HOUSE OF COMMONS

Digital, Culture, Media and Sport Committee

Oral evidence: Immersive and addictive technologies, HC 1846

Wednesday 15 May 2019

Ordered by the House of Commons to be published on 15 May 2019.

Watch the meeting

Members present: Damian Collins (Chair); Clive Efford; Paul Farrelly; Simon Hart; Ian C. Lucas; Brendan O'Hara; Rebecca Pow; Jo Stevens; Giles Watling.

Questions 809 - 1014

Witnesses

I: Claire Lilley, Child Safety Policy Manager, Google EMEA; Marco Pancini, Public Policy Director, YouTube EMEA; Rich Waterworth, Marketing Director, YouTube EMEA.

II: Karina Newton, Head of Public Policy, Instagram; Vishal Shah, Head of Product, Instagram.

Examination of witnesses

Claire Lilley, Child Safety Policy Manager, Google EMEA; Marco Pancini, Public Policy Director, YouTube EMEA; Rich Waterworth, Marketing Director, YouTube EMEA.

Q809 **Chair:** Good afternoon. If I could call this meeting to order of the Digital, Culture, Media and Sport Select Committee. This afternoon is a further evidence session in our inquiry into immersive and addictive technologies. We are delighted to welcome representatives from YouTube to give evidence to the Committee this afternoon.

You may be familiar with the White Paper the UK Government have published on online harms. At the heart of that is the principle that social media platforms have a duty of care to their users to make sure they are protected as far as possible from engaging with harmful content. Do you accept that principle?

Marco Pancini: Thank you very much for the question. Let me introduce myself. My name is Marco Pancini. I am Director for Public Policy for YouTube for Europe, Middle East and Africa.

The White Paper is a very important piece of policy-making. For the first time, a Government came with a comprehensive overview on all the different online issues and came also with some interesting solutions. We believe, for example, that the idea of a regulator could be a great opportunity to build the trust between the institutions and the industry in discussing and tackling the specific issues related to the internet. We have seen that in other sectors, and it was a successful way to find concrete solutions. We are really looking forward to participating in the public consultation with our ideas.

In answer to your question about us taking responsibility—we want to take our responsibilities. It is in our business interests. It is a priority for us, as our CEO stated recently. He repeated it at several opportunities. Creating a safe platform, creating a safe experience for our users, is our responsibility.

Q810 **Chair:** You believe that it is your responsibility to ensure that harmful content, content that should not be there, is removed in a timely fashion?

Marco Pancini: Indeed, and let me describe our approach to that. We have a comprehensive approach to tackle harmful content online. We have invested in people. We now have more than 10,000 people working on safety on our platforms. We invested in the best of our technology to make sure that together with people, technologies can identify at scale violation of our policies. If we look, for example, at the first quarter of this year, of the 8.3 million videos that we took down, 6 million were taken down thanks to the machines that we are using in order to send us a referral of violation of our policies. That is very important.

Thirdly, we need to keep our policies updated. We need to keep on keeping track of the development of online harms, and that is what we are doing. In the last six months we updated something like 30 of our policies, and we constantly do that.

Claire Lilley: My name is Claire Lilley. I am Public Policy Manager for Google, with particular responsibility for child safety. In terms of the White Paper, we were really pleased to see the focus on one of the key codes of conduct that the Home Office is driving forward and pushing through, which is on child sexual abuse and exploitation. We really feel that this is an area where we have tried to push the boundaries and lead the industry towards providing some solutions free of charge, and some technological solutions to ensure that that sort of content is removed as quickly as possible.

It is true to say that because that content is clearly illegal, that is relatively straightforward in some ways because the definitions are very clear. The point I want to make is that we have not waited for regulation in order to create those technological solutions. We have been really, as a company, striving to drive forward and to push the boundaries of what is possible and to make that as available as possible.

Q811 **Chair:** I know we have some questions on child safety coming up later on, which probably lets us explore that in a bit more detail. I just wanted to stay on this principle of duty of care. Do you believe that duty of care relates to not just removing content that has been reported to you for being in breach of your platform policies, but also to proactively identifying that content for yourself, not waiting for referral?

Marco Pancini: That is a very important development. We made important progress in this area in the last two years. We have now an internal team that is proactively looking for new trends of online harms, studying continuously the challenges that are put forward by people that want to exploit our platforms. The head of this team is a former deputy intelligence expert from the FBI who is reporting directly to our CEO. We have a team, part of this broader team, also here in London, and that is a way for us to work in a proactive way and study how online harms are developing. That is an example.

Just to give you another example, when we identify a video that is against our policies—some of the ideas that are of your concern, of course—we take these down. We make sure that if the same video is uploaded again on our platform, we proactively block the video. We do not let the video go up online. I think that being proactive is part of the responsibility that the company likes to take.

Q812 **Chair:** Looking at the Christchurch terror attack and the way in which platforms like YouTube were exploited to spread footage of that, do you feel the systems you have in place are adequate?

Marco Pancini: First of all, it was a horrific moment for us. It was a shock. When we woke up in the morning, we took over from our colleagues in the region to work on this terrorist attack. We were talking with the New Zealand Ambassador a few days ago in Brussels, and exactly this was the question: what kind of measures can companies like YouTube put in place in order to make sure that we can be more prompt in making sure that these kinds of content are not spread?

To answer your question, of course what happened on the day was very specific, but still there are some areas where we were able to be on top of the challenge. First and foremost is working together with the authorities and the law enforcement. In the past years we developed a relationship and contacts, so now, when something like this happens, in a matter of minutes, law enforcement can get in touch with us and alert us that something is happening. For example, we identified the original video in a matter of minutes. That is working.

Another thing that also is working is the use of technology to identify the video itself but also part of the video, and that is what happened. At the end of the day we identified something like 20,000 pieces of the same video that were uploaded on our platform and we took them down.

The third area where we have good collaboration in place is between the different companies of the sector. We exchange information, but of course there are areas where we have to improve. One is our policies. We decided, because of the experience of Christchurch, that if events of this kind happen, instead of letting our machines identify pieces of a video that is depicting an attack and sending those videos for review to our team, while at the same time leaving the content to be uploaded online—now, if something like this happens and the machine identifies pieces of the video, we block these videos. We review the videos. Only after we check that it is an authoritative source that is reporting the facts do we let the content go up online.

Q813 **Chair:** I know we will ask some more specific questions on Christchurch later on, but I just asked at the beginning in context to this. Often when we have had evidence sessions with companies like yours, we are given lots of statistics about how much harmful content is identified and removed. Then something like Christchurch happens, and it appears that those systems, when they are required to operate at scale and at speed, are not operating at a level that we would find satisfactory. Do you feel that that has demonstrated that more needs to be done, and that, as they operate at the moment, that level of adequacy is not high enough?

Marco Pancini: One piece of evidence to show that we think something more has to be done is that today we are signing, together with the President of France and the Prime Minister of New Zealand, a pledge, a call to action, to make sure that we, as an industry, come up with solutions in order to be ready in cases like Christchurch. To answer your question, yes, we have a plan and we have measures that we want to

implement, including making it a privilege to livestream content on our platforms, because of this interaction.

