

Facebook, Inc.

Moderator: Emily Cain
August 11, 2020
10:00 a.m. PT

OPERATOR: This is Conference #9857609

Operator: Hello and welcome today's press call there will be prepared remarks and a Q&A to follow. To ask a question and after the prepared remarks conclude, please press "star" "1". And now I'd like to turn the call over Emily Cain who will kick us off.

Emily Cain: Thank you and hi everyone. Thank you for joining us to announce the sixth edition of community standard enforcement report. You should have received a copy of the report and the accompanying newsroom post via email ahead of this call.

Today you will hear opening remarks from Vice President of Integrity Guy Rosen and Vice President of Content Policy Monika Bickert. We will then open up the call for questions. There is no embargo for this call and it is on the record. With that, I'll go ahead and kick it over to Guy. Guy?

Guy Rosen: Thank you. Good morning, everyone. So today we're releasing our sixth community standard enforcement report covering how we're doing at finding and removing harmful content from April through June. For the past couple of years we have done these once a half and now we are moving to quarterly just like our earnings report.

I'd like to walk you all through some of the highlights. To start, COVID-19 continues to have a impact on our content enforcement. To protect the health and safety of our workforce, in March we sent all our content reviewers home.

We then operated on two parallel tasks; one enabling our reviewers to work from home as much as possible and, two, relying more heavily on our technology to review content. We expedited some work which was already underway to reorganize our content enforcement system.

We're now using artificial intelligence to create a ranking of the most critical content for our teams to review, regardless of whether it was reported by users or detected by proactive system. This enables our team to spend their time on the cases where we need their expertise the most. And it means there will be a shift towards more content being initially actioned by our automated systems.

Today's report gives us a fuller picture of the impact of the COVID-19 pandemic on content review. It shows that we do have the tools and processes to respond to emerging issues then to continue finding and removing content that violates our policies. But it also reinforces that content enforcement is not an either/or approach where we use either human reviewers or AIs. Sophisticated systems need both people and technology.

Let's say for example suicide and self injury and child nudity and sexual exploitation content. These are two areas where we rely on people to both review content as well as to constantly help improve the technology that proactively finds and removes these and more.

And with fewer reviewers we took action on fewer pieces of content on both Facebook and Instagram for suicide and self injury and on Instagram for child nudity and sexual exploitation.

Despite this decrease, we prioritized and we took action on the content that is the most harmful such as live videos. But reviewing this content continues to be challenging. It can't be done from home due to its very graphic nature. We want to ensure it's reviewed in a more controlled environment and that's why we starting bringing a small number of reviewers, where it's safe, back into the office.

The number of appeals is also lower in this report because we couldn't always offer them. We still gave people an option to tell us that they disagreed with

our decision on a piece of content and our teams looked at these signals in aggregate to find potential issues and restore content where appropriate. You can see that in the report counted as content restored without an appeal.

Now lastly because we prioritized removing harmful content over measurements of certain efforts we were unable to calculate the prevalence in a few areas, in violent and graphic content and adult nudity and sexual activity.

We anticipate that the metrics from Q3 will be available in our next report. Now despite the challenges due to the pandemic, in some areas we have made improvements to our technology that's enabled us to take action on more content.

For example, in hate speech, the amount of content we took action on increased from 9.6 million in Q1 to 22.5 million in Q2. Of that, 95 percent was detected proactively before anyone reported it to us, up from 89 percent in Q1 and 24 percent when we first started publishing these reports.

In Q1 we made improvements to our proactive detection technology and expanded automation in Spanish, Arabic, and Indonesian. In Q2 we followed up by expanding with more automation in English, in Spanish, and Burmese.

On Instagram, we took action on 3.3 million pieces of content in Q2 for hate speech, 84 percent of which was detected proactively. And that's up from 45 percent in Q1. Additionally the amount of terrorism content we took action on increased from 6.3 million in Q1 to 8.7 million in Q2.

We've made progress in combating hate on our apps. But this work is never done. And as Mark said on our earnings call the other week, we don't benefit from hate. We don't want it on our platforms and fighting it requires a combination of enforcement and policies. Monika will speak to this more.

In the past months we've prioritized work around harmful content related to COVID-19 that could put people at risk. If some misinformation poses imminent harm, we remove it. So from April through June, we removed over

7 million pieces of harmful COVID-19 misinformation from Facebook and Instagram.

This is things such as posts that push fake preventative measures or exaggerated cures that the CDC and other health experts tell us are dangerous. For other misinformation, we work with independent fact checkers to display warning labels.

