

FACEBOOK, INC.

Moderator: Andrea Saul
May 23, 2019
1:00 p.m. ET

Operator: This is Conference # 4285738

Operator: Hello and welcome to today's Facebook Press Call.

There will be prepared remarks and a Q&A to follow. To ask a question after the prepared remarks conclude, please press "star," "1."

Now, I'd like to turn the call over to Andrea Saul, who will kick this off.

Andrea Saul: Hi, everyone. Thank you for joining with us today. I'm Andrea Saul on Facebook's Communications team and today we're going to be talking about the third edition of our Community Standards Enforcement Report that we released today. In addition to hearing from Mark Zuckerberg, you're going to hear from Guy Rosen, Monica Bickert, and Justin Osofsky and then we'll take your questions.

This call is on the record.

With that, I'll kick it to Mark.

Mark Zuckerberg: Hey, everyone. Thanks for joining today.

For the last few years we've been focused on rebuilding our system for removing harmful content so we can find and take down more of this content proactively instead of reactively once someone reports it to us. And as part of this, we've focused on measuring the amount of harmful content and tracking

how effectively we're removing it. Now as we've gotten more confident in our measurements, we've started publishing them so people can hold us accountable for our progress and today we're releasing the third of these transparency and enforcement reports.

Starting next year, we're going to start publishing these quarterly because I think the health of the discourse is just as important as any financial reporting we do so we should do it just as frequently. Understanding the prevalence of harmful content will help companies and governments design better systems for dealing with it and I believe that every major internet service should do this. In our next report we're also going to start adding Instagram as well.

When you're working on important social issues like speech and safety, the process needs to be as transparent as possible. So in a few minutes Guy Rosen is going to take you through the figures we're publishing, Justin Osofky's going to talk about our content review teams and Monica Bickert is going to give you an update on our policies then we're going to take some questions.

But first, I want to give some context around our work here. Every day people around the world are using our services to express themselves and giving everyone a voice is a principle that we believe in and hold very deeply. We have a responsibility to protect people's freedom of expression in everything we do but at the same time we also have a responsibility to keep people safe on Facebook and to prevent harm from playing out across our services.

Last year I wrote a long note about a blueprint for content governance and enforcement that would describe some ideas on how we think we should approach this. And these are hard tradeoffs to make and people can come to very different conclusions about the best way forward. But as we build our systems and update our policies, we have to remember that different countries and cultures have very different traditions around free speech.

It's impossible to build a single system that's going to work perfectly everywhere and that will catch everything. But there's a lot more that we can

do to improve accuracy and to stay ahead of new threats that might emerge. On the governance side, we update our policies continually because the issues are always evolving. Our team consulted experts across society to make sure we're taking as broad of a view as possible when setting these policies. The system isn't perfect. Whenever you're drawing a line around what content is acceptable, there will always be some people who think you've got it wrong. But overall I'm proud of the diligence and thoughtfulness that goes into this process.

At the end of the day, our teams are always going to have to make some judgment calls about what we leave up and what we take down. If the rules for the internet were being written from scratch today, I don't think most people would want private companies to be making so many of these decisions about speech by themselves. Ideally, we believe that this should be part of a democratic process which is why I've been calling for new rules and regulations for the Internet. This is something I've been talking to policymakers a lot recently, including with President Macron in Paris a few weeks ago.

I imagine that these rules will have to be different by country, depending on each country's traditions around speech. In the U.S., for example, where we have very strong constitutional protections on free speech, it might make more sense for there to be regulation requiring an industry body to set standards governing the distribution of harmful content and measuring companies against those standards. In other countries, it might make more sense for regulation to set clear baselines and require companies to build effective systems to proactively combat harmful content.

Now, in the meantime before any new rules are put in place, we're also moving forward on our own with plans to establish an independent Oversight Board so people in the community can appeal our content decisions. We know that our systems can feel opaque and people should have a way to hold us accountable and make sure that we're enforcing our standards fairly. So this independent Oversight Board will look at some of our hardest hit cases and the decisions it makes will be binding. That means that even if our teams or even if I disagree with its decisions, we won't be able to overturn them.