Q814 **Chair:** I suppose the question a lot of people ask is: why does it take a terrorist atrocity like that to make companies like yours sign a pledge?

Marco Pancini: We fully understand the question. What happened on the day for us was a huge learning. We learn from the failure of some of the systems that we put in place, and that is why I am saying that changing our policy is a way to address the risk that something like this can happen in the future, and make sure that the content that is related to a video uploaded of this nature is blocked as soon as possible.

To show that that is something that we are working on seriously—for example, when we saw what happened in Sri Lanka, in a situation that had some similarities, we were able to be up to the threat.

Q815 **Chair:** From our point of view, when we last had evidence from YouTube in February last year, we raised questions about harmful content and the removal of harmful content. We were given impressive numbers and told that everything that could be done is being done, and then something like the Christchurch attack happens, and then the response is, "There are new pledges, new commitments, offers to do more". That leaves us wondering, why does it always take some terrible event to make the companies recognise that they need to do more in this case? I have some more questions on that later so I will move on.

We are talking about harmful content and a duty of care towards content itself. Do you think the duty of care should also mean the company has a responsibility to make sure people aren't watching for excessively long periods, particularly younger users? Do you think there should be a function where people are prompted to take a break from watching?

Rich Waterworth: I can talk to that. To introduce myself first, I am Rich Waterworth. I am the Marketing Director for YouTube.

I will come to some specifics about products and tools that we have built to do exactly what you are suggesting, but to set the context, to emphasise, as Marco was saying, we really, truly believe that we have and want to have a responsible duty to users to make sure they all have a really positive and healthy experience on YouTube. The vast majority of users do. As part of that responsibility, we recognise that there may be instances where people need some help in managing their relationship with any media, including YouTube.

We have built and now have a broad range of tools across YouTube, across Google properties and across Android, the mobile platform that Google runs. On YouTube specifically, we have a series of wellbeing tools, one of which is exactly what you are suggesting. Users can set a timer to deliver a notification after a set period of time. The video will pause and prompt them to think about whether they need to take a break. That notification can be turned on by users in settings. The default for that is

that that would be delivered after 75 minutes of continuous viewing, but the user can adjust that to set it for a higher or lower time. That is one tool that we have.

We also offer users a dashboard so that they can see how much time they have spent watching YouTube by day or by week, so you can get full daily visibility of how much time you have spent on the platform. We do absolutely recognise that tools like that may be valued by users, and we have built those and we think that is an important thing to offer.

Q816 **Chair:** I have seen it reported quite widely that 70% of the film time that people watch on YouTube is not films that they have searched for but films that the platform has recommended to them. Is that correct?

Rich Waterworth: Yes, that is correct. That, as you say, has been reported externally. Around 70% of consumption is coming from recommended videos, and the remainder, about 30%, is coming from a combination of people searching for videos, people going directly to videos on channels or on their homepage that they know, or people subscribing to channels—that is an indication of saying that you want to follow a channel, in which case that channel will show up in your feed on YouTube. A user can also choose to be notified when a channel uploads a new piece of content. That is the breakdown of consumption.

Q817 **Chair:** The functionality is that—perhaps you can confirm this is correct—you use artificial intelligence to assess the viewing history of the user against the profile of the new piece of content, or any piece of content on there, and use that to make a match that is then used to make a recommendation to the viewer. Is that how it works?

Rich Waterworth: That is broadly how it works. Up Next is a tool that takes, as you say, information about the video that is being watched and will find a range of suggested next videos to recommend. If a user is logged in, then the algorithm will use information related to prior content that has been watched or that it knows that user might be interested in.

We have made significant changes to that algorithm over time. Over the last year or so, that algorithm has evolved from being focused on recommending a piece of content that someone would go on to watch—users gave us feedback that they were finding that there were too many recommendations in the list that were too similar. We do surveys on the site and we look at other signals such as how much of the content that they are being recommended they watch. Do they like that content? We realised that a narrow set of recommendations was not creating as high satisfaction among users as we wanted, so we have evolved that algorithm to recommend a broader range of content.

We have also evolved that algorithm to demote certain types of content, such as click-baity type content. There are a number of ways in which that algorithm has evolved and continues to evolve.

Q818 **Chair:** It is quite an astonishing statistic that 70% of what people watch

is content that is recommended or prompted to them by YouTube, rather than things they are actively looking for. It would really demonstrate the power of the platform to keep people watching for as long as possible. They maybe go on to look for one thing or they are being recommended one thing, and then they are watching more and more and more as it is recommended to them by the platform. It would suggest that the platform is remarkably successful in keeping people on for as long as possible.

Rich Waterworth: Up Next is a feature that users value really, really highly. There are a number of different ways that it is used. A very common use case on YouTube or type of consumption on YouTube is around music. It is very common for people to listen to music on YouTube, and that can be for a period of time, in a similar way that they may listen to radio. Users find that the ability to have effectively a linear stream, where one track or one piece of content plays after another in a similar way to a radio station, is a very useful thing. In a similar instance, there are lots of examples where people are using YouTube to do cooking, when they have their hands busy, and being able to have a device and watch one programme after another again replicates the experience of a linear television channel. It is a highly valued feature for users.

Q819 **Chair:** Since YouTube has introduced the feature whereby AI is used to select videos on the homepage that people are most likely to want to watch, what has been the impact on viewing patterns since those changes came in in terms of the amount of consumption?

Rich Waterworth: Our main focus has been on how satisfied users are with the experience.

Q820 **Chair:** Yes, but I did not ask that. Are they watching more as a consequence of this feature being introduced?

Rich Waterworth: We would have to take that away and get some detail on specific patterns, because YouTube consumption has grown over time from a combination of factors.

Q821 **Chair:** This is from a new book, Chris Stokel-Walker's book, "YouTubers". I don't know if you are familiar with it. It states that in the three years since Google Brain began making smart recommendations, watch time from the YouTube homepage has grown 20-fold. Is that correct?

Rich Waterworth: I would have to check if that number is correct.

Q822 **Chair:** Does it sound right? You are a marketing director. It is pretty core to the business, I would imagine.

Rich Waterworth: I would have to go back to that specific and we can get details and come back to you. I do not know if that is—

Q823 **Chair:** If that was the case, that is a pretty dramatic shift in viewing patterns as a consequence of using AI to more accurately guess what

people want to watch. You might say it is clearly a feature people like. Evidently so, but I just wonder if people are being drawn into ever-lengthier sessions of watching because the machine is so clever at constantly tempting people with new things.

Rich Waterworth: It all comes back to whether users are finding content or being recommended content that they enjoy and having a good experience. There are a number of factors that impact on how much YouTube is consumed. One of the really big factors over the last number of years has been the growing availability of smartphones, and YouTube consumption on mobile devices is now a very large majority of consumption. People have that availability. There are a pretty broad range of factors that have driven and continue to drive YouTube usage and growth, and that is one of them.

Q824 **Chair:** A 20-fold increase in three years is pretty dramatic, even in a fast-moving market.

Rich Waterworth: If that is the case. I don't know if that number is correct.

Q825 **Ian C. Lucas:** Can I just ask you, Mr Pancini? Does YouTube now accept that it is legally responsible for the content on its platform?

