNITY STANDARDS ENFORCEMENT REPORT
Q2 2021
August 18, 2021

quarter going forward together with the Community Standards Enforcement Report. Anna is joining us today and she will tell us more about it shortly.

On the Community Standards Enforcement Report, this is now our 10th report, and I'd like to maybe take a step back and look at our progress over time. We now include 13 policy areas on Facebook and 11 on Instagram. We have expanded coverage of our metrics about content action, our proactive rates, appeals, restores, and most importantly prevalence.

Prevalence is our primary metric, and it is the percentage of views on our platform that are violating content, and it matters because it captures not what we took down but what we missed and what was ultimately seen by people.

For example, on hate speech prevalence on Facebook has decreased now for three quarters in a row since we first began reporting it last year. In Q2, it was 0.05 percent, or five views for every 10,000 views on Facebook.

Our content moderation systems used to rely mostly on user reports. We've changed that over the past years. We've invested heavily in artificial intelligence to help us enforce at scale. The proactive rate metrics helps track that as it measures how much of the content that we took action on was detected by our systems as opposed to needing to be reported by users.

For example, on hate speech in 2017 when we first started reporting, that proactive rate was 24 percent, meaning that over three-quarters of our removals were facilitated by someone needing to report the content to us. Now, it's over 97 percent.

Even with this progress, there will always be examples of things we missed, and with the scale of our enforcement there will be examples of things that we take down by mistake. There is no perfect here, and we believe that the progress on this work should be evaluated through the metrics in this enforcement report. These are the same metrics that we operate on internally, measuring progress on keeping prevalence of violating content down while minimizing mistakes on the content we removed. We hope to continue expanding the transparency that we provide into this work.

Now on COVID-19, as the pandemic continues to develop we're focused on three things. First, understanding vaccine acceptance on Facebook. Secondly, promoting reliable information and encouraging vaccinations. And third, taking action against misinformation. Let's talk about those for a moment.

On the first, understanding vaccine acceptance, we are focused on outcomes which we believe are the right way to evaluate the end result. For example, for people in the U.S. on Facebook, the vaccine hesitancy has declined by 50 percent. We've similarly seen vaccine acceptance rising globally.

FACEBOOK, INC.
COMMUNITY STANDARDS ENFORCEMENT REPORT
Q2 2021
August 18, 2021

Our COVID-19 trends and impact survey, which we conduct in partnership with Carnegie Mellon and University of Maryland has since the beginning of the year shown vaccine acceptance rising, for example, by 35 percent in France, 25 percent in Indonesia, 20 percent in Nigeria. This is all movement in the right direction and we'll continue watching this data closely and working with health authorities around the world to help promote vaccine acceptance.

Now on the second point, we want to work to make sure people have access to reliable information. And thanks to products like our COVID-19 Information Center, we've connected over two billion people globally with resources about COVID-19. In the U.S. alone, we've helped over 4 million people get vaccinated through our Vaccine Finder, which connects them with appointment information, directions and contact information.

In the second quarter of this year, the vast majority of COVID related content that was widely viewed on Facebook was from authoritative sources including, for example, from the WHO and UNICEF. And it's most important for people to see their friends and family supporting vaccination. Since the beginning of the pandemic more than 18 million people globally have used Facebook profile frames supporting vaccines and Instagram COVID-19 vaccine stickers have been used by more than 7.6 million people globally.

More than 50 percent of people in the U.S. on Facebook have already seen someone use the COVID-19 vaccine profile frame which we developed in collaboration with the U.S. Department of Health and Human Services and the CDC.

Now on the third point and very importantly, we're also continuing to enforce against COVID-19 and vaccine misinformation on our platforms. That enforcement is built on that content moderation capabilities we built over the years and on our partnerships with health experts and with our network of over 80 independent fact checking organizations globally that cover over 60 languages.

Since the beginning of the pandemic, we've displayed warnings on more than 190 million pieces of COVID related content on Facebook that our fact checking partners rated. And we've removed more than 20 million pieces of content from Facebook and Instagram globally for violating our policies on COVID-19 related misinformation.