So we've spent the first half of this year working with experts on speech and safety, running workshops around the world and asking for public input on how this could work. And we're going to be publishing a report with all the feedback we've gotten so far at the end of June. Our transparency report which we'll be discussing today focuses on the enforcement side, finding and removing content that violates our policies. And it's important that we aren't grading our own homework here, so we've asked an independent data transparency group of experts to check that our methodology is sound.

Today's report shows that we're making progress in areas like hate speech and graphic violence. We're increasingly catching it before people report it to us. We proactively identified 65 percent of the hate speech we removed before anyone reported it to us, up from about 24 percent a year ago. That gives a sense of the progress but also how much there is still left to do.

We also proactively identified about 83 percent of posts and comments trying to sell drugs. And we're taking down more fake accounts than ever, 2.19 billion in Q1, up from 1.2 billion in Q4 last year. This is due to an increase in automated attacks by bad actors who try to create large amounts of fake accounts at once. And most of these accounts were blocked within minutes of their creation, before they could do any harm. So they were never considered active in our systems and we don't count them in any of our overall community metrics.

On bullying and harassment, we have a lot of work still to do. In this cycle, we proactively found less bullying and harassing content, down from 21 percent last quarter to 14 percent in the first quarter of this year. For child nudity and the sexual exploitation of children, we took action on less content than previous quarters. And although we detect more than 96 percent of adult nudity on our services, we didn't proactively identify as much as we did in previous quarters because in this cycle we asked our teams to focus on more severe content and prioritize proactively identifying that instead.

We know that there's a lot of work ahead, not just on these specific challenges but on content issues more broadly. We're very focused on this and I'm

personally very focused on this. And we're going to continue working as openly as possible as we continue making progress.

And now I'm going to hand it off to Guy to talk about the report in more detail.

Guy Rosen: Thanks, Mark, and morning, everyone. My name is Guy Rosen and I lead the product and engineering team focused on safety and integrity here at Facebook. Today, I want to walk you through some of the highlights from this third edition of our community standards enforcement report.

Over the past two years, we have drastically changed our approach to all of this space. First, thanks to the growth in our business, we've been able to massively invest in all things safety and integrity. The development in AI in recent years means our systems can detect more of this content proactively before anyone reports it to us; sometimes before anyone even sees it.

But most importantly, we've adopted a methodical approach of understanding and quantifying the experiences that billions of people around the world actually have on our services -- using that to identify the biggest gaps and executing on a multiyear roadmap to improve. As we've instrumented our systems and developed metrics that we use internally, we now include those same metrics in this report so that people can hold us accountable for our progress.

In today's report, we've added some new datapoints that I want to flag right at the top. First, we're including data on regulated goods; specifically, firearm sales and drug sales. This is in addition to the eight policies that we also included in the previous report. Second, as Mark noted, we're sharing data on how often people appealed our decisions and how often we restored content that didn't actually violate our policies.

We think these numbers are important because they show where people think we made the wrong calls, where people disagree with our policies, and where we corrected our mistakes. These new numbers add to the three main measures we included in past reports. Rather than diving into all of the details of each, I'll quickly go over the metrics and highlight some standout numbers.

The first and perhaps the most important is prevalence; in other words, how much bad content did people actually see? You can think of this like an air quality test to determine the concentration of pollutants in the air. Just as an environmental regulator might periodically sample air quality to calculate what percent of the air we breathe is, let's say, nitrogen dioxide, we periodically sample content that's viewed on Facebook to calculate what percent violates our policies.

And we focus on how much content is seen; not just how much sheer content is out there that violates our rules. In that way, a post that's seen, let's say, a million times is a million times -- is a million times more likely to be sampled and that's a good thing. Today's report shows that for every 10,000 times people viewed content on Facebook, 11 to 14 views contained adult nudity or sexual content, 25 views contained violent or graphic content.

