

Community Standards Enforcement Report Press Call

November 9, 2021
11:00 a.m. ET

Operator: Hello, and welcome to today's Community Standards Enforcement Report 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 one followed by the four on your telephone.

As a reminder, this conference is being recorded Tuesday, November 9, 2021. Now I'd like to turn the conference over to Sabrina Siddiqui who will kick this off, please go ahead.

Sabrina Siddiqui: Hi, everyone, thank you for joining us. You should have received embargoed materials ahead of this call. We are on the record and this call is embargoed until 10 am pacific time. Today you will hear opening remarks from Vice President of Integrity Guy Rosen, Vice President of Content Policy Monika Bickert, 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're also sharing additional integrity and transparency items.

Including the Transparency Report for the first half of 2021, The Oversight Board Quarterly Report for the second and third quarter of this year, and the Wisely Viewed Content Report for the third quarter as well.

With that, I'll kick it over to Guy.

Guy Rosen: Thank you, Sabrina. Hey, everyone. Today we're publishing our Community Standards Enforcement Report for the third quarter of 2021. This report provides metrics on how we enforced our policies from July through September of this year.

We're sharing some new numbers today, which include prevalence for hate speech on Instagram, prevalence for bullying and harassment across both Facebook and Instagram, and the full set of metrics including prevalence for violence and incitement across Facebook and Instagram.

With these new numbers we now report on 14 policy areas on Facebook, and 12 on Instagram. This is our 11th report, which as people here know shares more data and information than any of our peers in the industry. Our work in this space is a multi-year journey.

And we use this report and this call as a quarterly touchstone to update you about our progress. We were the first in the industry to release prevalence metrics, which some other companies now do as well. We believe prevalence is the best metric to evaluate our progress in this work.

Let's talk about some of the new metrics we have today. First on violence and incitement where we've added the full set of metrics for the first time here. There's a range of content that we might remove under this policy. And due to the potentially harmful nature we over index on safety and will removed such content even if it's unclear whether the content maybe was said in jest.

This could range from something serious such as instruction on how to use weapons to cause injury all the way to a joke where one friend maybe says to another something like I'll kill you. In instances where necessary we also will work with law enforcement when we believe there is a genuine risk of physical harm or direct threats to public safety.

In the third quarter we removed 13.6 million pieces of content on Facebook and 3.3 million pieces of content on Instagram under this policy. Prevalence on Facebook was between .04 to .05 percent and on Instagram it was .02 percent.

Our second new area is bullying and harassment where we've added prevalence. This is a unique policy area because it often requires context. It's very difficult to know what is a bullying post or comment and what is perhaps a light hearted joke without knowing the people involved or the nuance of the situation.

That's why in some cases we will require a user report from those who may experience this behavior in order to even remove something which means we may not take action proactively in those cases. And you can see that reflected in the proactive rates in this area with by definition will be lower.

In Q3 the prevalence of bullying and harassment was between .14 to .15 percent on Facebook and between .05 to .06 percent on Instagram. Tracking this prevalence metric over time will aid our work to reduce this type of content on our platforms. And we're also publishing a post today with more details on how we approach our work in this area.

On hate speech, prevalence on Facebook has continued to decrease for the fourth quarter in a row now and in Q3 it was .03 percent, that's down from .05 percent in Q2. We've added prevalence for Instagram which was .02 percent in the third quarter. We continue to see these improvements due to the many ways we work to reduce prevalence (on our apps). There's many tactics we

use and that includes things like removing accounts or entities that violate our policies.

It includes filtering problematic content and entities from recommendations. It includes reducing the distribution of likely violating content or content that's borderline to our policies. It also includes things like adding warning screens (at) post time. As well as prioritizing content reviews so that the most (inaudible) severe content is reviewed faster and of course are tactics also include removing content where we're confident it violates our policy.

These efforts all operate globally and we have a review team that review content in over 70 languages and A.I. technology for hate speech in over 50. We continue to build technology that we have mentioned on these calls in the past that can help us identify harmful content faster across languages and content types.