Marco Pancini: If we look at the legal framework in which we are operating, which is the Commerce Directive, there are different responsibilities according to the role that the platform is playing in relation to the content. I will give you two examples. If a user is uploading a video on YouTube, the user is the content provider—they are the author of the content and responsible for the content. We are also carrying a responsibility, which is to make sure—

Q826 **Ian C. Lucas:** That person may be paying YouTube or you are securing income from that process, so you take that responsibility. You are happy to take that responsibility. My question to you is: are you responsible for the content of your platform? I want to be clear on this because I think it is a very important question.

Marco Pancini: Indeed. Our important responsibility in relation to the content that is uploaded on our platform is to make first and foremost available tools that allow every user, everyone that is interested, to report the content. Whenever the content is reported to us we need to take action, but on top of that we have also responsibility, which is sometimes it goes beyond—

Q827 **Ian C. Lucas:** This is why I am asking you the question, because there was a time when platforms were saying, "The content of the platform is not our responsibility". That was largely the position that we were being told perhaps back in February of last year. That position seems to me to have changed. That is not your position anymore.

Marco Pancini: Let me clarify because it is a nuance, but it is an important nuance.

Q828 **Ian C. Lucas:** It is important, but I want a clear answer to the question.

Marco Pancini: There is a difference between liability and responsibility. Liability means that if I am posting content on a platform online, on every platform, which contains an illegal activity, something that is illegal, I am responsible for that as an individual. The platforms carry some responsibilities as well, so I am liable, just to make it very clear, to stick with my distinction. I, as a user, am liable for the content of my video. The platform is responsible for a different aspect, which is not the content itself, but is making sure that first and foremost we provide a safe experience for users.

Secondly—

Q829 **Ian C. Lucas:** No, no, no. Hold on. This is not a pick and mix process. Are you lawfully responsible for the content of what is on the YouTube platform? That is what I want to know. Yes or no?

Marco Pancini: We are not liable for the content, no. We are not liable for the content unless we are producing the content. For example, we have a service called YouTube Premium. We produce the content and we are liable for the content. Users or any individuals that are uploading content online are responsible for the content.

Claire Lilley: I think, perhaps, if I may, the shift that you are referencing is something that I have felt within the company, which is a real shift towards responsible innovation and being a leader in making sure that we are being responsible in the way that we are designing and developing the tools and the way that we are thinking about our users and the need to protect them.

If I think about children, for example, we have really tried to be thoughtful and responsible in the development of tools that are specific to children. For example, there is our YouTube Kids app, which I can talk you through, or our Family Link app, which enables parents to manage their children's use. I would say that responsibility is very much a theme that is running through the company at the minute and it is running from top to bottom.

Q830 **Ian C. Lucas:** That is good to hear, but you used to be a teacher. My wife is a teacher, and she talks to me a lot about social media and the impact on children, and I know your background in the NSPCC. Isn't it a bit late at this stage to be considering the responsibility towards children? What I see in my constituency surgeries are increasing numbers of parents who are concerned about the impact of social media on their children and on their children's mental health. It seems to me that an awful lot of people have made an awful lot of money before we started to think seriously about the impact of social media on the mental health of our children. Do you agree?

Claire Lilley: I think that is a very legitimate concern but I do genuinely think it is an area where Google has really tried to push the boundaries

and to lead the industry. The creation of YouTube Kids, for example, was a recognition that YouTube main is not appropriate for young people. They need a more contained, safe environment. The investment that has gone into that is not to be underestimated. It is significant, and rightly so. It is also a sign that we want to show what is possible so that other players in the industry can do the same and provide children with good, quality content that is suitable for their age and stage of development.

There are a huge range of issues. The Chair asked before about the impact of screen time on young people and so on. The evidence base is very much still developing. We certainly want to take a responsible approach to young people and their mental health. For example, we saw a report earlier this year from the Oxford Internet Institute, which looked at the experiences of 300,000 adolescents and found that the association between their technology use and their wellbeing was less than 0.5%. In fact, wearing glasses had a greater negative impact on their mental health.

It is a very nuanced debate. There are a lot of factors at play here. What we as a company try to do is be very evidence-led in our approach, but also really understanding our responsibility to protect young people in particular.

Chair: Just a note for the record. The statistic I read out earlier about the 20-fold increase in use of the platform since Google Brain technology was being used to recommend videos on the homepage was also reported by the Verge in 2017 in an article that was focused on YouTube. It also included the statistic about the fact that 70% of videos are recommended. It is another source for it. It doesn't seem to have been a fact that was challenged by YouTube either at the time or since.

Q831 **Simon Hart:** We have moved on but a quick reference to Mr Waterworth's comments earlier on about a timer prompt process, which you referred to. I wondered if you had any indication of how many people, particularly young people, use that. How effective is it?

Rich Waterworth: It is a fairly new feature in the last year, just under a year. We do not have usage data on it.

Q832 **Simon Hart:** You have no idea if it works at all?

Rich Waterworth: As the marketing person, I do know that we have made tens of millions of people aware of that through promotional messages.

A couple of things I do know that give an indication of the use of these kinds of tools. Within the Android set of wellbeing tools there is a wind-down feature, which enables Android users to set effectively a time limit when they want to be either asleep or going to sleep. You set that time, and from the time you set that, you do not get any more notifications on your phone, and the phone screen goes to black and white, to greyscale. We have seen just over a quarter—27%—reduction of usage among

people who enable that setting. We do see that these settings have usage, and I can say that we get qualitative feedback from users that they really appreciate the fact that we have developed this and are telling them about this.

Q833 **Simon Hart:** As yet, you do not have anything you can give the Committee on the particulars? I maybe got you wrong, but did you say a 75-minute limit?

Rich Waterworth: Exactly.

Simon Hart: You have no idea whether that actually—

Rich Waterworth: I am happy to take that away and see if we can find any more detail on that, but I do not have that with me now.

Q834 **Simon Hart:** That would be useful. Very quickly on that, is the fact that you are introducing it a sort of admission that you think there may be something in the evidence that suggests that long periods of time onscreen may be having a detrimental effect on your customers? What is prompting you to go down this road?

Rich Waterworth: The reason behind creating tools like this is really linked to a very core design principle within the company, which is about giving users transparency and control over their usage across many different products, and it is definitely the case in YouTube. It is one of many ways we give users control over—

Q835 **Simon Hart:** It must be in response to something, and that is what I am trying to get at. I am suspicious that you are responding to a growing awareness perhaps that elongated periods of time onscreen may be doing a certain amount of harm. I wondered if these new innovations were a reaction to that or is it really just a gesture of goodwill on your part?

Rich Waterworth: As Claire was saying, the evidence base of screen time and impacts and so on is growing and developing, and we take a very keen interest in that and we think that is very, very important, but building those tools is not a response to a particular thing. We are very regularly building and releasing new tools and features of all different kinds, and that is one wellbeing tool I mentioned among a number within YouTube and a number within Google and Android. These things are developing both within the company and within the industry at large.

Claire Lilley: I don't know, I am just surmising, but it is potentially because mobile devices have become in the last five years so ubiquitous in our lives, and we all need better ways of managing that. That has been a theme running through the tech sector generally over the last two years, I would have said. We saw other companies introduce similar tools.