We've also taken action against people who repeatedly post content that breaks our rules. Since the beginning of the pandemic, we have removed over 3,000 accounts, pages and groups for repeatedly violating our rules against spreading COVID-19 and vaccine misinformation.

FACEBOOK, INC.
COMMUNITY STANDARDS ENFORCEMENT REPORT
Q2 2021
August 18, 2021

Our work is never over and we're continuing to look at content to assess emerging trends as this pandemic evolves and to engage with health experts and work with authorities around the world on this issue.

One final update about the report regarding child safety. In previous reports, we shared metrics for one policy area, child nudity and sexual exploitation of children. In today's report and going forward, we're sharing metrics through two categories under this broader topic of child endangerment.

The first is nudity and physical abuse, and the second is sexual exploitation. To be clear, this isn't a change in our policies or in our enforcement, only in how we're reporting this. And we're producing this more granular reporting in order to provide a more detailed and transparent views into our efforts in the space for the benefit of child safety experts and researchers and, of course, the general public.

Over the past few months, we've also introduced a number of new product updates to strengthen children's safety and privacy across our platforms. For example, on Instagram, young people who sign up are now defaulted into a private account and we're notifying those who are already on Instagram and have public accounts to encourage them to switch to private.

On both Facebook and Instagram, we've introduced restrictions to prevent contact between teens and potentially suspicious adults. And no adult on Instagram can DM people under 18 who don't already follow them. Across all our services, we continue to report on any apparent instance of child exploitation to the National Center For Missing and Exploited Children.

With that, I'll turn it over to Monika.

Monika Bickert: Thanks, Guy, and hi, everybody. Thanks for joining today. As Guy mentioned, this is our 10th report. So, that's our 10th report in a process that I believe has held us to a higher standard. It has kept us accountable to the goals that we've set for ourselves.

As many of you know, I lead the team that writes Facebook's policies. I want to take a moment to add just a few thoughts to what Guy shared about our ongoing efforts around COVID-19. Especially because this is an area where my team has had to react quickly to developing guidance from health authorities over time.

Since the start of the pandemic, we've been working to identify and remove false COVID-19 and vaccine-related claims that public health authorities have debunked, and identified as contributing to a risk of imminent physical harm.

FACEBOOK, INC.
COMMUNITY STANDARDS ENFORCEMENT REPORT
Q2 2021
August 18, 2021

We now have more than 65 specific claims that we remove from our platforms around COVID-19 and vaccines because they are false and may contribute to the risk of imminent physical harm during the pandemic. We're continuing to add to this list as new trends emerge.

For instance, in the past month, we added to our list claims that COVID-19 vaccines cause Alzheimer's, that the vaccines cause magnetism, and that being around vaccinated people could cause secondary side effects to others.

Finally, I'll just add that while we're encouraged by the progress that we're making we also know that progress and the fight against misinformation will lead to new challenges. For example, we've seen some efforts to evade our misinformation enforcement through the use of coded language to spread vaccine misinfo. While the need for coded language is an indication that our enforcement is working, it also highlights why it's important for us to continue to evolve our policies and our enforcement to address these attempts to evade our enforcement.

Now moving onto the oversight board, it's now almost 10 months since the oversight board has been up and running. And we've now gotten into the rhythm of responding to their recommendations. The board has selected 19 cases since they began their operations last year and they've issued about 60 recommendations.

So, they have their binding decision about whatever the specific piece of content is that they're addressing, and then they also have the recommendations that they issue alongside their binding decisions. We've responded to each of those recommendations within 30 days and we've committed to action on a majority of them.

I really think that these recommendations are going to have a lasting impact on our content moderation across the platform. For example, in response to the board's recommendations we have launched and are continuing to test more specific messaging to send to people when we remove their content.

We made progress of the specificity of our hate speech notifications by using an additional classifier that's able to predict what kind of hate speech is in the content, meaning is it violence, dehumanization, mocking hate crimes, visual comparison, claims of inferiority, contempt, cursing, calls for exclusion or slurs.

As a result, people who are using Facebook in English now receive more specific information if they violate our hate speech policy and have content removed. We intend to roll out more specific notifications for hate speech violations to other languages in the future.