For countries that are experiencing or at-risk of conflict we have made multiyear investments to build teams that work with local communities, develop policies, improve our technologies and respond to real-world developments on the ground.

Final update on the Community Standards Enforcement Report itself, (I wanted to) update here on the independent audit that we committed to so that people can verify the accuracy of the metrics here.

The audit will cover the fourth quarter of this year. We expect to release the results in the spring of next year. We're currently the only one of our peers voluntarily undergoing this type of assessment.

Two more things before I hand it off. First, we're also releasing another addition of the Widely Viewed Content Report for the third quarter of 2021 here. There's been a lot written about what kind of content succeeds in these feeds on Facebook and this report sheds more light on the types of content that reach the most people.

It spans things like entertainment, cooking, family, but still these reports have also highlighted gaps in our enforcements and demotions which we're taking steps to address. And we have Anna Stepanov here available to answer any questions about this latest report.

And one last thing. The U.S. midterm elections are exactly one year from yesterday. So I wanted to mention that we're focused on these elections as well as several others around the world.

Our strategy builds on the approach that we built over the past several years during elections, including the 2020 U.S. election. No two elections are alike, though, which is why we're working to understand the (top) challenges and identify where we can play a role.

We're also talking with federal government partners as well as local and state election officials to make sure that we're properly preparing for different scenarios. We'll have more to share about this over the next year.

And here, as for all of our integrity work, abuse on our platforms is never static and neither are we in how we address it.

And with that, I'll turn it over to Monika.

Monika Bickert: Thanks, Guy. Hi, everybody. And thanks for joining today. As many of you know, I lead the team that writes our policies on which content is and is not allowed on our platforms.

This is an international team, writing those policies and managing those policies, that works every day in partnership with local experts and organizations to identify how we can best refine our policies to keep up with and to get ahead of new shifts in global dynamics.

To that end, last month we launched a new policy that helps protect people from coordinated efforts of mass harassment and intimidation. And we also broadened the protections for public figures under our existing bullying and harassment policies.

With today's report, we're also introducing prevalence for bullying and harassment on Facebook and Instagram, which we hope will provide greater transparency into how we're enforcing on our policies – how well we're doing that. And we'll also the public to hold us accountable for continued progress.

Guy talked a little bit about how violence and incitement is inherently contextual, and I'll just echo that. And that's something that we see every day.

That people use colloquial speech where they'll say something innocent or even good – like you threw me a surprise birthday party, you're a dead man or I'm going to kill you for that – and there's nothing that is actually violent that is intended. Nevertheless, it's very hard for us to discern intent.

And bullying and harassment is another very context dependent area because distinguishing humor can be so hard, for instance there are words that could be derogatory labels in some instances. So for instance picture some words that can be used to say that somebody is sexually promiscuous, but those same

words are also used sometimes by friends in a joking or even affectionate manner when they're speaking to one another.

We, under our policies, remove derogatory terms related to sexual activity as bullying and harassment regardless of intent because we want to ensure that baseline of safety for all members of our community. So in that sense our policy may cover well-intended uses of such words.

Next I'd like to take a moment to look forward and discuss next year's election in the U.S. Like Guy said, we are already focusing on the midterms, which means we're taking steps to combat election interference and misinformation while also working to help people vote. We're enforcing our policies against removing voter interference content and we'll continue to refine our strategy to combat content that discusses – excuse me – to combat content that discusses the legitimacy of voting methods, like voter fraud claims.

And this is all building on our efforts during the U.S. 2020 elections and we'll have more to share as we get closer to next year's elections.

I know that during the last election there was a lot of focus on the actions we took to remove militias and QAnon from our platforms. Since then we've continued to enforce this policy, and so I want to share some new numbers with you. As of September 14th, we've identified over 1,013 militarized social movements and in total we've removed about 4,000 pages, 20,600 groups, 190 events, 54,900 Facebook profiles, and 8,300 Instagram accounts.

We've also removed about 3,900 pages, 11,300 groups, 640 events, 50,300 Facebook profiles, and 32,500 Instagram accounts for violating our policy against QAnon.