Marco Pancini: One general principle of design of the Google product is to make technology transparent, to make sure that technology is there when we need the technology, not to become a slave of technology. That

is something that our CEO, Sundar, has repeated several times. I think these kinds of tools are going in the direction of making sure users have full control over their day-to-day technology. They can activate technology when they need it but they are not hooked on or a slave to any technology.

Q836 **Jo Stevens:** Mr Pancini, can I just ask you a question? Earlier you were asked about whether or not you have systems to proactively take down content, and I think you said that you do. Where you have taken content down previously and it pops up again because somebody has uploaded it again, you have people who will then take it down again. Can I just clarify? Are those the only circumstances in which you proactively take down content?

Marco Pancini: I gave you two examples. One example is the case when we take down a piece of content and the same content is coming up again. That is a proactive use of our detection tools.

Jo Stevens: No, I understand that bit.

Marco Pancini: The second is, for example, the work of the Intelligence Desk. Maybe I can give you—

Jo Stevens: Yes, so you do have another system. That was what I wanted to clarify.

Marco Pancini: Correct.

Q837 **Jo Stevens:** You mentioned that you have that headed up by a former FBI official. How many people do you employ on proactively looking for content that is harmful so that you can take it down?

Marco Pancini: To clarify this point, the people that we have, these experts of the Intelligence Desk, are not just facing the platform looking for the illegal content because that would not be effective. This is work where indeed the machines make scaling a much more effective way. Just to give you an idea, of the 225 minutes of content that we took down in the first quarter of this year, 90% of it was flagged by machines.

Q838 **Jo Stevens:** Yes, but can you answer my question about how many people you employ? They may not be solely employed on identifying content to proactively take it down, but how many people do you employ where that is part of their role?

Marco Pancini: There is no one that is looking actively for violations, not because we do not want to have someone doing that but because it will not be effective. The experts that I mentioned to you before search for information about new trends of online harms using different sources. Something that they use is the observation of our platform. For example, the flags that we are receiving, but also they use other sources. For example, the interaction with law enforcement. The interaction with the experts and the trusted flaggers that are reporting content to us. They are experts in their own domain, so they look in academic papers. They

speak with other experts. They collect all this information. They understand, "Okay, this is a new trend of online harm that is happening on our platform. Let's start to find the solution with our machines to identify these and fight against them".

Q839 **Jo Stevens:** How many people do you employ who do that work that you have just described?

Marco Pancini: In general, we have 10,000 people working on security. Part of them are working in this team.

Q840 **Jo Stevens:** How many?

Marco Pancini: I do not have this information with me. I am very happy to come back to you to clarify this information.

Jo Stevens: That would be very helpful. Thank you.

Claire Lilley: Could I outline some work that we proactively do in relation to child sexual abuse imagery? Just very briefly, we have three specific types of technology that we use to identify this content. We have one machine learning AI, whatever you want to call it, which identifies images that are already known to authorities, to the IWF and so on, and we run scans and remove those. The second one is a product we pioneered called CSAI Match, which looks for video content, which was industry-leading at the time. The third one is a new piece of technology that we just announced last year, which will look for new images.

The reason it is important for me is these are images that have never been seen before, so probably they are less likely to be historic images. They are more likely to be images of children who are being abused today, and therefore offer us a greater chance of rescuing children who are being abused today. Those are the three things in relation to images.

Fourthly, we also do proactive work in order to identify in comments any words that paedophiles might be using, and sometimes these can appear to be innocuous terms and sometimes they are very context-laden, but sometimes they might be different in different languages. We work with a network of what we call trusted flaggers across the globe who feed us the on-the-ground intelligence about what the latest terms are that are being used so that we can feed those into the classifiers, proactively identify these comments and remove them as quickly as possible.

Q841 **Clive Efford:** Mr Pancini, can you just explain to me so I understand a little bit better, how do you monetise time on your platform? How does it make money for you?

Marco Pancini: The business model of our platform is linked to advertising. Advertisers that find it interesting to promote their products on the platform can buy the possibility to display their advertising on our platforms. All the revenues that we get from advertisers, we share them. Actually, we give to our creators the vast majority.

Q842 **Clive Efford:** Let's say I spent an hour on YouTube. Is it that you get paid for every advert that pops up on my screen? How do you make your money out of me seeing shoes or double glazing being advertised?

Rich Waterworth: The amount of money made is not linked to the time that you spent watching something. The way that advertising works is that some videos you watch will have usually a video advert before them—not all of them—and that will depend on whether an advertiser wants to reach a user such as yourself at that point in time with a certain product. A majority of videos will not have one of those, but the most common form of advertising you will see will be a video advert before a video you are going to watch. There are two types of video adverts. There is a very short video advert, which is usually six seconds, and there are longer video adverts. The longer video ads can be skipped by the user after seven seconds. If they are not interested in that, they can skip that and go straight to the content.

There are also other advertising opportunities on the site. For example, sometimes there will be what we call a companion ad, which might be below the video that you are watching, which might offer information about a related product or service. Those are the most common types of adverts.

Q843 **Clive Efford:** Yes, but what is it that triggers payment to you?

Rich Waterworth: Every time an advert is served and viewed by a user, the advertiser will pay as part of the package that they have.

Q844 **Clive Efford:** The longer someone spends on your site, the more you are likely to earn?

Rich Waterworth: As I said, ads being served is not linked to time spent on the site, but that is how the advertising payments work.

Q845 **Clive Efford:** Mr Pancini, you told the Home Affairs Committee that your artificial intelligence is very precise, so how can you be confident that it is precise enough to deal with critical events such as the Christchurch shootings? What are you doing to improve it?

Marco Pancini: Indeed, if we look at what happened on the day, the machines that we have in place were able to identify all the different pieces of the original video. First and foremost, the original video was identified together with the police and we were blocking it, so it was not possible to upload the same video on the platform.

What also happened is that our machines were identifying when users were uploading different pieces of the video, but some of these pieces of videos were also included in reportage from the BBC, from other sources, authoritative sources. Indeed, for us it is important to have on our platforms also authoritative sources, if a user is coming to YouTube and wants to know what is happening in Christchurch and searches for Christchurch, we want the user to see authoritative sources reporting of

the facts. For this specific reason, when the machines were identifying the pieces of the original videos, the videos were going online and were flagged for review to our people. Because of the amount of content that was posted, it was not possible for us to keep up with this. That is why we changed our policies, and now if something happens again these videos will not go online. They will be reviewed by our people and only after approval they will go online.

Q846 **Clive Efford:** Just so I am clear about that, how will you know to delay posting that content? How do you know it is harmful? What is triggering it for you?

Marco Pancini: Let us go back to this example. When a video is reviewed and we see that it is against our policies, we apply a fingerprint, so a technological watermark, to the video. We put this video in our database of videos that are against our policies, that we identified as against our policies. When a user is trying to upload the same video, the system recognises the match between the original video that we took down and the video that somebody is trying to upload, and then the video is blocked.

For example, if somebody is just uploading a part of the video, so not exactly the same video but a part of the video, there is a light that comes up and says, "Wait a minute. We need to review this video because it is linked to something that we saw before and we decided it was against our policies".