FACEBOOK, INC.
COMMUNITY STANDARDS ENFORCEMENT REPORT
Q2 2021
August 18, 2021

In response to another board recommendation, we're now running tests to access the impact of telling people whether automation was involved in our enforcement decision. In June, we updated our Dangerous Organization's policy language to provide more public detail. And in the new language, we explained our three tiers of content enforcement for different types of dangerous org designations and we defined key terms like praise, support and representation.

This update follows earlier updates where we've shared details about our strike system so that people know what actions our systems will take if they violate our policies.

All of these updates influenced by the Oversight Board's recommendation are examples of how the board is helping us make meaningful and lasting improvements to our work.

I also really hope that people will take some time to read through some of the decisions that the board has published on their website. In those decisions, they dive into the complexities of speech issues and they shine a light on some of the difficult trade-offs that are the very nature of content policy and content moderation.

With that, I'll hand it over to Anna.

Anna Stepanov: Thanks, Monika. So as Guy mentioned, today, we're releasing our first widely viewed content report, a quarterly update that provides an overview of the content that reaches the most people on Facebook. This report includes public content, meaning content that can be viewed by everyone on Facebook. This initial report contains four sections that each cover the top 20 most viewed pieces of content across posts, pages, links and domains in news feeds in the past quarter.

While it does not cover ads, it does include content recommended by Facebook within news feed recommendation unites like suggested for you. So we already shared some information about public content through our insights tool CrowdTangle, which includes engagement data, meaning the number of interactions. So things like reactions, comments shared, the content received.

But CrowdTangle does have its limitations. It contains no information on (reacher) views, focuses specifically on engagement and represents a limited sample of all pages, groups and accounts. As the result, we wanted to use this report to take a step further and provide people with an even more comprehensive view of the content that is seen on Facebook.

FACEBOOK, INC.
COMMUNITY STANDARDS ENFORCEMENT REPORT
Q2 2021
August 18, 2021

This new report together with CrowdTangle will provide even more transparency and context about the experiences people have on Facebook in a way that's easily accessible by both stakeholders and consumers alike.

So some key takeaways from this first report include the content that gets the most engagement isn't necessarily what most people (see). Likes and comments don't equate to reach, as our ranking models consider thousands of signals well beyond engagement alone.

For example, they can incorporate survey results, such as when we ask people directly whether they feel a post is worth their time. As a result, posts that are engaged with in the form of likes and comments may not reach as many people as other posts that people tell us they find valuable.

Moreover, news feed is a deeply personalized experience. Said another way, while the content in the report was seen by the most people, it isn't in fact the content that is most seen by people in their news feed.

Throughout the report, we provide view-based data to give it – to give a sense of scale and again content. So for instance, posts represented in the WVCR make up less than a tenth of a percent of all post content in news feeds.

Building on this, the majority – so about 57 percent of posts that people see are from their family and friends. This is in line with changes we made in the past such as our meaningful social interaction change, or the MSI change, that prioritizes this type of content.

Zooming out a bit on the report, what we see is that there is quite a bit of diversity across the content that gets the most distribution across things like fun memes, mainstream media news domains, and authoritative health content. For example, the top post was a first three you see crossword, the top domain was YouTube, and the top page was UNICEF.

Moving forward, we will release these reports on a quarterly basis alongside the CSCR. This effort is a work in progress and a key part of our deep commitment to transparency. We're very interested in your feedback on this report as we continue to iterate on both format and scope.

Ultimately, this report is about holding ourselves accountable, and we think it will foster the right discussions around the experiences people have on Facebook and how we can continue to improve those experiences.

And with that, I'll turn things over to Schroep.

Mark Schroepfer: Thanks, Anna. It takes the work of tens of thousands of people to enforce community standards across our platforms, and their work is supported by A.I.

FACEBOOK, INC.
COMMUNITY STANDARDS ENFORCEMENT REPORT
Q2 2021
August 18, 2021

systems that can scale their impact to billions of users. Those systems have made extraordinary progress in recent years. Today, I want to highlight how we see that progress evolving.

A.I. is an incredibly fast-moving field, and many of the most important parts of our A.I. systems today are based on techniques like self-supervision that seemed like a far-out future just years ago.