While each election will bring its own unique set of challenges, we're working diligently to apply the lessons we've learned from previous years to elections in the U.S. and other countries in '22 and beyond – 2022 and beyond.

And finally, just to touch on the oversight board, in July we shared our first oversight board quarterly update, which covered the first quarter of 2021, and this update covering the second and third quarters of 2021 were reporting on the details of cases that Meta referred to the board and updates on our responses to the board about the recommendations that they made, along with the decisions in their cases.

So as you may recall, they can make a binding decision on whether or not content should be up or down; they can also make a recommendation to us which is advisory but we respond publicly to those recommendations.

So for instance in this report we discuss actions that we've taken to increase transparency by releasing additional details about our policies and how they're enforced, and we're announcing changes that we've made to our technology to better take context into account when detecting certain types of potentially violating content.

Through these recommendations the board has pushed us to be more transparent about how our systems operate, how we evaluate tradeoff (in) content moderation enforcement, and how we write our policies. Moving forward we'll continue to include our quarterly updates about the oversight board, along with the community standards enforcement report.

And with that, I will hand it over to (Schroep).

Mike Schroepfer: Thank you, Monika. Guy opened this call by talking about how integrity isn't a static issue. The problems we're dealing with are always evolving, and so is the way we approach them. This is just as true for AI systems as it is for the teams they support. So today I wanted to share a little bit about how our AI is evolving to enable the kind of ongoing improvement in enforcement that you've heard about today.

One of the most significant changes is the move towards generalization. Individual AI models that were built to do one very specific task, like check if an English language text post contains hate speech, are being replaced by systems that can operate across multiple problem areas in multiple languages and even in multiple formats like text images and video. And importantly, doing all of this at once.

For example, this year we deployed an A.I. system that can recognize three different related category violations; hate speech, bullying and harassment, and violence incitement.

These issues are often related. So having a more generalized A.I. system that can work across all of them has helped us to catch more violating content and it is outperforming the individual classifiers it replaced.

This contributed to the significant reduction and the prevalence of hate speech in our platforms this quarter as Guy mentioned at the beginning of this call. The drop is even more dramatic over the last 12 months. The prevalence of hate speech on Facebook has more than halved.

It's worth noting that journalist systems like this were almost unimaginable just a few years ago. They're only possible today because of the breakthrough A.I. research that pushed the entire industry forward and a lot of these

breakthroughs took place in our A.I. research labs. It's having a significant impact.

It's also helping to expand the reach of our state of the art enforcement systems. Previously it would take months of training to create separate cost suppliers for our new market. But when we recently replaced some of these with our cross problem systems, we rolled them out to multiple markets in just weeks.

This helped drive better enforcement. For example, the non bullying and harassment content that was (actioned) and (assured us) that was detected proactively each increased this quarter, following the expansion of these classifiers to new markets.

That kind of speed and agility in handling problems on a global scale is going to continue to drive progress in our integrity efforts for years to come. And it's a great example of the industry wide trend of the A.I. generalization and consolidation.

We've all experienced this kind of technological consolidation before. Think about how your smart phone replaced what used to be a standalone camera, music player, alarm clock, GPS navigation systems.

The benefits of consolidation into a single system are similar when it comes to A.I. It's a computing and energy resources because the efficiency of running everything at once in one place.

And it's easier to maintain and improve one consolidated system versus dozens of standalone ones. It means we can quickly deploy research breakthroughs in our operational systems and that research to production pipeline will get more important in the coming years as the path to progress accelerates.

For example, our teams are currently exploring breakthroughs in an area called few shot learning or zero shot learning. This means building A.I. models that can learn to recognize something from just a small number of training examples or even a single one.

It's a very different approach than most of the A.I. in operation today, which has required a lot of label training data before being able to operate reliably on the new tasks. Many of the hardest problems facing our integrity teams are issues that didn't even exist just a few years ago.

And in some cases just a few months ago. So this can become a major factor in improving our integrity efforts. It could make our systems much more

flexible and agile, capable of very quickly enforcing new policies at scale without extensive retraining and calibration.