Q847 **Clive Efford:** Correct me if I am wrong, that sounds like how you respond once you have taken the initial content down.

Marco Pancini: Correct.

Q848 **Clive Efford:** How do you identify and how quickly can you identify harmful content of that nature, or of any other nature for that matter, to take down the initial posting?

Marco Pancini: In the context of terrorist content, it is a matter of less than one hour usually, in the relationship that we have with law enforcement, with the relationship that we have with experts like the Counter-Terrorism Internet Referral Unit.

Q849 **Clive Efford:** Is that in response to being contacted by authorities or is that your own monitoring? How are you generally alerted to the presence of content on the platform?

Marco Pancini: Usually it is happening in real time. With the kind of relationship that we have with the Internet Referral Unit or Counter-Terrorism—

Q850 **Clive Efford:** You wait to be told to take it down.

Marco Pancini: It is happening in real time, so the content is identified, is sent to us and we take action against the content.

Q851 **Clive Efford:** Just to be clear, there are instances where you act without being asked or required to take harmful content down?

Marco Pancini: For known videos, for example, of terrorist content, we act. If we recognise the video, if the video is already in our database, we block the video. For new content, we work together with law enforcement and take prompt action.

Q852 **Clive Efford:** If an individual has a history of posting, let's say, racist material on YouTube, is it just the racist content that gets taken down, or does that individual get taken off the platform?

Marco Pancini: First and foremost, if the single violation itself is so bad that we need to take action, we take action directly on the channel of the user, so the user is not able to use YouTube anymore. If it is a normal violation, then after the third violation the user is kicked out from the platform.

Q853 **Clive Efford:** The person who carried out the atrocity at Christchurch urged people to sign up to a site, PewDieWee.

Marco Pancini: PewDiePie.

Q854 **Clive Efford:** Yes. That individual has been accused in the past of posting anti-Semitic and racist material. Is he a suitable person to be on YouTube?

Marco Pancini: Every user has to comply with the same rules. In the context of PewDiePie, as with any other user, we apply the same set of policies. When we identified that PewDiePie was uploading content that was against our policies, we took action against his content. We apply exactly the same rules that we apply to every user also to PewDiePie.

Q855 **Clive Efford:** Against the content but not against the individual. He is still able to post on YouTube.

Marco Pancini: On this specific instance we took action against his content. He never happened to be in the position to have his channel suspended.

Q856 **Clive Efford:** Is there a wider moral issue, if people know the views of an individual that uses YouTube regularly, that you are normalising people who hold those extremist views? Does that cause concern for YouTube?

Marco Pancini: When we look at the content on our platforms, we make an analysis on each and every piece of content that is uploaded. That is the way for us to maintain an objective approach to the behaviour of the users of our platform, which means that we have a very strong set of policies. We apply the law. If we looked at the content that is uploaded on our platform, if the content is against our policies or of course against the law, we take action against the content. We do not look at the views of the individuals that are on our platform, even if in some cases we can disagree with them or they can be controversial. We look at the content.

Q857 **Clive Efford:** Is the viewing and following of extremist content profitable?

Marco Pancini: In all the cases that you made, so a user that has abused our platform—

Q858 **Clive Efford:** Is there a big customer base that uses YouTube that watches extremist material? Is that a disincentive for you to take it down?

Rich Waterworth: I can answer that as well. It is not in our interests and it is not good for our business. It is not good for our brand. It is not something that our advertisers want and it is not something our advertisers are serving ads against.

One of the very important principles that guides how we deal with transgressions against community guidelines is that monetisation on YouTube is a privilege. It is not a right to everyone. Not all creators are able to monetise and monetisation is not granted to all. That is not something we profit from or want to, and it is not something that our partners want. It is absolutely not part of our commercial—

Q859 **Clive Efford:** If you do not profit from it, why is it there?

Rich Waterworth: YouTube is fundamentally, first and foremost, a content platform. We are proudly an open content platform. That brings enormous benefits to the world. People are learning things from people they have never met. They are finding new voices. Diverse users and voices are able to access audiences around the world. There are enormous benefits from that, but it is very, very important to qualify that it is not openness without any limits. We do have limits. We have guidelines that Marco has talked about.

We introduced in 2017 a category of content called limited state. That is a category of content where someone may not have gone past our community guidelines, so a piece of content may not be violating community guidelines and it may not be taken down, but it might be close. It might be highly controversial. In that case, we would put that piece of content in what we call limited state, which means that it cannot be monetised, there are no comments around it, there are no likes, and it cannot be served up as a recommendation to any users. We recognise that there are lines, and we try and we draw a line where the openness of our platform ends. As Marco described, we work with third parties, with society and with academics to review those policies and make sure we have the lines in the right place, and then we are investing significant amounts to enforce those very aggressively.

Q860 **Clive Efford:** In answer to earlier questions about how many people are monitoring this, you could not answer.

Marco Pancini: Ten thousand people.

Q861 **Clive Efford:** Ten thousand people, but when we asked how many of

them are watching the content, enforcing these new policies that you have put in place, you cannot tell us. Why is that?

Marco Pancini: No, that is the 10,000. Just to clarify, the number of 10,000 is about the number of people who work every day to keep the platform safe, including people who are reviewing flags we have received from users and flags that are generated by our machines.

Q862 **Clive Efford:** Can I ask you about your algorithms? YouTube is yet to reveal how its algorithms fully operate. Is there an opportunity to be more transparent and thus improve both the platform's algorithms and the trust in them?

Marco Pancini: Transparency is absolutely very important. For example, when we look at Google, on Google we publish a website that explains how search works, where we explain in very plain words how the algorithm works. This approach is something that we are considering also for YouTube, taking into consideration also a piece of legislation that was approved recently by the European institutions called platform-to-business, which will imply specific obligations of transparency to platforms that provide services online. That is the direction we are taking. Taking into consideration the information we gave you today, for example, on how recommendations or Up Next works, our intention was to let you understand that is not secret. That is generating why we are suggesting a video, but it is based on previous views or previous videos that the same users or a user that has similar behaviour have seen before.

Q863 **Clive Efford:** How do we ensure that the algorithms are not biased? Without transparency, how can we be sure that that is the case?

Marco Pancini: Transparency is very important, so again—

Clive Efford: Why haven't you—

Marco Pancini: That is the direction of travel, transparency, of course. Also in terms of algorithms, we made huge improvement in terms of looking at the data collecting implications of AI. I am very happy to come back to you, consulting our internal expert on all the works that are done.

Q864 **Clive Efford:** Do you accept that transparency in this area is absolutely essential for people to have confidence in the datasets that they are trained on and how the algorithms operate? Do you accept that?

Marco Pancini: I accept that transparency is important.

Q865 **Clive Efford:** You have yet to make that data available.

Marco Pancini: I am accepting that transparency is important. I will tell you more that in terms of the dataset, we are pioneering as Google, as the wider family, the work in this direction by making something called TensorFlow. Most of the datasets that we are using for our machine learning analysis are available for researchers to look into them.

Q866 **Clive Efford:** If we are going to regulate in this area, how is it going to be possible to regulate if we do not have that transparency? How do we know what we are dealing with?