We have moved the ideas from research to production at unprecedented speeds and they're helping drive the ongoing improvement in enforcement that today's numbers show. One example of this is reinforcement learning that we introduced last year, which we've named RIO.

The classic approach to training A.I. uses a fixed dataset to train a model that's then deployed to make decisions about new pieces of content. RIO does something new. It guides the model to use online metrics as reward signals to optimize across all aspects of development, data, feature, architecture, and parameters. It constantly evaluates how well it's doing its job, and it learns and adapts to make our platform safer over time.

This kind of end-to-end learning is incredibly valuable for enforcing community standards, because the nature of the problem is always evolving alongside current events. When new problems emerge, our systems need to be able to quickly adapt. Reinforcement learning is a powerful approach to help A.I. meet new challenges when there's a shortage of good training data.

We expect our RIO system to continue driving down the prevalence of hate speech and other unwanted content long into the future, which is very encouraging for such a new technology.

That's far from the only cutting-edge A.I. approach that drove improvements this half and we'll continue improving. In particular, the trend of replacing single-purpose bespoke systems with more generalized ones continues at a surprising pace. Our multimodal A.I. models which operate in multiple languages, multiple modalities like text, images and video, and across multiple policy areas continue to deliver impressive improvements this half. These kinds of A.I. models were almost unimaginable just a few years ago.

They only exist today thanks to fundamental breakthroughs in A.I. research and those breakthroughs have quickly moved from our labs into an operational system. We see no signs of this research-to-production pipeline slowing down.

Our researchers have made huge progress over the last year in areas like zero-shot and few-shot learning which enables A.I. systems to recognize violating content even if they've never seen it before or only seen a very few

FACEBOOK, INC.
COMMUNITY STANDARDS ENFORCEMENT REPORT
Q2 2021
August 18, 2021

examples of it during training. This ability to quickly begin enforcing against content types that we don't have a lot of label training for will be a major step forward for us. It will help make our systems even more agile and responsive to emerging challenges.

Zero-shot and few-shot learning are one of the many cutting-edge A.I. domains where we've been making major research investments and we expect to see results in the coming year. So we're very encouraged by the results shown in the numbers today but even more hopeful for what's coming.

There's a lot more exciting research making its way into our systems that I look forward to sharing with you on future calls. But for now, I'll hand it back over to our host for questions.

Operator: Thank you. We will now open the lines for questions. Please press the one followed by the four on your telephone. Again, as a reminder, it is the one followed by the four on your telephone.

Our first question comes from the line of Queenie Wong from CNET. Please proceed with your question.

Queenie Wong: Hi. So when I look at the Community Standards Report, it mentioned that you took action on 7 million pieces of content that contained terrorism in Q2. And given, sort of, the events around Afghanistan, I was wondering if any of that content was posted by the Taliban or is that not information you currently ...

Guy Rosen: Hey...

Queenie Wong: ... have?

Guy Rosen: Hey, this is Guy. Hi, Queenie. Thanks for the question. We don't have that information here with us today. I think that's something we should follow up with you on. We're certainly paying a lot of attention to the events that are happening over the past several days, the past week, and our thoughts are with everyone that is affected.

Similar to other crises, we have a dedicated team including Afghan nationals, and native Dari and Pashto speakers working 24/7 right now to monitor the situation in real-time and address emerging challenges.

Monika can maybe speak to how we're thinking about classification of content within this current series of events. But in terms of how we're approaching this overall, the overall crisis, that's some of the background on that.

Monika Bickert: Yes. And can I just say that the Taliban as an organization do violate our dangerous organizations policy and they have. I mean I've been in this job for

FACEBOOK, INC.
COMMUNITY STANDARDS ENFORCEMENT REPORT
Q2 2021
August 18, 2021

eight or nine years and certain when I joined the team, they were already designated. So we've been removing any accounts that represent the Taliban as well as any content that praises or supports them, that's something that we've been doing for many years and we're continuing to do that.

Like Guy said, we have a team of people that are actively monitoring the situation as it manifested in speech on our platforms and making sure that we're doing our best to remove that content.

Operator: Thank you. Our next question comes from the line of Ina Fried from Axios. Please proceed with your question.