It's pretty exciting and one I'm looking forward to sharing more detail on very soon. And with that, I'll hand it back over to our host for questions.

Operator: Thank you. We will now open the line for questions. To ask a question press the "1" followed by the "4" on your telephone. As a reminder it is the "1" followed by the "4" on your telephones to ask a question. Our first question comes from the line of Cyrus Farivar with NBC News, please proceed with your question.

Cyrus Farivar: Hi, good morning. Thank you for taking the time. This is Cyrus Farivar from NBC News. I wanted to ask just a couple of – two quick questions. First, why disclose this information concerning bullying and harassment, and incitement to violence now as opposed to another time?

And then I also just wanted to better understand do these numbers with regard to bullying and harassment and incitement to violence, do those reflect only English language content on Facebook and Instagram or does that include multiple languages as well? Thank you.

Guy Rosen: Hey, this is Guy. So, we're constantly working to expand the metrics in this report. And when we have metrics that are accurate, and we believe represent the experiences that people are having on our platform then we strive to release them. And so, that's why we're expanding, these metrics were ready to be shared here.

On your second question, broadly for all the prevalence metrics here, this isn't just – not just for bullying and harassment and violence and incitement, the numbers do cover what we believe is a global representation. So, they're not just about English. We strive to make sure that this represents the global experience that people are having on our platforms.

Operator: Thank you. Our next question comes from the line of Alex Heath with the Verge. Please proceed with your question.

Alex Heath: Hey, Alex Heath here. I was wondering if you guys could talk about the status of the XCheck program and any conversations you've had with the oversight board about it since details of that came to light from "The Wall Street Journal" reporting? Are you considering any kind of substantial changes to XCheck? Thanks.

Monika Bickert: Sure, this is Monika. And first for folks who aren't familiar with what the cross-check program is, basically this is a means for us double-checking,

making sure that we are getting the answer right when we go to enforce a policy on an account that is very high visibility.

And that's something we try to do because if we know that a lot of people are likely to see a piece of content we want to make real sure that if it violates our policies it comes down. And similarly that if it doesn't violate our policies it doesn't mistakenly get taken down.

So, that – the intent behind the program is something that I think is really important. Now, there's been criticism of our execution, and I think some of that criticism is fair. And that's also one of the reasons why we want to make sure that the oversight board has a say in how we think about this program.

So, we proactively referred – in addition to the oversight board hearing cases from users, we can also refer decisions to them. And we referred questions about the cross-check program to the oversight board and they have now accepted that and said that they'll hear that and weigh in.

So, at this point, we're waiting for them to come back with guidance about the cross-check program. And that's part of the reason that we set up the oversight board, and a very positive part of the relationship.

Next question.

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

Glenn Chapman: Oh, hi, everybody. Thanks, good to chat again. I'm sorry if this is a little nebulous but it just seems to be one of the challenges that you potentially face, I'm interested in your perspective on it. Is that even this morning get another email from Media Matters about engagement or the time that certain views are – certain kinds of posts are viewed.

So, I guess how do you – how does this – how does your enforcement activity – how does – how does your AI, in the case of Mike, how do deal with the fact that part of the concern seems to be that there – the more – the more troubling content and the bullying content or the harassment content seems to get so much more engagement. I guess I'm trying to wrap my head around, is that just something that's just so organic and user driven that you can't do anything about it? Or is that something you can address in some form or fashion?

Guy Rosen: Look first of all we should be clear. We have absolutely no incentives whether it commercial or otherwise to do anything other than make sure people have a positive experience. I mean people don't want this kind of bad

content, we don't want this bad content. Clearly our advertisers don't want this bad content.

But the question is, OK, so we not incentivized how do we act on it. And I think it's also just true that our algorithms are just optimized to squeeze engagement. We're constantly refining how we do ranking in order to tackle these problems. And we've been working on ranking specific approaches to integrity for a number of years. And it's had real impact. And it's (demotions) that are designed to reduce the distribution of potentially harmful content or content that's likely to violate our policies, content that's borderline to our policies.