Marco Pancini: I think here the key is to find the right balance between transparency and maintaining something that I think is also a fair principle, especially for other companies that do not have our size, industry secrets. Finding the right balance between transparency to final users and to the wider public and protection of important industrial secrets is the future in this.

Q867 **Clive Efford:** One final question. Do you use Google search to inform your algorithms to recommend films or videos on your site?

Rich Waterworth: If you are logged into YouTube, your usage on other Google products, including search, could be taken into account, in particular around the advertising that you might be served. We have a dashboard called myaccount.google.com where you can go and see precisely the data that YouTube or Google has and uses about you in relation to all of our products and services. For each and every one at a granular level you can turn that on or off. For example, you could see that YouTube is saying you are interested in cars because you have watched a piece of content around cars or you have consumed car content somewhere else. You could choose individually and say, "I do not want to be targeted by advertising around cars", so there is a very, very granular level of control and visibility of what signals are being used.

Q868 **Chair:** Just a couple of quick questions on Christchurch, just to clarify a couple of things. You talked about the removal of content. Have you closed channels or cancelled accounts of anyone who was sharing the footage?

Marco Pancini: Yes, we have done. We closed also a series of channels. I do not have now the numbers. We are very happy to come back to you. As I said, 20,000 pieces of content were taken down by the end of the day, and we closed also a few accounts of people who were spreading misinformation or even glorifying the attacks. We did, and I am happy to come back to you.

Q869 **Chair:** If you could write to us with that information.

Marco Pancini: Yes, of course.

Q870 **Chair:** Thank you. Do you know what the maximum number of views for any one film was that was shared of the Christchurch attacks?

Marco Pancini: You are talking about the original video or other pieces of—

Chair: Any of the films, so any of the films that were shared, that were taken down, what the maximum number of views for one of those films was.

Marco Pancini: No. I do not have this information. I can search for this information.

Q871 **Chair:** Also, would the YouTube algorithm have recommended that film to other users as well? People were using basic search terms to find the footage, but is it possible that some customers would have had a recommendation that prompted to them to watch?

Marco Pancini: The only possibility that some content related to the attack would be recommended is if this was an authoritative source like the BBC or other kinds of broadcasters. Of course, none of the videos that were identified by our machine for review would be eligible to be recommended. That could not happen.

Q872 **Chair:** Have you done an audit looking at the way in which the footage was shared, maybe to identify networks of people who take a particular interest in this sort of footage—the sorts of people who are interested in terror events and may be of interest to the authorities? Have you done a full audit of the platform?

Marco Pancini: We made an analysis of what happened on the day that I am very happy to share with you, including the information that you asked about.

Q873 **Chair:** Thank you very much. Just one final thing on the algorithm. One of the concerns about the algorithm is it seems to direct people towards odd content, and there are big concerns about the way medical conspiracy theories and anti-vaccine campaigns have been promoted through YouTube. Also, on the day of the Notre Dame fire, there were reported instances of YouTube recommending films relating to the 9/11 attacks in New York, which have nothing to do with the Notre Dame fire but would suggest that somehow it might be a terror attack or might be linked. Why would the algorithm be doing that?

Marco Pancini: Let me take a step back. Why are we working specifically on creating some additional information when users are looking for breaking news? We are doing this because we want to make sure that when a user is coming on the platform and is looking for something that is happening or news-related information, we provide to them authoritative sources. Also, if somebody is looking for events that could trigger some conspiracy theories or some questionable results, we want to make sure that still the users can see what authoritative sources, like Britannica Encyclopaedia, are providing in relation to what the user has searched for.

Going to what happened in relation to Notre Dame, the algorithm linked the picture of the fire in Notre Dame to September 11. It is a mistake. Machines make mistakes. Humans make mistakes. What is the solution? The solution is to make sure that humans and machines work together. In this case, what happened is that the combination that was triggered—again, the picture of the fire—was solved by humans looking at these

mistakes made by the machine and solving it. After one hour this bug was solved by our engineers.

Q874 **Chair:** Yes, I understand that. The concern or the suspicion is that the algorithm has a bias towards conspiracy theories, sensationalist news and, therefore, it makes these connections.

Marco Pancini: No, it is the other way around. Let us say a user was searching for September 11, the idea is for us to provide a panel, a box near the content saying, "Look, this is what happened", and you get alternative sources. The problem is that this box was triggered. In order to address your concerns, we wish to ensure that if some conspiracy theories are present in the search page—something that we are also trying very hard to avoid—there are also alternative sources that are displayed prominently.

Q875 **Chair:** I understand that. The reason I asked the question that way around was, if you are not just looking for information about 9/11 but you are looking at another news event, and then alongside that subtly YouTube is promoting information about another atrocity, or another thing that took place, there could be an inference or suggestion that somehow there could be a pattern between them.

Marco Pancini: No, no, that was just a mistake on this single case.

Q876 **Chair:** Yes. But the concern is that there seems to be a lot of mistakes where conspiracy-theory footage is promoted through the platform, and there is a slight concern as to whether that is because there is a bit of bias. Because some people might find conspiracy theories more sensational, they might drive bigger viewing figures and, therefore, the platform directs you towards some of that footage.

Rich Waterworth: I completely understand the question in terms of what that looks like. Just to be really, really clear, the watch-next algorithm is not designed to recommend more extreme or extreme content. Also, back to my earlier point, it has evolved and continues to evolve. The way that it is designed is to recommend a range of content. Some may look or feel more extreme, some will look less, some will be closely related to what is being watched and some will be more distant.

Q877 **Chair:** It is quite simple, it seems to me, that the next-up function is, "If you liked that, you will like this. If you take an interest in that, you will probably be interested in this". If you have an interest in extreme content then you would be recommended more extreme content.

Rich Waterworth: I was going on to also say, we have made changes to the algorithm to reduce what we call click-baiting content, and we have made changes to promote authoritative sources so that it will positively promote authoritative sources. We are very conscious of managing that dynamic, and we are absolutely designing that product and that feature so that it promotes authoritative sources, so that it promotes a range, and so that you do not drive people into a single source of truth. We are

very consciously managing, building and evolving that algorithm to ensure that precisely the scenario you are suggesting is not the case.

Marco Pancini: Just to give you an idea, the majority of the views that go to broadcasters on our platform, so coming from Up Next in the recommendation that we are providing, it is in our interests to make sure that quality content and authoritative sources are promoted.

Claire Lilley: Just one more thing from me on the recommendation system. One of the use cases is parents watching YouTube and then co-viewing with their children. If they are watching what we categorise as family-friendly content they will continue to just get family-friendly content. If the parent then later that evening watches a cooking video, or a DIY video, or whatever, while they are watching the family-friendly content it will never go to the DIY content or the cooking content. It will stay within that backing. That is one of our more recent developments to really cater for how people are using the platform in the real world.

Q878 **Jo Stevens:** I am going to go back to removal of content; I have just one question about that. When content is removed and you have identified it as being harmful, is that removed by a person or is it removed by an algorithm?

Marco Pancini: No, no, it is always removed by a person.

Jo Stevens: Always by a person? Always? Okay.

Marco Pancini: Always.