Ina Fried: Hi, yes thanks. So you talked about a couple different stats with regards to COVID-19 misinfo but – and I may have missed it, it might be in the report but can you talk about the prevalence rate and obviously things like the frames that people can put to be pro-vaccine are very visible but sometimes the language and ways that people show anti-vaxx preferences, obviously, it isn't just putting a banner around it. And I'm curious what you've learned and how you've been – how able to grow the ability to detect sort of information that might lead to be more vaccine hesitant?

Guy Rosen: Hi, this is guy. Thanks for the question. So I appreciate the pointer to prevalence because we do think as we've repeatedly used in this report that prevalence is generally the way we think is right to measure these kind of efforts.

It is also a difficult metric to get right in across the areas. And as folks may remember, was an area like hate speech, people used to say that Facebook is awash with hate speech but last year when we published our prevalence number, which did take awhile, it's been establish that the use of hate speech are with this report for sharing five views out of every 10,000 on Facebook. It did take us a long time.

And when it comes to COVID though, things are evolving even more quickly. So it does make prevalence even more difficult to define and measure. Language is evolving and the kind of content that people are worried about is evolving.

Now as you've noted, there's a lot of authoritative and quality information on the platform and we know the vast majority of the most viewed content about COVID is from authoritative sources and care experts around profile and profile frames and stickers and things of that nature are also very important in this.

We do believe in the long term. We need to – prevalence is the right way to articulate and measure the space. We don't have that number here today. No

FACEBOOK, INC.
COMMUNITY STANDARDS ENFORCEMENT REPORT
Q2 2021
August 18, 2021

other company shares prevalence about this kind of misinformation, which kind of speaks to how complicated this is.

And the thing though that we're focused on really because it's nuance to precisely categorize content, measure the impact of any single piece of content. As you've said, some are more prominent than others, clear than others. The right way to measure is to focus on the outcome. And that's at the end of the day are people more or less accepting of the vaccine. And that's where I go back to the trends and impact survey that we run and that we conduct in partnership with Carnegie Mellon and with the University of Maryland. It's the largest survey of its kind, and those numbers are showing that globally as well as for people in the U.S. on Facebook vaccine hesitancy is declining. People are becoming more accepting of the vaccine. And we think that outcome-focused data helps ensure that we have insight in the totality of what's happening both on Facebook and elsewhere.

Operator: Thank you. Our next question comes from the line of Tim De Chant from Ars Technica. Please proceed with your question.

Tim De Chant: Yes. Hi. You note earlier that this was for public content. Does that include posts made without public-facing content shared between friends and family or is this only public content that is posted with the visibility set to public?

Anna Stepanov: Yes. This is Anna here. So this is content that has its visibility set to public and accessibility set to public.

Operator: Thank you. Our next question comes from the line of Issie Lapowsky from Protocol. Please proceed with your question.

Issie Lapowsky: Hi. Hi. Thank you for taking my question. I think my question is also for Anna. You mentioned that in comparing the most viewed content and posts and pages to CrowdTangle that CrowdTangle shows us engagement that is limited. I wonder if you could expand on sort of how you value a view, right? Are views just content that I am scrolling past in my news feed or pages I'm scrolling past in my news feed? And if that is the case, then why do you value views as heavily as engagement where you have people sort of actively interacting with posts, which might indicate a stronger attachment to the content?

Anna Stepanov: Yes, this is Anna again. So just to clarify the report – the WVCR report that is anchored on view data is not so much about the value of the view or whether we view – whether we value engagement or how we value engagement, but it is weighed to most accurately in a sense to represent the actual experiences people have, so it is what people are seeing in their – as they engaged with and use news feed, and the app.

FACEBOOK, INC.
COMMUNITY STANDARDS ENFORCEMENT REPORT
Q2 2021
August 18, 2021

Operator: Thank you. Our next question comes from the line of Shirin Ghaffary with Recode. Please proceed with your question.

Shirin Ghaffary: Hi, thanks for taking my question. My understanding is that the majority of what people share on Facebook is through their private pages amongst friends and family, so why not include some aggregated version of that in this list of the most viewed content?