Our ranking generally considers thousands of signals, much more than just how much people engage with something. And like user feedback surveys that gauge people experiences beyond an immediately reaction in the moment. Negative signals as well. And so by building on the one hand ranking that does so responsibly.

Focusing, by the way, particularly on sensitive areas such as civic or political discourse, health discourse where it's valuable and important to take a more responsible approach there. And on the other hand ensuring we're building our systems to tackle, remove, reduce the distribution of harmful content.

We believe that really enables us to build this community and the results are ultimately that we see the prevalence of these types of content and the vast – what it says really these prevalence numbers is the vast majority of content on Facebook doesn't violate our policies and is perfectly good content.

Monika Bickert: If I could just add to that. I mean echo all of that. But in an effort to be more transparent about the work that the teams do put in to demoting that sort of content we published, for the first time, gosh I want to say it was a few months ago, Guy correct me if I'm wrong, our content distribution guidelines.

So if you go into our transparency hub and you go into demotion, how we rank content, you'll see that we actually describe the source of content including sensationalist content and (clickbait) content that those teams are working hard to identify and reduce the visibility of. That's something that my team works alongside Guy's team in figuring out how to tackle these really tough challenges.

But as Guy says, that numbers, the prevalence steps show that we're getting better and better at this.

Operator: Thank you. Our next question comes from the line Alexandre Picard with La Monde. Please proceed with your question.

Alexandre Picard: Hi. I was asking myself why would you not publish figures per language or per country, number of content being taken down, number of human reviewers? And do you plan to do that? That was one question.

And also you just mentioned before the fact that the prevalence was calculated on some global representation of the various languages. How do you calculate that? And can you say a bit more about this? Thank you very much.

Guy Rosen: Hey, this is Guy. So we – let me start perhaps with your second question. The way we think about prevalence in (the figure) here, it is a global figure but it – that means that it factors in a sample that is globally representative. We're confident it's reflective of the overall community on our – on our application.

What we do is we sample across several countries and languages, that includes both large and small regions around the world. We do also endeavor to understand the prevalence, particularly concerning countries where there may be more risk of current events – more risk of conflict.

And then we look at that and ensure that it is generally in line with that global prevalence metric, both to make sure that that number is accurate and of course to make sure that we are also operating correctly in terms of our policies, our enforcement and our product work. And these spot checks do indicate the prevalence in these countries is generally in line with the global metrics.

Our prevalence calculation does – it relies on manual review of samples (of views). We're always working to make sure that we can publish the most detail we can.

And right now, we're publishing it at a global level with a confidence interval around it. That's why you sometimes you see these numbers as between one number and another, because there's some statistical uncertainty there.

And we're constantly working to refine (that) methodology to make sure that we can provide more visibility as well. The same goes, really, for publishing additional figures, to your first question.

We're always working to refine the content in this report in additional metrics. We know people have been asking for more country level metrics. That's not possible for all of these metrics but it's definitely something that's top of mind for us and something that we're – that we're going to look into in the future. And we know that it's something that's valuable to people and we're looking to share more.

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

Shannon Bond: Hey, guys. Thanks for taking the time. Going back to the prevalence of bullying and harassment, so as I understood, Guy, you were saying that this is an area where you're actually relying a lot on user reports, right, because some of this is very contextual.

But then in the blog post about bullying and harassment, it points out that the prevalence metric captures only bullying and harassment where we do not need additional information, such as a report from the person experiencing it.

Does that mean that the prevalence number is potentially low, like, if that's not including user reports? Like, I'm just trying to understand, like, how to understand that figure, what that means about the user reporting and I guess why it wouldn't include the user reporting.

Guy Rosen: Well, the way you can think about this is there are certain violations that our reviewers can basically determine without context by looking at a post or comment and say this definitely violates, this doesn't violate.

There are certain other categories of the policy where something basically wouldn't violate unless it is reported by a user. And those are some of the examples Monika gave where someone might be saying something that could be said in jest, but if someone reports it then we will actually take it down. And you can see in our policies, if you go to our transparency center, it actually details the specific subcategories of the policy where we also need that user report.