From April through June, we put these labels on about 98 million pieces of COVID-19 related misinformation on Facebook. We're also very focused on the upcoming U.S. election. And we spent the past four years building up systems and teams to fight misinformation, to stop election interference and to push out accurate information.

We've already broadened our policies to ban more content that would mislead people about voting or try to intimidate them so they don't vote. And from March to July, we removed more than 110,000 pieces of content in the U.S. for violating those policies.

We're also building up our election operation center to continue our work with state election authorities that we can quickly respond to and remove false claims about polling conditions in the 72 hours leading into election date.

Earlier today we also introduced new ratings and updated labels for our fact checking program. We've heard from fact checkers that more clarity and distinction between ratings is important and better reflects what they're seeing on the platform.

As Mark also announced in June, we're going to use our election operation center, which I just mentioned, and we're going to be rolling out our voting information center soon. It will be a one stop shop for people to find accurate authoritative information about the election.

Now, finally back to this report. We know people want to be confident, but the numbers in this report are accurate, so we're planning to undergo an independent third party audit of our content moderation systems to validate the numbers that we publish.

Now, this week we are issuing a request for proposal to external auditors and we hope to conduct this audit in 2021. Now, I'd like to turn it over to Monika with more detail on the people and the policies that back our enforcement numbers.

Monika Bickert: Thanks, Guy, and hello, everybody. Guy made reference to the progress that we're making on identifying and quickly removing hate from our services. And these numbers are important because they help people track our progress over time and they serve as a metric of accountability.

But it's equally important to understand the people, and the policies, and the product work that goes on behind those numbers. First the people, there have been questions recently about the diversity and the makeup of the teams that write our content policies and weigh-in on the content decisions.

I lead the team that writes our policies. It's made up of about 200 people based in 11 offices around the world. And that team includes experts in everything from child safety to cybersecurity to hate organizations, and this team also has a diverse range of backgrounds, experiences, and political views.

Beyond that, we work regularly with hundreds of organizations and experts around the world to help us understand diverse perspectives in crafting our policies. This is a regular part of what we do and has been for years.

We're always looking to broaden the input that we receive. And so in that spirit we are setting up a diversity action council that will include a cross-section of employees from our company to provide input based on their lived experiences on a variety of topics that touch policy and product development.

Turning to our policies, we continue to refine our policy lines as speech and society evolve. In our fight against all forms of hate, for example, we've consulted with many of our external group partners over the past year to better understand (the implicit) speech that has historically been used to disparage, intimidate, or exclude people based on protected characteristics like race or religion.

So that consultation, which included more than 60 outside experts including social psychologists, historians, folklorists, and groups that represent a diverse range of religious, ethnic, and racial communities, has helped us write rules for the removal of more implicit hate speech, such as content depicting blackface or stereotypes about Jewish people controlling the world.

This type of content has always gone against the spirit of our hate speech policies but it can be really difficult to take concepts, especially those that are commonly expressed in imagery and define them in a way that allows our content reviewers based around the world to consistently and fairly identify violations.

Those challenges will remain, but we are committed to improving our policies wherever we can and we'll continue to work with experts around the world to get better at doing (this).

We also continue to explore new ways to combat dangerous organizations, like terror and hate groups. So first, we continue to ban these types of groups from having any presence on our services and we remove content that represents, or praises, or supports them.

As part of implementing that policy, we routinely evaluate groups and individuals taking into both -- taking into account both their online and their offline behavior. To date, we've identified and banned hundreds of groups around the world.

And last year, we took an enforcement strategy that we had successfully used in a different area, which is combating coordinated inauthentic behavior, and we started using that to focus on dangerous organizations who attempt to evade detection or return to our services after we've banned them.

Since October of 2019, we've used this tactic 14 times to remove 23 different banned organizations. The first strategic network disruption we conducted last October was against three U.S. based white supremacist groups, (the Right Stuff, Patriot Fronts, and Identity Dixie).

(In here sense), we've disrupted networks tied to hate groups including the Northwest Front, The Proud Boys, The KKK, Combat 18, and (The Bing).

The last thing I want to mention is the product work they're tapping to ensure that we introduce technology, tools, features, and products that are fair and equitable. We've set up two new teams.

At Instagram, the equity team and at Facebook, the inclusive product council; both of which our tasked is making sure we since check key moments in product development with diverse perspective both internal and external to the company.