Anna Stepanov: This is Anna again. So this first version of the WVCR we chose to focus on public content and, once again, we'll continue to build this out over subsequent reports across languages, surfaces, and other types of content.

Guy Rosen: Hey, this is Guy. Just to add there, obviously sharing information about private sharing that happens between peoples friends is a highly sensitive thing we have to conduct this in order – in a way that respects people privacy. We're certainly trying to understand the right way to share information and the aggregate and through other mechanisms in ensure that we can share a broad view of things.

But I want to make sure it's clear that there is always – there is attention here between ability to share content and information that is public and making sure that we are respecting what's happening privately between people.

Operator: Thank you. Our next question comes from the line of David Gilbert with Vice News. Please proceed with your question.

David Gilbert: Hi, everyone. Thanks for doing this. I just wanted to ask – because you spoke about transparency and helping researchers and journalists with these reports. Has there been any update on the situation with the researchers at New York University, Laura Edelson and Damon McCoy in relation to their accounts being reinstated after they were banned for the rollout of the Ad Observer tool?

Guy Rosen: Hey, thanks for the question, this is, Guy. This is a complex topic and look we've been working with these researchers for almost a year to try to find ways to enable their research. We've – in addition to giving them access to research tools like our Ads Library, (API) and CrowdTangle. We also offered them access to an initiative we have called the Facebook Open Research and Transparency Program, which is exactly designed to create a mechanism through which we care share – we can share data in a way that's privacy protective.

And we offered them a data set which details political ad targeting information. They however declined to join this program. Now when they started scraping again without informing us, the responsible thing we had to do was enforce our terms until we can ensure that the risks are mitigated.

FACEBOOK, INC.
COMMUNITY STANDARDS ENFORCEMENT REPORT
Q2 2021
August 18, 2021

This, to be clear, wasn't the path any of us preferred here. But the methodology does create privacy risks, which as a company we have to prioritize.

Now browser extensions like the one that this team uses allow access to whatever data someone who installs it can access including their friends data and private messages. And as much as it's very well intentioned and we want to support this and we know others have looked into the code and vouched for it. Even if the code avoids collecting private data today, if there is a bug or a compromise or maybe an unrelated change in our code that changes, for example, the rendering of a page and their code parses it the wrong way, it could lead to an unintentional collection of personal data. And we will be held accountable for the consequences of that.

So we have to protect people's privacy and security and we have to work with them to find a way to do this in a privacy protective way.

Operator: Thank you. As a reminder, to ask a question please press the "1" followed by the "4" on your telephones.

Our next question comes from the line of Emily Birnbaum with Politico. Please proceed with your question.

Emily Birnbaum: Hi, thanks for having the call today. I just – I had a quick question. I know that some of the pushback to the CrowdTangle data was that it promoted the idea that right-wing content did – does well on Facebook and Facebook uplifts right-wing content. So, is this most used content report, would you say is it an attempt to push back on the narratives that Facebook is dominated by right-wing news?

Anna Stepanov: So, this is – this is Anna, again here. So, the – this report is really part of, again, our broader commitment around transparency. I will say that the, again, the content that does get the most engagement isn't necessarily the content that people most see.

So, again, we're taking a number of different signals from other things like survey results. And so, this is – this isn't a commentary on any particular type of content. It is, once again, just our next step in creating greater transparency about what people see.

Guy Rosen: Hey, this is Guy, just to add. There's a few gaps in the data that's being used to date. And the narrative that has emerged is quite simply wrong. CrowdTangle is focused on interactions. CrowdTangle only has a limited set of certain pages, groups, and accounts.

FACEBOOK, INC.
COMMUNITY STANDARDS ENFORCEMENT REPORT
Q2 2021
August 18, 2021

We are creating a report that provides a broad view and represents as Anna I said earlier, an accurate representation of what people's experiences actually are on our platform. We think it's really important for this data to be out there and it's a report – that's why we're doing this report, and that's why this report is something we're going to continue to do as well.

Operator: Thank you. Our next question comes from the line of Kurt Wagner with Bloomberg. Please proceed with your question.