Now, if you think about the way prevalence works, we sample – we randomly sample a number of – samples of content and then our reviewers review those, that means that in that process since those have been sampled enough and they're not reported, we'll only be able to determine whether something is violating or those cases where – for those cases (apart from the policy) where we can make that determination on our own.

So that's what that prevalence metric represents, it represents that (sort of – that set of things). The additional content that's also reported and taken down will be reflected in the content action number of course, and you'll be able to see that come through there.

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

Issie Lapowsky: Hi, thank you. Two quick questions. One is you note in the report that the – that you'll begin reporting instances of newsworthiness exceptions at the end of 2022. Given what you've said about midterms and your focus on that and the salience of the newsworthiness exception during elections, I wonder why it will take another year to report those instances?

The second is you note some of the challenges that Facebook has had in keeping up with the oversight board recommendations. What are the steps Facebook is taking to address that and does the board need to sign off on any changes you would make? Thanks.

Monika Bickert: Thanks, Issie, this is Monika, and on the newsworthy notices, there's a lot to think through about how that looks in terms of experience for people and we're still working on that.

I don't have anything further to share right now, but in terms of the broader question on the recommendations – I'm sorry, can you – can you repeat the second part of the question for me? Recommendations from the oversight board, is it how we're responding to them or what goes into the process? (Inaudible).

Issie Lapowsky: Oh you know what, (I'm not sure if you can still hear me).

Monika Bickert: Yes, yes, I can. I can.

Issie Lapowsky: Can you guys hear me? You can.

Monika Bickert: Yes.

Issie Lapowsky: OK, the question was given the concerns that you have about the current response times to the recommendations the board makes, what are the steps you're taking to try to fix that process? And does the board need to sign off on whatever changes you make under their bylaws?

Monika Bickert: I see. OK, yes, so let me just explain a little bit about what goes into that process. When the board has a recommendation we have a group across the company that comes together and looks at the recommendations and evaluates whether or not they're feasible for us.

So sometimes the board – and this might be an example of maybe the most lightweight sort of recommendation from the board – sometimes the – lightest weight meaning easiest to implement – we might have an idea from them about how we should adjust a policy (line).

And that's one thing, and we can look at that in fairly short order. Something that we would do in that case is we would go to our policy forum process, which is a process whereby we get input from stakeholders and experts and people across the company, experts outside the company, and we would bring the board's recommendation into that process.

But that's a defined process and that's something that we can go through in relatively short order. Some of the trickier recommendations are things around how to use technology, how to adjust the way that we use human review, building additional transparency into the product; some of those questions – first of all, they might be really tough for us to actually accomplish, and they necessarily include trade-offs.

And so for instance there might be something that we think, yes, we'd love to do this in terms of the product but we also know that we're trying to build this other very safety focused tool right now and we would have to deprioritize that to prioritize building this other thing that the board would like us to build.

So we find the recommendations very helpful and it still means we can't always get to things right away. So we're trying to get faster in the way that we consider those requests and more streamline in the way that we respond to them and we'll continue to do so but those recommendations are never going to be something that we can just implement or even immediately respond to without taking that considered approach. Next question.

Operator: Thank you. Our next question comes from the line of Queenie Wong with CNET. Please proceed with your question.

Queenie Wong: Hi, there's been concerns raised, including by some of your former employees that maybe prevalence isn't the adequate enough metric. Have you considered any other metrics and how did you settle on prevalence being the metric to put in these transparency reports?

Guy Rosen: Hey, this is – this is Guy. So prevalence represents the experiences people have on our platform. It represents the amount of violating content that people actually review, that actually shows up on someone's screen. And that means that it captures not just what we caught but also – but sorry – but specifically captures what we missed.

And so it's a good metric that represents and is able to hold us accountable for continued work. Now we've been reporting for three or so years now and we're also very glad to see other companies reporting similar metrics, which is violated (view rate), which is potential prevalence. So (it's easy) to think you'll see a couple of other companies reporting that.

It is a massive step forward from where we were several years ago only reporting on things like how long it took to take content down or things of that nature, which don't really represent the right kind of thing that you would like a team like ours to be – to focus on and to improve at.