So this could mean something simple like making sure that the stickers in our library reflect the diversity of our community. It could also mean addressing more complicated questions such as how we'll ensure fairness in our algorithms which are used to help us understand things like whether companies are likely to violate our policy.

We're committed to identifying concerns about equity, equality, and civil rights as early as possible so that we can make a difference in investments where needed and engineer stronger more impactful products.

As we continue to make progress, not just on policy enforcement but also on the people policies and the products that back that enforcement, we will communicate openly and often about what we're doing and why this is where we are today with the CSCR in this press call. And with that I will turn it back over to the moderator for questions.

Operator: We will now open the line for questions. To ask a question, press "star" followed by the number "1". Your first question comes from the line of Steven Overly from Politico. Please go ahead.

Steven Overly: Thank you. Hi. I just wanted to ask on the hate speech numbers where each of the (inaudible) is 22.5 million pieces of content compared to 9.6 last quarter, is that rise due completely to changes to technology or where there other reasons for that number growing so significantly?

Guy Rose: Hey, this is Guy. Thanks for the question. The change is largely driven by the increase in proactive technology – proactive detection as driven through the technology that we’ve been working on.

Operator: Your next question comes from the line of Laurence Dodds from the Daily Telegraph. Please go ahead.

Laurence Dodds: Hi (inaudible). Can you hear me all right? Thanks very much for the question answer – (inaudible) information. I have a question about your policy against images of black folk. In the Netherlands, it seems there’s been some sort of (travesty) about how this applies to (inaudible), which has been already very controversial (inaudible).

I wanted to just check with you – pardon me, check whether the policy would also apply English Traditional (inaudible) who have a similar tradition of (inaudible) (using black face) called that I believe. I like – (for us to check that it doesn’t apply to them and with) – (that’s a good picture of them).

And then also just if you – I wonder if you could explain a little about your rationale for including those things within the category of (things you banned) from my time sitting on conference standard forums before. I know that sometimes it would be considered whether they should be a cultural exception – the different policies.

I would imagine that might be considered – in this case – (inaudible). I wonder if you could extend more of your rationale and what the process was in deciding against them. Thank you.

Monica Bickert: Absolutely. So as you know, often when we talk about hate speech, we’re focused on more explicit attacks on people based on a protected characteristic, but over the years, we’ve -- and especially through our work with outside experts, we’ve identified often it’s images or it’s sort of mean content or generalization that is more implicitly attacking people.

And that over time has led to us actually having a list in our community standards where we say here are some stereotypes that we’ve removed as hate speech.

And so the ones that we're adding now today were arrived at very much in the same process, and you've seen it in the Content Standards Forum. For those who don't know, that's a meeting we have every two weeks where we look at potential revisions to our policies and we look at feedback that we've received from groups around the world, and we did that here.

In fact, we worked for the better part of the year in consultation with partners around the world to understand not only how can these stereotypes be used in harmful ways, but also what is the kind of speech, as you point out, what are the kinds of speech that might actually unintentionally be swept in to this policy if -- depending on how we draw the line.

And so, for instance, when you think about our policy that is prohibiting black face that we are launching now. That is designed to stop people from using black face to target or mock black people. But you can imagine somebody sharing an image that shows black face where they are doing so to say look at this politician who wore black face, everybody should know this and they're trying to raise awareness.

And you can also imagine that there could be other circumstances where somebody might happen to be sharing images of black face in it but they're not doing it for hateful reasons.

And so, those are exactly the sorts of nuances, including the example you raised in the Netherlands and the U.K. that we are looking at to understand now and have taken into account as we define our (in-communication) standard. We are announcing the policy today, it is not live yet in terms of enforcement.

Operator: Your next question comes from the line of Sarah Frier from Bloomberg. Please go ahead.

Sarah Frier: Hi, thank you. My question relates to just a follow-up from the civil rights audit. In that audit, they noted that content reported by users as voter interference is only evaluated and monitored for aggregate trends, so it doesn't tend to get looked at by a human moderator.

I'm wondering whether you have addressed that concern and are you changing the way you review voter suppression ahead of the election?

Monica Bickert: Sure. Yes, so we do actually have clear policies against voter suppression and we have people that review content that we either detect or get reports about violating those policies.

So we do have people look at those and make those determinations. And our voter suppression policy, for instance, ban misrepresentations about how to vote or voting logistics or methods of requirements or whether a vote will be counted or a threat to violence related to voting.