And as we already reported in our earnings, we estimate that fake accounts make up around 5 percent of our monthly active users. All three of these metrics are up a bit from the last quarter. We're focused on addressing this trend and making sure we improve our ability to catch more in the coming quarters.

For the first time in this report, we're also including prevalence numbers for terrorist propaganda and for child nudity and sexual exploitation. These are some of the most egregious violations so we're focused on making sure that as few people as possible ever encounter them. The prevalence for both areas is actually too low to measure using our standard method. But using a different method, we've been able to estimate that in the first quarter of 2019, for every 10,000 times people viewed content on Facebook, less than three of those views may have contained content that violated each policy. As opposed to the other areas, this is an upper bound; the actual prevalence may be even lower.

The second measurement is how much content we took action on. This includes removing content, applying a warning screen or disabling accounts. This metric reflects how often people violate our policies and how much of it

we identify. It also means that it can fluctuate a lot. For instance, if we find and remove a handful of viral memes that violate our hate speech policy, that might drive the number up one quarter and then the number can drop the next quarter if people no longer try to share those memes.

For child nudity and sexual exploitation of children, we take action on less content than previous quarters. This may be driven by a number of factors, including fluctuations in viral content, tightening of our guidelines for what content we include in the database of hashes that we check all images against. And additionally, a bug in one of our systems impacted our ability to add new hashes of videos that we had already removed. We fixed this bug, and we're working to remove content that may have been missed.

We also saw a significant increase in the number of fake accounts we took action on. We disabled 1.2 billion accounts in the fourth quarter of 2018; almost 2.2 billion in the first quarter of this year. And these numbers are driven largely by automated attacks by bad actors who try to create large volumes of accounts. Most of these accounts, blocked within minutes of their creation before they can do any harm, and they were removed so soon they were never considered active, and they don't show up in our monthly active user numbers nor in our prevalence.

Lastly, we measure the rate of proactive detection. This is a way to reflect how effective AI is in a particular policy area, as it shows how much of the content we took action on was detected by our systems before someone reported the content to us. In six of the policy areas we include in this report, we proactively detected more than 95 percent of the content we took action before someone reported it. And when it comes to hate speech, we now proactively detect 65 percent of the content we remove; that's up from about 24 percent a year ago.

We've made real progress, but we have more work to do when it comes to enforcing our policies and fighting abuse on Facebook. And we're focusing, not just on catching more, but, more importantly, on missing less, as measured by prevalence. And that's we're investing heavily in AI that can catch more violations before people see it. This technology has helped us make huge

strides in the past few years, but AI is also not a silver bullet. We need a combination of technology and people, whether it's the people who report content to us or the reviewers on our teams who review content.

And to share more on that, I'll now turn it over to Justin.

Justin Osofsky: Hi, everyone. My name is Justin Osofsky and I lead our global operations team, which includes around 15,000 people across the globe who review content and enforce our community standards. This is critical and challenging work, and so I want to share a bit about how we're supporting these teams, which is also an absolutely essential element to how we invest in keeping our platform safe. So today, I'm going to share a bit about how we're supporting these teams and the practices we're putting in place to make sure they can make informed and accurate decisions.

First, on support. Last week, we announced that we will be increasing pay for our content reviewers. This means at least \$22 per hour in the Bay Area, New York City and Washington D.C., \$20 per hour for those living in Seattle, and \$18 per hour in all other metro areas in the U.S., and we're working to develop similar international standards. We're also rolling out some new programs and tools to promote wellbeing and resiliency.

We'll require partner companies who run our content review sites to provide on-site counseling during all hours of operations; not just certain hours of each shift. And we're improving the tools that reviewers use to view and make decisions, such that they can customize how they view content. This is just the start.

We're also building more direct communication channels with partners and reviews, because we want to hear from them -- on what's working and what's not. Content review at this scale is no easy feat, but it's on us and our partners to put the wellbeing of our reviewers first as we continue to grow and improve our operation.