Q879 **Jo Stevens:** Thank you. I want to move on then to the recommendation system that we have talked about. Do you recommend videos with political and current affairs subjects?

Rich Waterworth: Yes. Those videos are included. We have seen over a number of years the consumption of current affairs content has grown on YouTube, alongside uses around music and all sorts of other things. We have seen a growth in that, and yes, that is included. That is one of the main reasons why we have taken some of the steps that we have to really closely manage Up Next and what is being recommended.

Q880 **Jo Stevens:** Yes. Do you profile users politically for those recommendations?

Rich Waterworth: No. The recommendations are based on what is being watched. If you are logged into YouTube it will take and include information on past content you might have watched. The recommendations are designed to recommend a range of videos, not, "You have watched 10 videos like this; we must recommend to you 10 more like that". The recommendations are designed to provide a range of options.

Q881 **Jo Stevens:** Okay. Because the platform has come under criticism for conspiracy clips relating to politicians. When we spoke to Juniper Downs

when we did our evidence hearing in the States at the beginning of 2018, she accepted that there was still work to do and progress to be made in terms of improving authoritative sources of news and moving away from that conspiracy-theory stuff.

Rich Waterworth: Yes, and that has been a big area of investment and focus over that period of time. Not to repeat myself too many times, but we have introduced specific changes to the algorithm to manage that as closely as we can. We think that is a really important area. One of the recent areas that I do not think we have mentioned yet is that conspiracy-type content is demoted by the algorithm now. That is a change that we have specifically made, so that if it is anti-vaccine content, or other sorts of content that have been flagged and are clearly of concern, those are demoted.

Q882 **Jo Stevens:** Okay. Another criticism that you have had is from a former employee, Guillaume Chaslot. He said the algorithm he worked on, "heavily promoted Brexit, because divisiveness was efficient for watch time, and watch time leads to ads." How does that correspond with what you have described as your broader societal aims and the values behind your company and your responsibilities?

Rich Waterworth: Yes. I think it very much speaks to the journey that we have been on and are on with both Up Next and YouTube. The absolute focus for us as a company from the leadership all down through the company, and the way we design products, is around our responsibility to our users, providing a high-quality and good and positive experience for users. The changes we have made exactly talk to this fact that we do not, in that algorithm, just prioritise watch time. We prioritise users having a good and valuable experience. That means that we take feedback from users, qualitative feedback where they will fill in a survey while they are using a site, telling us whether it has been a good experience for them, as well as other signals relating to their engagement with the content. Those are the things that determine whether we feel that tools and features are working or not, or are they just driving watch time.

Q883 **Jo Stevens:** Okay. Obviously, we have elections a week tomorrow; the European elections. You have talked about the steps you have taken around authoritative news. When did you introduce those steps?

Marco Pancini: We will go back to your question, but for us the European elections are very important, and we found the findings from your report on misinformation extremely interesting and in line with the work that we are doing together also with the UK, participation in the context of the code of practice on this information with the European Commission. Some of the steps that we are taking go in the direction of providing more transparency around political advertising. We are now making available every week information on how much money is spent for political advertising in Europe, with a breakdown country by country. If you go on the record you can see how much money was spent on our

platforms in the UK, and even the different advertising that was proposed.

Q884 **Jo Stevens:** When did you introduce that for the UK?

Marco Pancini: I think that was launched three weeks ago. It is for all of Europe, including the UK, the transparency record, for advertising and the authoritative sources. Going back to your question on when it was launched, I can come back to you, but is one month ago, or something like that.

Q885 **Jo Stevens:** One month ago. Okay. Have you done any assessment of whether or not in that month the feature has been successful? Obviously, we are in the regulated election period now. Have you done any assessment?

Marco Pancini: To go back on the process that we are going through with the code of practice, in the code of practice we have to provide every month to the European Commission and to all of the member states a report that is stating all the initiatives that we are taking and all the product that we are launching to fight misinformation and to make sure that the electoral period is going on—there you can see in a very transparent way every month. Yes.

Q886 **Jo Stevens:** Yes, I understand you are reporting that data to the European Electoral Commission.

Marco Pancini: No, no, to everyone. We document it to the Commission but this document is public so everyone can consult this document.

Q887 **Jo Stevens:** However, you are not doing anything internally, as far as you are aware, to assess whether or not the system that you have put in place has worked?

Marco Pancini: Let me go back and see what kind of findings we can already provide to you in relation to our findings.

Q888 **Jo Stevens:** Okay. You mentioned that you had rolled this out in the UK and in other countries across Europe. Did that happen to all the countries at the same time or were they done in a particular order?

Marco Pancini: No, no, the transparency register is for the 27 plus one, 28 countries, and was launched at the same moment, again, a few days before the end of the month in April.

Q889 **Chair:** Just to be clear on political advertising, do you disclose the amount of money being spent on advertising?

Marco Pancini: Correct.

Chair: Do you disclose who is spending it?

Marco Pancini: Yes. Yes, in two dimensions. You can see on the report both how much money and who invested for political advertising. It is also important for you to know that in order to do political advertising at

the moment you need to be registered. You need to provide to us evidence that you are doing this as a carrier of interest, or as a political party registered in the member state. Again, all this information is tracked and made public, but also every member who wants to know who is paying for a political advertisement can click the three dots on top of the political advertisement and see who is paying for it.

Q890 **Chair:** What would that say? If I had set up an account to serve these adverts it might tell you which account had done it, but would it tell you who the real person was or who was providing the money? Or is it just a self-declared piece of information from the person who served the adverts?

Marco Pancini: You can see in the report that we provided to the Commission—in order to be registered and to make political advertising on our platform you need to go through our registration process. You need to provide evidence that you have an interest in doing political advertising, including registration of your party or institution that wants to make a political advertisement, and including responsible people for the campaign, and the countries in which you want to make it. There is registration at the beginning of the process, and then the advertising is proposed and, of course, it has to be in line with our policies and transparent to the public and to the institution—

Q891 **Chair:** Is the information that is provided to you when a registration is made the same information available to the public?

Marco Pancini: The same information that is provided to us and then made public to the users, yes.

Chair: It is all of that?

Marco Pancini: Yes.

Q892 **Chair:** On profiling political messages, you have an advertising tool on YouTube that is rather similar or the same to the sort of tools that Facebook have. With Facebook it is called Lookalike Audience, and you have Customer Match. You can supply data to help target adverts and you can have a data set and then match to people who are similar on YouTube. Is that tool available for political advertising as well?

Marco Pancini: We need to look into the tool that the advertisers use in order to provide political advertising. However, let's just make it very clear, we do not provide those making advertising any information on the political preferences of our users.

Q893 **Chair:** Okay, but if I want to run a political ad on YouTube and I want to use Customer Match to help expand the audience for it, in the literature that you provide online it is quite clear that the idea of Customer Match is to help you to grow your audience. If you are using Customer Match to run political ads you are using data about people to guess whether they are interested in receiving that ad, just as Facebook does.

Marco Pancini: Our method is very different to Facebook. Unless you have an answer on this, I can come back to you with a screenshot.

Rich Waterworth: Yes. I can come back to you with details about whether political advertisers have exactly the same access to Customer Match or if there are any differences. I think we need to come back with details on that.