And in fact because of the civil rights audit and some of our more sustained engagement with the broader civil rights community, we've actually gone further recently and expanded our policies to also remove content that says that participating in the vote will or could result in law enforcement consequences, so you'll be arrested if you show up or deported if you show up and vote today, and also posts that indicate an intent or encourage others to bring weapons to polling places.

So, the audit was helpful in helping us the other ways that we can expand those policies but we do use people when we enforce them.

Operator: Your next question comes from the line of Queenie Wong from CNET. Please go ahead.

Queenie Wong: Thank you so much for taking the time to answer my question. So, Guy, you mentioned that Facebook is using AI to create a ranking of the most critical content for the company to review. And then you mentioned in an example that you're prioritizing content that is most harmful such as live videos.

Can you explain in more detail like how this ranking system works because like if somebody says for example like I want to kill myself and they put that in a live video versus just putting it in a regular post that seems like it would have the same amount of potential harm or damage.

So, is the system like ranking content based on like live video versus photos versus stories or something like how does the whole system work because it seems like trying to determine what is more harmful than something else is pretty -- it could be difficult.

Guy Rose: Hey, thanks for the question. So, the way this AI works is it tries to understand and evaluate how severe the content of a certain post might be. Live videos was one example I gave which is something we've always wanted to make sure we get to very fast.

But you're absolutely correct that a post -- the text post that has some sort of imminent -- someone is imminently implying they might be taking their life is something that we're -- that would absolutely also be ranked very highly because it's really important for us get to that.

If you step back and sort of think about the way technology helps other content review historically it's helped in proactive detections, so finding things people don't report or finding them sooner. It's helped in automation which is enabling one decision by a reviewer to be extended to many more pieces of content.

And it helps with prioritization which is increasingly where we've been putting more focus so that the time a reviewer spends is on the content that's the most harmful or content that's being seen by the most people and is really important for us to get to fast.

Operator: Your next question comes from the line of (Izzy Wachowski) from (Protocol). Please go ahead.

(Izzy Wachowski): Hi, guys. And I'm glad I'm following Queenie because my question is related. So, Guy, you have been saying that there was a dip in the amount of self injury and suicide content and child sexual abuse material content that was taken down in part because of the reduction in human moderators who were available due to COVID.

I wondered if you could talk a little bit about why that type of content requires more human review because to the point you just made you would expect that most egregious content to be the most proactively action.

And you guys obviously took a lot of proactive action on hate speech. So I'm wondering why child sexual abuse material or self injury would prove more difficult without human beings.

Guy Rose: Hey, thanks for the question again. The -- so, the reason this content is challenging is because its graphic content that honestly at home is very hard for people to moderate with people around them and so forth.

And so, we want to be very careful and the people and the environment that people have in order to look at that content. In some of these areas it is a combination of people and technology working together with people making some of the more nuance decisions and technology helping to amplify their impact significantly.

And so what was really important to us is to make sure that the reduced workforce that is able to work that's more difficult type of content is able to focus on the most severe and the most harmful categories within that pool of content that's available to review and that we're taking action on that content.

Operator: Your next question comes from the line of David Uberti from Wall Street Journal.

David Uberti: Hey all thanks for taking the time. I had a similar question just related to this. You've spoken a bit in the past about sort of like the push and pull, give and take between AI and human reviewers on counterterrorism context and how some of these groups basically try to (just trick the) systems or avoid the systems which requires more human reviewers. I was curious if you could try to flush that out a little bit for some of these other categories?

If there's particular types of areas or particular types of content or user behavior the AI is more adept at picking up over the last few months in particular? Or areas where the human reviewers are still key?

Guy Rosen: I think overall this pandemic and the situation has really (reinforced us) that it is always people and technology working together and you can't let the technology just sort of go entirely on its own.

We always need people who also look and measure and help to tune automation to ensure that we're always up to speed and always up to date with sort of how content is evolving in different – be it more adversarial things in the nuances that you mentioned and an area where perhaps someone is trying to constantly change how they're expecting things in content.

Or just keeping up to date with the latest kind of trends and culture and the way people are expressing themselves. (Now this is – there's) always a feedback loops and so the way it work is when reviewers make a decision that decision is then used by automation to help train and learn and detect more similar or identical pieces of content.

And so we need to keep that sort of cycle going which is why also where we focus our workforce is working on the most harmful kinds of reports but also making sure that we're – that they're focused on measuring and improving the quality of the different kinds of systems so that they can continue to be up to date over time.

Operator: Your next question comes from the line of Mike Isaac from New York Times.