Kurt Wagner: Hey, thanks, good morning. I think this is for Monika. I just was – wanted to follow up on the issues that are happening in Afghanistan. I know you said the Taliban has long been considered a dangerous organization. Are you revisiting that distinction at all? And are you waiting on any public signal as to whether you should – you should recognize the Taliban as a – the formal government there in Afghanistan? I'm just wondering if basically, things could change from the way you've been doing them?

Monika Bickert: Thanks, Kurt. What I can tell you is there – they are prescribed under U.S. law. They're sanctioned under U.S. Treasury – under the U.S. Treasury list. So, they're not allowed to have a presence on our services and our – we don't make determinations about who is able to make government in a place. We're not a – we're not a body that recognizes governments. We look to the authority of the international community in setting our policies. And part of that is following U.S. laws.

So, they will not be allowed while they are prescribed under U.S. law and even if they were not prescribed under U.S. law, we would have to do a policy analysis on whether or not they nevertheless violate our dangerous organizations policy.

Operator: Thank you. Our next question comes from the line of Julie Jammot with AFP. Please proceed with your question.

Julie Jammot: Hi. I would like to come back to the figures about the vaccine, saying that the vaccine hesitancy has declined by 50 percent, and that people have seen their contact with colleague vaccine profile frames. I'm just wondering how do you get that figure of declining by 50 percent? And how do you know that this really has impact that people are actually getting vaccinated from seeing the frames or vaccine hesitancy declining? Thank you.

Guy Rosen: Hey, this is Guy. Thanks for the question. So we've since early in the pandemic have been running a large-scale survey called the COVID Insights and Trends Survey that asks people across the globe a number of questions, including – we started this actually asking around symptoms related to COVID and it's been a useful tool in order to actually predict based on

FACEBOOK, INC.
COMMUNITY STANDARDS ENFORCEMENT REPORT
Q2 2021
August 18, 2021

people's reported symptoms, the actual prevalence of COVID in – on a geographic bases for example.

As the vaccines rolled out we also added questions around people's acceptance of vaccines. Now this is a survey that we run in partnership with Carnegie Mellon and the University of Maryland. They actually do the data crunching. So what happens is we – because we have such a large user base, we're able to sample and for the random sample of users to take the surveys and the answers actually all go to those universities. We don't see the answers. And they crunch the data and understand people's actual attitudes here.

This is, we believe, the largest public health survey anyone's actually done in history. We have something like over 17 million responses across 200 countries and territories and the data can be broken down by a number of factors. This is all – the data is public, can be found – we have links from our website that goes to report – the reports and the data crunched by those universities.

And so, it – the data that we quote is based on what we're seeing there, where we actually see the people's responses to our questions around vaccine acceptance improving. So more people are saying they are accepting of the vaccine.

If people – people who don't like the vaccine love – would love to fill in surveys and tell us how much they don't like it. And so we see this as a very encouraging fact that people are filling in these surveys and responding that they are accepting of the vaccine and that is improving. We see this tracking well. We see this tracking well in the U.S. and we see this tracking well globally as well.

Operator: Thank you. Our next question comes from the line of Marty Swant with Forbes. Please proceed with your question.

Marty Swant: Hi, thanks for doing this. Yes, so what was Facebook's (own) efforts, I'm wondering if the company has plans to support the Senate legislation to create the Health Misinformation Act that was introduced earlier this summer that would hold social media accountable – social media companies more accountable for health related misinformation? Any updates on that?

Monika Bickert: Hi. Yes, we've long been supporters of common industry standards and Section 230 reform generally. And we do think that clarification on the really difficult and frankly urgent questions about health related misinformation would be helpful. So we look forward to working with Congress and the industry as we consider what options for legislative reform could look like.

FACEBOOK, INC.
COMMUNITY STANDARDS ENFORCEMENT REPORT
Q2 2021
August 18, 2021

Operator: Thank you. Our next question from Pavel Kasik with Seznam. Please proceed with your question.

Pavel Kasik: Hello, can you hear me?

Guy Rosen: Yes, we can.

Pavel Kasik: I'm sorry, I had trouble with my microphone. Can you hear me now?

Guy Rosen: We hear you.

Monika Bickert: Yes.