I also want to take a moment to share the progress these teams have made on one of our toughest policies to enforce -- hate speech. As Guy mentioned, AI has really changed the game when it comes to detecting violating content. But

what it still can't do well is understand context, and context is key when evaluating things like hate speech. A slur for instance is often an attack on someone based on race, national origin or sexual orientation -- but it can sometimes be a joke, one used self-referentially or could be employed to raise awareness about the bigotry someone has experienced.

We've invested a lot to understand linguistic and cultural nuance around this kind of speech, and we've also trained our AI to detect similar content to what our teams have already removed. The numbers we're releasing today make clear that this approach is working, but that there's also more we can improve. And I want to share some of the steps we're taking to make even more progress.

First, we ask our content review team for more details after making a decision, which can help strengthen our AI. When a reviewer looks at a piece of content to decide whether to leave it up or take it down, they are also asked to label which policy guided their decision. This information is then shared with the engineers on Guy's team and used to power machine learning algorithms that will thereafter help detect other similar content that violates the same policy.

Second, we're continuing our work to give reviewers more context when they review posts that may contain hate speech. Rather than just showing the image or comment that's been reported, we try to give a fuller picture. For instance a screenshot of hate speech might be accompanied by a caption that clarifies the user is condemning this kind of speech. But our review tools have not always presented the reviewer with this clarifying caption; we're working to change that so reviewers have what they need to make the right call.

And last but certainly not least, we're launching a pilot program where some reviewers will specialize in hate speech. Right now, most of our reviewers look at content across the spectrum. By focusing on hate speech enforcement, these reviewers will establish a deeper understanding of how it manifests and be able to make more accurate calls.

We think these steps, paired with the advances Guy's team is making in our products and tools will go a long way to helping us catch and remove more bad content that breaks our rules. We'll keep sharing progress on this work, and as we do we also need to explain what those rules are and how they're developed.

For more on that I'll turn it over to Monica.

Monica Bickert: Thanks, Justin, and hello, everyone. My name is Monica Bickert, and I lead the team that writes the rules for what is and isn't allowed on Facebook -- what we refer to as our community standards.

I'm going to talk about how we write these rules, and then I'll share some updates and changes that we've made in recent months. In shaping our policies we consult with experts around the world in public service, law enforcement and the nonprofit sector. We're focused on three core principles -- giving people a voice, keeping people safe and applying the rules fairly and consistently to the people who use Facebook. We know that the rules we write have broad implications for online speech; that's why we think it's important to be as transparent as we can as it relates to the writing, revising and enforcement of these rules.

So over the past year and a half, we've made some big moves in that direction. We've started publishing the minutes from the policy development meetings that we have every two weeks where we discussed and debate changes to our community standards. And earlier this year, we launched a searchable archive of the updates to our policies so that people can actually track our policy development process over time.

We continue to build out the committee standards enforcement report. As Guy noted, this third edition includes numbers around appeals and restored content. For instance, with hate speech, between January and March of this year, we received appeals for 1.1 million pieces of content that we had removed, and of those, we restored 155,000 pieces of content after reviewing them a second time.

This is really important. Not only does it tell us where and potentially why we're making enforcement mistakes, but it also helps my team understand where we might need to continue working with subject matter experts to make sure our policies make sense and where we need to provide more clarity on what is and isn't allowed on the site.

The appeal numbers that we shared are broken down into two categories -- content restored after appeal and content restored on our own, even when it wasn't appealed. So when it comes to the latter, there's a few reasons that might lead us to restore content even where it wasn't appealed. Think for example of a case where we removed multiple instances of the same post. Once we know that we've made a mistake, we will restore all the shares of that content and we don't require an appeal for every single one.

This happened recently when people were posting a famous photo of a scantily clad Burt Reynolds as a tribute after his passing. Initially, we took down posts that featured that photo that violated our nudity policy, but when someone appealed our decision, we realized that the photo did not in fact break our rules, thanks to some strategic arm placement. So we restored the original photo, as well as other instances and shares even if they hadn't been appealed.