Mike Isaac: Hi there. Sorry my question is probably boring. But I'm just trying to figure out – so like you're talking about how the COVID stuff is impacting how content moderation sort of works and wanting to sort of keep it safe and not have people review things from their homes.

So I guess like I was just wondering what the partnerships with the third party contractors look like at this point? Are you still – like the, I don't know, the (picture) of the world do you still have those going on but they – they aren't handling all the content? Are those sort of closed too? I guess I'm just trying to figure out what that – what that even looks like right now.

Guy Rosen: Hey, Mike, thanks for the question.

Mike Isaac: Sure.

Guy Rosen: We continue to work with our – with our partners to help us scale the workforce especially in many places around the world, many of which are in different phases of the pandemic.

And we're working with them literally everyday to understand where we can bring some people back into some offices with the right safety precautions. What's the right environment and the support and well being support that we can provide their people whether it's at home or within offices.

And we're also working with our own full-time employee workforce where we have people who are trained on content moderation, they typically help to sort of train and manage that larger workforce and they have picked up more of that work in the past months to ensure that we are on top of the most severe and the most harmful content.

Operator: Your next question comes from the line of Alex Heath from The Information. Please go ahead.

Alex Health: Hey, Monika and Guy. I was wondering, could you elaborate on Facebook's approach it's taking to QAnon -- the QAnon movement especially going into the election.

There was a story the other thing about how there's a lot of QAnon content on Facebook and (inaudible) that movement with regards to content moderation and any kind of more systematic action you might be taking or not taking against QAnon. Thanks.

Monika Bickert: Sure. I'll take that. We have removed -- unfortunately QAnon is not for us. We've removed pages and groups and other QAnon content where it's violated our policies and also where they're behavior has violated our policy.

So as you know there are times were a content might cross the line and then there's times where the accounts may be inauthentic or the way that they are behaving on our Facebook violates our policies. So we've taken action

against them in April. We did it again in July and we'll keep looking at other ways for making sure that we are addressing that content appropriately.

Operator: Your next question comes from the line of Mark Sullivan from Fast Company. Please go ahead.

Mark Sullivan: Thank you. I'm still focusing on this large increase of the number of hate posts that you took action on from Q1 to Q2. You give a couple of different reasons for that. One of them is that you say that you expanded to take action on Spanish, Arabic, and Indonesian or posts in those languages. And you say that it's also because of the English detection technology improvements.

But I wonder if you could give an idea of removing those other companies from -- or those other languages from the equation. Is there a way that you can quantify the increase in the number of actions taken just on English language post from Q1 to Q2?

Guy Rosen: Hey. Thanks for the question. We don't have that breakdown here so I'd have to -- we'd have to follow up with you. But generally it's the increase in proactive and automated systems that is driving this increase across content action for hate speech.

Operator: Your next question comes from the line of Raphael Balenieri from Les Echos.

Raphael Balenieri: Yes. Just picking up on what you said, except the improvement in languages, isn't it also because globally you saw more hateful and problematic content on Facebook than COVID because basically everyone was at home spending time online?

And here in France, I mean many French NGOs actually said there was just much more hate speech on the platforms during that time. So what's your comment on this? And then I had a question on the audit.

Who could be the organization that will audit your work on this? And perhaps a final question on China, given the current tension -- have you seen more disinformation from China on Facebook these days? Thanks.

Guy Rose: Hey there. Thanks for those questions. I'll try to remember three different ones. So again, I mean the question – the question on the amount of hate speech. The metric we really want to get to is the prevalence of hate speech on Facebook.

That's something we are working on, and we plan to have it in our next enforcement report; and it will describe how many – in the same way we have prevalence metric for a number of other categories.

It will describe how – how many views on Facebook are content that violates our hate speech policies. The – just to reiterate – the increase that we saw and that is reflected in this report.

If you – if you do the math in terms of the content that is removed, then it is indeed driven by the increase in automation as a result of our work on – across a number of different languages; and generally to make sure that we're automating and that we're expanding and amplifying decisions by humans or detecting and sometimes even taken action automatically on different kinds of posts.

Your second question on the audit – the -- so we're issuing a request for a proposal to external auditors. There's a number of large reputable firms – we're working with folks like them, and this is the kind of audit that's not been done before on this kind of system so it's not something their accustomed to doing as oppose to typical financial or related reporting, which the world certainly knows how to execute.

And so we're working with them to explain and to walk through, and obviously we want something that is independent and a – an independent assessment from a reputable firm that we'll be able to come in, ask questions, and pretty confident in the accuracy of the numbers that we are reporting out in this regular report. And for your third question, I'll turn over to Monika.