Pavel Kasik: In the – in the report, you talk about hate speech content removal has increased over 15 fold. Could you explain more of what it means? Like is it a percentage, is it relative to – because 15 folds could also mean that there's so much more hate speech that you're doing so much more removal or it could mean that there's just more content so it's staying the same. So could you maybe break down this number a bit more?

Guy Rosen: Hey, this is Guy. Thanks for the question. I think this was an important point to clarify. So the 15 fold refers to the amount of hate speech that we take down today versus what the amount we were taking down in the first quarter that we reported. So the pace of our removals has indeed increased very significantly based on improvements we've made, particularly through artificial intelligence, as we've said. Today, the vast majority of what we removed is detected by our systems before people need to report it to us.

Now you bring up a really important point with these metrics, content action numbers could go up because we're getting better and they could go up because there is more – more violating content to find. That's why we believe prevalence is a really important metric.

The prevalence metric doesn't talk about what caught but it talks about what we've missed and ultimately people have viewed. And so that's why we, in this report, across different areas have been reporting prevalence. And what we've seen, for example, for hate speech is prevalence of hate speech has been decreasing for three quarters in a row now, even as the removals have gone up.

Since we – this is – the prevalence has been declining since we've started reporting it. And our first report with this was last year and it is now at 0.05 percent. So five views for every 10,000 and that's down from the previous numbers we had.

FACEBOOK, INC.
COMMUNITY STANDARDS ENFORCEMENT REPORT
Q2 2021
August 18, 2021

This is why it's really important to look at prevalence numbers to ultimately hold us accountable and this is what we're – we use internally to hold ourselves accountable. It's not about how much we remove but how much we missed and how much people ultimately end up seeing.

Operator: Thank you. Your last question comes from the line of Shannon Bond with National Public Radio. Please proceed with your question.

Shannon Bond: Thanks for taking the question. So my question is, I guess, for Ana about the widely viewed content report. So in the section on the widely viewed post, I see the number three post is something that it was not – it's not available. It says this Facebook post is no longer available. It may have been removed or the privacy settings may have been changed.

Are you going to, in the future, be able to give a little more information when you have something like that because it seems to me there's quite a difference if it's something where somebody changed the privacy settings of a post versus it was maybe taken down for a standard violations. So are you going to be able to kind of share more information in cases where things have gotten a lot of views but are no longer available to be included in the report?

Anna Stepanov: Yes, that's – that's a great question, and you're right in this case the removal is on account of the change in privacy setting. Again, this is our first report and I think feedback and ideas like this are actually very helpful as we sort of form our next reports and refine how we talk about the details in those reports. So I would say that is something we would consider, yes.

Guy Rosen: One more thing to add is as we think about content and how the content here is shown in the widely viewed content report and we talked earlier in a previous question about private versus public or – private, we mean friends-only content, which is a lot of what is actually on Facebook.

One thing to clarify is when there is a public post that is, for example, re-shared by someone to their friends, and that is with a friends-only setting, the views on that re-share will count towards that public post that we show in this report.

Similarly, if there is a link and it is posted for example linked to a news site, and I post it to my friends and that gets views, that we will also accrue to the – to the links to the domain section in the widely viewed content report.

So in this way for the – for the public content that is circulating even in those friends-only settings, we are including that in this report, we think that's important, and so I just wanted to clarify that point on how we are approaching the data in this widely viewed content report.

FACEBOOK, INC.
COMMUNITY STANDARDS ENFORCEMENT REPORT
Q2 2021
August 18, 2021

Just as Anna said, this is our first one, we're definitely going to iterate, much as we have on the Community Standards Enforcement Report, which by the way similarly based on understanding people's experiences and its (use) on the platform, it's a really big focus of ours and making sure that we're reporting in these transparency reports what are the experiences people are having on our platform and being held accountable to those.

So we definitely look forward to feedback so that we can improve these reports and make them better and more useful for everyone over time.

Sabrina Nadia Siddiqui: Hey, guys, thank you everyone for joining. If you have any additional questions or you need any follow-up, you can reach out to me Sabrina or reach out to our press line. Thank you.

Operator: Thank you. This concludes the Community Standard Enforcement Report Global Press Call. Thank you for joining. You may now disconnect your lines.

END