As Mark noted, this system will never be perfect, and there will always be people who disagree with us, they'll think we've gone too far, or they'll think we've not gone far enough. But we will continue to develop thoughtful policies that balance voice and safety and we'll continue to invest in both proven and innovative ways of enforcing them. Finally, we'll continue to be open about our progress along the way.

And in that spirit, let's open it up for questions.

Operator: We will now open the line for questions -- we will now open the line for questions. Please limit yourself to one question per person. To ask a question, press "star" followed by the number "1."

Your first question comes from the line of Tony Romm from Washington Post. Please go ahead.

Tony Romm: Hey, everybody. Thanks so much for doing this and taking my question.

Mark, I wanted to ask you about the report that the transparency advisory group had put together that was released today with the transparency report. The report itself says that quote, "folks there did not speak directly with engineers," and then the reports found that they could not evaluate the extent that Facebook's daily operation deviate from the process that was described to us in oral briefing. So essentially it feels like the researchers here are saying that they can't actually determine whether Facebook workers are doing what they're supposed to be doing on content takedowns.

So how confident are you in this, and are you willing to open up that process to more external review?

Thanks.

Guy Rosen: Hey, this is Guy. Thanks for the question.

The data transparency advisory group is a really important part of our effort here. And the focus of our engagement has been on first and foremost, understanding the methodology. And so what we did is, our team spent a lot of time working with this group, providing a lot of details and information about all of the enforcement process, all of the measurement methodologies. And we've been doing this over -- since November or so to make sure that they have all the information to understand how we actually calculate these metrics.

We have more work ahead of us to continue sharing more information and making sure that they understand not only the methodology but how we're implementing it. The -- but the main goal here is to make sure that we are reporting metrics that are meaningful and accurate in terms of are they -- are they the right way to measure the progress of an internet company in tackling content issues.

Operator: Your next question comes from the line of Kurt Wagner from Bloomberg. Please go ahead.

Kurt Wagner: Hey, thank you for taking the question.

I noticed, obviously you talked about adding this new metric around regulated goods, sale of drugs and firearms. I'm curious, do you plan to expand that category in the future? I know there are other illegal things that are sold on Facebook and I'm curious how that is going to expand.

And second, as you guys move the company toward more privacy and encryption in the future, how is -- how is that going to change your ability to actually track this kind of stuff moving forward?

Guy Rosen: Hey, this is Guy.

Couple things. On regulated goods, we have a broad set of policies around the kind of -- the kind of content that is or isn't allowed, including around regulated goods. We've specifically added two of these categories in the report today around firearm sales and drug sales. Our overall goal is to continue expanding the categories that we report on so that we -- over a long period of time -- we are able to explain all of the different reasons why content might come down.

On your -- the second part of your question, within a private environment that we're working to -- that we're working to build, yes, there is a tradeoff between protecting privacy and protecting safety and that's something that every society grapples with. We do believe encryption is an incredibly powerful tool for privacy and we are working to detect bad actors through things like identifying patterns of bad activity or building better tools for people to report bad content to us. We have started to work on this building on the work that WhatsApp has been doing over the years and it's something we're incredibly focused on and working together with outside experts over the next year as we work towards this future plan.

Mark Zuckerberg: Yes, maybe I'll just add something. This is -- this is Mark.

Just to emphasize the last point that Guy made, which is when we made the decision to move all of the private messaging to be end-to-end encrypted, we

also made the decision that we were going to take at least a year and perhaps even a little longer for a consultative process to talk to safety experts and law enforcement and governments around the world in a number of different countries to make sure that we build up the right safety systems for an encrypted environment.

And we recognize that it's going to be harder to find all of the different types of harmful content. We'll be fighting that battle without one of the very important tools which is, of course, being able to look at the content itself. It's not clear on a lot of these fronts that we're going to be able to do as good of a job on identifying harmful content as we can today with that tool. But we think that this tradeoff of protecting people's privacy and giving people world class tools for privacy and security coupled with doing as good of a job as we possibly can working with stakeholders and experts around the world on building up these safety systems before we continue rolling out end-to-end encryption to more of our products is the right path forward.