Monika Bickert: Yes. In terms of disinformation campaigns, that's something that we continue to see globally, and we do – as Guy mentioned earlier – we do have a team now that's really focused on identifying this kind of behavior.

That team that's focused on understanding disinformation campaigns and what we call coordinated inauthentic behavior, they work with other companies and they work with researchers and security firms to understand this behavior, get a head of it, and take it down.

And we publish a report every month now where we detail exactly what we've found, so our recent one that we put out did not have anything from China. It did have takedowns from Brazil, the U.S., Yemen, Ukraine.

We have taken down operations from China before. I can remember one that was around this time last year, but if you go back and look at those reports you'll see that this is very much something we focus on globally.

Operator: Your next question comes from the line of Glenn Chapman from AFP. Please go ahead.

Glenn Chapman: Hello. Good morning, evening, or afternoon everybody. Just hopefully some quick points of clarification, can you say how much – specifically how much reduction in content review has been done by humans due to the COVID -- we're referring to the increase (inaudible) automated systems.

Can you say its cut in half, cut into a third? (Can we have) just a clarification on how much human reduction review has happened? And then, I apologize if this was previously announced but Monika, when would those changes to, like, you'll be deported if you show up to vote, when were those implemented? Is that, sort of, previously announced and -- or are those new, kind of, those kinds of announcements?

And then, Monika, also, if you can, give some more context on how you decide what kind of meme or implicit content it is. Like for example, the new ban on content, does it include Holocaust deniers? Does that go into that section of Jews controlling the world? So some more sense of how you figure that one out. Thank you.

Guy Rosen: Hey, I'll jump in on that first question and then Monika can take the second and third parts. So on human review workforce, we don't have exact numbers to change. Largely, most reviewers are back online working from home.

But the thing that we work through as we think about the system, the content enforcement, is not just how many people but how effective they are in a slightly different environment.

What kinds of content they are able to review given the different environment they're in and how we need to balance what kind of work they do to ensure that we're taking care of their well-being and their families. And so, it's a bit of a more nuanced question. But largely, mostly of our reviewers are online. But we continue to work through just the details of how our enforcement works.

Monika Bickert: And I can take the other two. On -- maybe I'll take them in reverse order actually. About the changes with the hate speech policy around memes around, for instance, Jewish people controlling the world or global institutions, that's the boundary basically of that policy change.

Although, I will say, even getting to that policy change we engaged with dozens of organizations and experts. And those relationships, we'll definitely keep. We'll keep those conversations going because we know that the trends around hate speech will continue to evolve and our policies will continue to evolve with it.

In terms of the changes to -- the expansion of the voter suppression policy, those are recent and we did announce them. I want to say it was late June, June 26th or so, in Mark's announcement. So just to be clear, our voter suppression policies where we remove misinformation about how or when you can vote, those have been in place since before the 2018 midterms.

But in response to our dialogue with our civil rights auditors and the broader civil rights community, we did extend the policy to cover more. And that now means that we will remove content that tells people they're going to get arrested if they go to the polls, or says I'm going to bring a weapon to the polls, or other people should bring weapons to the polls.

Operator: Your last question comes from the line of Musadiq Bidar from CBS. Please go ahead.

Musadiq Bidar: Hey, all. Thanks for the time. We're now about 80 days away from the election. I'm wondering if you have specific plans to ramp-up resources within that department for content moderation as far as fake news is related to the elections, vote-by-mail, and other issues. Do you have plans to ramp-up (each of these for) that department?

Guy Rosen: Hey, there. Yes, our -- I mean, our response to the pandemic is -- has largely been a result of all these investments that we've made over the years. And as we've prepared for the elections, we are absolutely committed and focused on working towards the U.S. 2020 election to make sure that we are protecting the integrity of those elections.

There's many different kinds of risks and we've worked over the past years to combat coordinated inauthentic behavior, misinformation, voter suppression, and we continue to work on those as well as evaluation in this fast changing environment we are always sort of red teaming and working with partners to understand what are the next risks?

What are the different kinds of things that may go wrong? And very much focused on ensuring that we're preparing for different scenarios that may -- that may occur. And building systems and building the teams to be able to handle those.

Emily Cain: Thank you so much everyone for joining us on this call. As a reminder there was no embargo and the call is on the record. Hope everyone enjoys the rest of their day. Thanks very much.

Operator: This concludes the Facebook press call. Thank you for joining. You may now disconnect your line.

END