But look, this is an evolving space and I think it's actually grate that we're having conversations about the next generation or next cut of metrics should be. I don't think we're done yet here even with prevalence, which clearly matters a lot.

And then you can see even today we added prevalence for a number more areas and we want to make sure we continue filling that out. I think there's ideas and there's definitely continued explorations that we're also doing, especially in areas where overall prevalence is perhaps is lower where we can measure sort of the next cusp of where people might still be seeing more of this content.

This is work we're doing and we also welcome and industry conversation about the best ways to evolve this measurement.

Operator: Thank you. Our next question comes from the line of Mark Albert with Hearst Television. Please proceed with your question.

Mark Albert: Hi, thanks very much for doing this call. I just wanted to circle back to something both Guy and Monika talked about on election spreading false claims, et cetera. Guy, you mentioned that you have been, I believe, talking to election officials but more to say about this in 2022.

Monika mentioned combating content that addresses voter fraud claims. I'm just curious if you have an update on any actions you took for the – for the election this month in 2021 and any preview of what we can see for next year. Thank you.

Guy Rosen: Hey, this is Guy. We continually work with election officials at both the federal and state level. I don't have anything to share in terms of who – the small number of elections that we had in the past couple weeks.

But as I said, we will be sharing more as we head into 2022 and ahead of the midterms. We're preparing for that very seriously, of course. And we'll be sharing more about the steps that we're going to be taking.

Operator: Thank you. Our next question comes from the line of Casey Newton with Platformer. Please proceed with your question.

Casey Newton: Hi, I wanted to ask about the Wisely Viewed Content Report. I noticed that some spammy-looking pages are still hovering near the top there. And I wondered if the company is thinking about doing anything in response or if it feels like those links look basically what it wants to see on that list?

Anna Stepanov: Hi, yes, this is Anna Stepanov. Great question; and I'll take it. So, we are in (inaudible) we talked about last quarter when we first released this report. We are using the data and the experiences that this report surfaces as ways to improve our product and the experience.

And the specific case of the pages and links that you're referring to, they don't violate any of our community standards. But we are continuing to examine and understand whether that is the content that our users do find most meaningful.

Monika Bickert: And I think we've got time for one more question.

Operator: Thank you. Our last question comes from the line of Kaley Hayes with Business Insider. Please proceed with your question.

Kaley Hayes: Morning, everyone, thanks for taking some questions. I wanted to circle back to the bit about the oversight board. It seems like from the embargoed report there were maybe some changes coming to how you all would be dealing with the recommendations specifically over the oversight board.

So, I'm hoping you guys could be a little bit more specific about the changes that you're looking for maybe to streamline the process. Or any conversations that you're having with the board currently about changes that might be coming to the way recommendations are given and acted upon?

Monika Bickert: This is Monika. No, I don't really have anything else to add other than what I said earlier about the process and why it takes some time, and how we'll continue to work to streamline and get faster at that.

Guy Rosen: Hey, this is Guy. I can maybe just add there. The – particularly as the board's recommendations pertain to products-related changes, technology-related changes the kind of – the cycle times for those things and the efforts required are quite long.

And the current ways the boards mechanisms were designed, and it's what you'll read in that report, are really designed – focused more on (served fees), one-month processes to evaluate.

And then continued updates where for some of these more technology-centric things, these things can sometimes take many months to better understand and

validate. The way we think about work of our engineering teams that build these things within typically three or six months cycles for some of these larger efforts.

And so, what potentially the report is sharing here is we're having this conversation on how to better structure that relationship to ensure that it can account for that (flavor) of work.

Because of the nature of the time it takes, and the kind of effort that we want to put in to make sure that we're diligently looking, assessing, and responding to these recommendations from the oversight board.

Operator: Thank you. I will now turn the call back to Sabrina Siddiqui for some closing remarks.

Sabrina Siddiqui: Thanks everyone for joining the call. And if you have additional questions just ping our press line or ping me. Thanks.

Operator: Thank you. This concludes the Community Standards Enforcement Report press call. Thank you for joining. You may now disconnect your lines.

END