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

Queenie Wong: Hi. Thank you so much for taking my question.

So I was wondering if you could provide more details about who was creating the huge volume of fake accounts and why? Who were the bad actors? Was it election related?

Guy Rosen: Hey, this is Guy.

So on the fake accounts number, the -- we don't have specific attribution. The larger quantities of fake accounts are driven by spammers who are constantly trying to evade our systems and they're trying to, on an automated basis, create thousands or millions of fake accounts through automated means. The numbers -- as the numbers show, we are catching many of those and that's why the number actually is large.

When you think about fake accounts, there's actually two things we should be looking at within the metrics. The first is how much -- how many accounts

we took action on and that is the number that was around 2.2 billion in the first quarter of this year. There's a few reasons why something might be there. It could be -- it could be that we're removing accounts when they sign up to Facebook. And it could be that we're removing accounts that are already on Facebook.

There's also another category of accounts that we prevent from even created in the first place; for example, by blocking full ranges of I.P. addresses. In that case, there's actually -- the spammers or whoever it may be are actually not able to even contact our systems. And so, we can't actually -- we can't count the number of accounts that we've prevented through that means.

So the first number is content -- is the amount of -- the number of accounts that we have taken action on. The second one -- and it's the one that I think we should all focus on -- is around prevalence. And that is, how many accounts actually made it through and became active accounts on our system; that's the number that we've also reported here.

Operator: Your next ...

Mark Zuckerberg: I think your point on spammers is really -- is really important because, I mean, I think the assumption that a lot of people have is that it's not necessarily commercially motivated.

But a lot of the harmful content that we see, even including things like misinformation that are not as obviously commercially motivated are actually in fact commercially motivated. So one of the best tactics for removing and preventing that downstream is, if you can eliminate the fake accounts or remove the incentive to create fake accounts upstream then a lot of this harmful content doesn't even get created in the first place.

Operator: Your next question comes from the line of Jason Koebler from VICE. Please go ahead.

Jason Koebler: Hi.

In February, Facebook banned white nationalism and white separatism. I'm just curious, how is Facebook going to judge enforcement success on those categories?

Monica Bickert: Sure.

Since we've put that policy in place, one of the things that we've been doing is talking to the money groups -- this is Monica, by the way -- talking to the many groups that we work with around the world about our hate, and violence, and terrorism policies. And so, part of the way that we will assess how we're doing is talking to these groups and their constituents about what their experience is on Facebook.

Another part will, of course, looking to expand the categories of data that we provide in our enforcement report. Right now, we provide some data that is around global terror organizations, and that includes ISIS and Al Qaeda and their affiliates. And that's because when we started launching technology to look at these problems, those were the groups that were the most prolific in producing organized propaganda and pose a large global threat.

Over time though, in last fall I would say, we started expanding our use of technology to also go after certain hate groups. And we have more than 200 white supremacist organizations that we have kicked off of our site. And as we get better and better at using technology to remove that content, that's the kind of data that we'll be able to build into the community standards enforcement report.

Operator: Your next question comes from the line of Edgar Alvarez from Engadget. Please go ahead.

Edgar Alvarez: Hi, thanks for taking my question.

I know that you mentioned that the number of fake accounts you've taken down -- you've been able to do it before they're considered active. But I'm wondering if you working on something that maybe spots them before they're being made or as they're being made?

And then related to that, the ones that are -- actually manage to evade your systems, how were they able to do that?

Thank you.

Guy Rosen: Hey there. Guy.

So as I mentioned earlier, there's a few different levers we use to catch these fake accounts. Some of them we do block before they're even able to create anything or contact our systems by completely blocking that bad actor's access to our systems. For example, a spammer that may be using many different devices to try to open a large number of accounts in an automated fashion. So we're definitely doing that.

Some of the -- some bad actors will manage to get through those systems and try to create accounts. We will catch another large number of accounts at that point of creation and we will remove many of those within minutes of creation. And as you've mentioned, some do slip through and we report the number that slip through and are able to be active.

The way we do this is we consider the activity on the account and we have systems that are constantly running and evaluating the activity that -- and the account, whether it's a time of creation or an active account is performing, in order to try and understand if it's a fake account that is there really just to post spam, for example.

The important thing as we think through this is also to balance that with understanding how legitimate users are behaving on our systems. So for example, sometimes real people might sign up and they may behave oddly because perhaps they're a person that's completely new to the internet or someone joins Facebook, they maybe send out a lot of friend request sand it can look like they're a spammer.

But actually they're very a social person. Maybe they're a young person in a certain country and they're rapidly adopting social media. So we build our system to try to balance, making sure that we're catching the right number of fake accounts but also that we're letting real people use our services

effectively. Because we also don't want to in a place where real users are running into problems and not able to create or use their accounts on Facebook.

Operator: Your next question comes from the line Dave Lee from BBC News. Please go ahead.

Dave Lee: Hi there. Thanks for taking my question.

Mark, you'll have seen the White House recently a launched a tool inviting people to submit examples of unfair censorship on social media. Do you have any thoughts on that move and are you engaging with the White House on that?

Mark Zuckerberg: Sure.

So in general we want to make sure that our tools can be used to help people share any kind of perspective. We want to be a platform for all ideas and freedom of expression is just an incredibly foundational principle that we hold here. We work hard to enforce our rules impartially. We have processes undergoing, including this audit that we're doing with Senator Kyl, to make that there isn't any kind of bias in the systems and if the audit turns up anything, to make sure that we can adjust our processes for doing that.

And of course we're taking steps like setting up the independent Oversight Board to externalize and create more of an independent institution some of the content review to make it so that people can have confidence that a lot of these most important decisions aren't even being made by our company directly.

So this is certainly an issue that we take seriously, I personally deeply care about freedom of expression for all people, and it's an ongoing effort to make sure that all of the systems and processes that this company tries to execute uphold those principles.

Operator: Your next question comes from the line of Issie Lapowsky from WIRED. Please go ahead.

Issie Lapowsky: Hi, thank you for taking my call.

So I was wondering, Mark, if you might be able to respond to the calls from Chris Hughes and also a number of democratic candidates urging a breakup of Facebook? I know Facebook has responded to this, but I think a lot of people would like to hear from you.

Mark Zuckerberg: Sure.

Well, in addition to all of the actual questions around competition which is really what those calls focus on, and I would imagine that everyone on this call uses many competitors to the services that we use whether it's iMessage for those of you in the U.S., or Snapchat, or YouTube, or Twitter, or TikTok, or any of the different folks who -- I mean, the average person here uses -- I think it's seven or eight different services to communicate and consume different kinds of content.

So I think it almost goes without saying that we exist in a very competitive and dynamic environment where new services are constantly coming up. And of course on the advertising side, where we're I think less than 10 percent of the global total ad market, I think on the order of 20 percent of the digital ad market. So I think arguments that we're in some sort of dominant position there are -- might be a little stretched.

But look, I mean we're on a call today talking about safety, and I actually think one of the biggest questions that I -- and reactions that I have when I hear this question is what problem are folks trying to solve when they -- when they raise the question about what -- about what kind of regulation should exist around the internet?

I certainly think that there should be regulation. I mean, that's why I'm out there calling for it because I think we're reaching a point on some of these questions around speech and safety where I don't think that companies by themselves should be making all of the decisions about what should be acceptable speech on the internet. So I fully am behind regulation.

I think the question is what regulation do you want to exist, and what problems are you trying to solve? If the problems that you are most worried about are ones about balancing -- making sure that we address harmful content, making sure that we prevent election interference, making sure that we have the right privacy controls and at the same time that people have the ability to bring their data to other services for innovation and competition and research -- those are to me the areas that I think are the most important social issues right now.

And I don't really think that the remedy of breaking up the company is going to address those; I actually think it's going to make it a lot harder. One of the things that I look at now is the amount of effort, the amount of capital that we are able to invest in all of the safety systems that go in to what we're talking about today, our budget in 2019 is greater than the whole revenue of our company in the year before we went public in 2012 -- just earlier this decade.

In one decade, the success of this company has allowed us to fund these efforts at a massive level. I think that the amount of our budget that goes towards our safety systems, I believe is greater than Twitter's whole revenue this year. So we're able to do things that I think are just not possible for other folks to do. And I don't think it's because companies like Twitter, or others who are really big and important players in the space face qualitatively different issues than we do.

So I think when you look at it, we really need to decide what issues we think are the most important for -- to address and to focus on because in some ways, some of the remedies cut against each other in terms of making progress. And I at least -- I really believe that the fight against harmful content is an incredibly important one. We're fully invested in this and we'll continue to do even more. And that's my view on this.

But I think regulation here is going to be very important. I want to keep working with folks like President Macron and other world leaders to try to build the best and most effective framework as possible for this. And I think you'll continue to see us keep doing that.

Operator: Your next question comes from the line of Damien Leloup.

Damien Leloup: Hey, everybody.

About the proactive detection of bad content, the numbers given in the report are averages. Since your tools use deep learning, I guess they work better when you have more data. So how big of a difference is there between the detection numbers for widespread languages like English or Spanish and others that are less prevalent like Dutch, for instance?

Guy Rosen: Hey, this is Guy.

So the numbers there, you're right, are the overall numbers. We have been working to expand our proactive detection on hate speech to a large number of languages. Right now, we are able to proactively detect 40 different languages. It's important to note that the policies and our enforcement overall are global and span any languages. People can report things to us, our team will review it, and if something violates, it'll be taken down.

We have been focused on expanding and rolling out our proactive detection to more languages and we're at 40 languages at this point.

Andrea Saul: And with that, we're going to have time for one more question.

Operator: Your last question comes from the line of Ed Ludlow from Bloomberg Television. Please go ahead.

Ed Ludlow: Hi, guys. Thanks for taking the call.

You talked about this record of removing fake accounts, the shift towards privacy, and you wanting to increase the involvement of regulators and government. What -- on the ground, what kind of pressure is this having on kind of ad targeting? Because you've given guidance on some of it but what are you seeing on the ground with all these initiatives? Because they're -- there's a focus on shifting towards encryption and privacy, you're dealing with the malicious content, and you want to increase the parties who are having oversight.

So how is that affecting the ad business, that compromise between the three?

Mark Zuckerberg: So I think the specific things that you're talking about, I don't think have that large of an impact on ads.

On the encryption side, certainly we won't be able to see the content of messages but we already aren't using the content of any conversations between people for ads anyway. So it's not that -- so, if we move towards encryption, that's not going to actually impact our practices on the advertising side. Similarly, the multiple billions of dollars that we invest a year in content safety and all these safety systems certainly impact the company's profitability overall because that -- those are -- that's an important line of expenses that have grown quickly.

And I've talked about this on our earnings calls, this is one of the areas that we're ramping up investment for the fastest and where we expect that to continue because this is important and we expect there to continue being a lot of work to do. But I don't think that taking down harmful content or proactively identifying nudity or terrorism affects ad targeting in any particular way. If anything, while this isn't a targeting issue, I think the feedback that we get from business partners is that they want to make sure that Facebook and Instagram are safe environments and making sure that this kind of harmful content isn't there is deeply-aligned with what they want to see us doing.

Of course, that's not the reason why we're -- why we do it. But if anything, I would expect it to be somewhat helpful on the business on that side. But it certainly is a big impact on profitability. And there's -- as always, when you're running a company, some of the biggest decisions that you make are around prioritization. And the opportunity costs are massive. And the people we have working on this are talented, and the teams are deep and good. And we have a real choice every day, of whether we put our best people on fighting and working on some of the biggest issues we face, or whether we put it on other things.

And it's certainly where we've shifted a lot of our best people and a lot of our resources onto addressing these major social issues, which I think is just absolutely central to the responsibility that our company has.

Andrea Saul: Great. Well, thank you, all, for joining us today. Have a good day.

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

END