Chair: Okay. I think this is fairly basic stuff about the way the platform works. We are getting up to Facebook levels of writing back at the moment.

Q894 **Clive Efford:** Are you aware of what the Home Affairs Select Committee did with a fresh laptop with no content on it whatsoever? They went on YouTube and they searched, "News, politics and sport", and they were recommended a right-wing YouTuber. They clicked on that, and then were recommended a far-right Canadian video. They clicked on that and were recommended a UK channel from someone who uses racist, homophobic and violent content. Within three clicks they got to someone who made remarks about rape in relation to a woman member of this Parliament. That is a fresh piece of kit, no history on it whatsoever, immediately directed to far-right content. What is driving your platform when people are looking at political content?

Marco Pancini: We are looking into this specific example in this moment. We received the letter from the Home Affairs Committee yesterday, so we will come back to them with an analysis of the case that they have brought to our attention. We can go back to the description of how Up Next works, because I think it is addressing this.

Q895 **Clive Efford:** Okay. Just so you get the chance to have the full picture, the Chair of the Home Affairs Committee then put this content into her personal iPad, so these US and Canadian channels that the laptop was directed to. When she put them into her iPad, of the top videos that were recommended to her the first one and the fourth one were both of Stephen Yaxley-Lennon, a far-right extremist. Doesn't that cause you concern, that your platform is being driven unbeknownst to you? Is there any way that someone is hacking into the system and forcing it in this direction?

Marco Pancini: Again, we are looking into the case that has been brought to our attention by the Home Affairs Committee in these hours, because we received the letter last night. However, on the specific point of Stephen Yaxley-Lennon, we do not promote any video of his channel. His videos are not allowed to be shared, commented on, or promoted by our platforms. In two cases that are in the letter, as far as I can remember—because again, we received it last night—reference is made to one report speaking about this individual, and secondly about a video of a third party, the Canadian media, which is not directly managed by Stephen Yaxley-Lennon.

Q896 **Brendan O'Hara:** We have talked a bit today about almost rogue

YouTubers and how you take down content that breaches your rules. YouTube has also been pivotal in the creation of this influencer culture, and there are a handful of people who probably earn more money and have more recognition factor than some Hollywood stars. They are hugely important in driving particularly young people, I would imagine, to your platform. How do you deal with those influencers and those YouTubers who have this mass following, who are creating or propagating disinformation or spreading harmful content?

Rich Waterworth: I can talk a bit about influencers. We call them creators in general. They are a group of people that we work very closely with and we are enormously proud of. I talked earlier about why we are proud that YouTube is an open platform; it enables new and diverse voices to express themselves creatively and to find audiences. The other really important part to keep in mind when thinking about creators is that these are often people whose job is running a YouTube channel. They are creative entrepreneurs. We see, particularly in the UK, enormous growth, and we should be very, very proud of the strength of the creator economy in the UK. These are people from the length and breadth of the UK who are finding ways to express themselves that have not always been easy to do. They are finding audiences around the world. They are often building very, very strong communities around their channels. Those can be communities about learning how to do something, communities of shared interest, or communities of support around their personality or interests.

First and foremost, I think we are and should be very, very proud of those creative entrepreneurs, in the spirit of the UK having a proud history across the music industry and television industry. We also take seriously our responsibility to those creators as they build their businesses and manage themselves. We do a lot working closely with them to help educate them as to how to be successful creatively and commercially, but also to help them manage some of the challenges that come with being an individual, creative entrepreneur. Some of them find that very challenging. We work closely with them. We also take very seriously the fact that they are very influential, as you mentioned, often with younger audiences but often with older audiences. We take a strong position in terms of harnessing the positive influence that those creators can have.

We have invested over the last few years in a programme called Creators for Change. That is 68 creators and influencers from around the world. There are eight of them from the UK. They are using their channels and their influences to tackle really challenging topics, which could be bullying, could be xenophobia, or could be hate of different kinds. We have created a network to encourage and really lift up that group of creators to use their voices and use their influences to help tackle some of those issues. There is a really important role in the emergence of influencer culture and creator culture. There is a really important role in knowing how to harness that for positive benefit to users and society.

Q897 **Brendan O'Hara:** Do you as a platform feel you have a responsibility for protecting truth?

Rich Waterworth: As I say, we are an open platform, and people express themselves and communicate on the platform. We have community guidelines, and so if a piece of content is contrary to the community guidelines that will be taken down. We also have, as I mentioned, limited state, where something may not be against our guidelines but it may be highly controversial and damaging, in which case that video will not be recommended and all features will be taken away.

Q898 **Brendan O'Hara:** Would this be what is known as borderline? That is your borderline content?

Rich Waterworth: Yes.

Q899 **Brendan O'Hara:** This is all very well and good but 30 million people have watched a show that questions whether 9/11 happened. Thirty million people. There is a channel with 19 million subscribers that has just released a documentary saying that the earth is flat. These are your influencers. These are the big, big players on YouTube. Where is your responsibility to protect people and defend truth?

Marco Pancini: That is a very good example. For this specific reason, now when a user is searching for information on September 11 or information on other important historical facts that are sometimes subject to conspiracy theories, we provide very evident authoritative sources that put what happened in the right perspective; in this case, Encyclopaedia Britannica, or in other cases Wikipedia. It is very important for us to make sure that, of course, users have access to a diversity of views, and sometimes those are incorrect views, but always in the full respect of historical fact. Therefore, we provide a so-called knowledge panel, so some additional information that we proactively provide with context on these facts.

Q900 **Brendan O'Hara:** Surely even by your own evidence you are saying that you teach or you encourage and mentor these people on how to be successful. It would appear that these two channels alone have over 40 million subscribers. The lesson would be, the more outrageous you are, the more controversial you are, the more successful you will become, surely, which is not a good message to send.

Rich Waterworth: Our actual experience in looking at the most successful channels on YouTube is that they do not tend to be extreme. They tend to be more mainstream. They tend to be dominated by music, by comedy, and by vlogging. These are the kind of things that are the most popular content on YouTube. Education is a huge and really growing category. Our experience is that those are the types of content that really succeed on YouTube.

Claire Lilley: You mentioned in particular young people, so if I may, we feel like we have a responsibility to young people, both as creators of the

content, but also as viewers of the content. If children are in the content we will do a lot in order to try to protect them. Every time you upload a video to YouTube you will be presented with a question to say, "Is there a child in this?" and then be directed to information that will help you understand what your responsibilities are to any child in that video; complying with local employment law in terms of paying your children, in terms of working hours and all those sorts of things. We have teams of people who reach out to creators, particularly where there are children and where there is a family dynamic, because these are very popular channels, and we do a lot of work to reach out to them.

However, it is also true that in terms of the children who are viewing the content, a lot of children—43% of children in the UK—use YouTube to do their homework. That was a stat that came up. We really like to try to promote really good content. There are some creators like Physics Girl and Veratasium who are promoting really good STEM content, a lot of it particularly aimed sometimes at young women, which we really feel proud of as things that young people are engaging with and that are getting really good followings.

Q901 **Brendan O'Hara:** Rightly so. Those are things of which you certainly should be proud. Would I be right in saying that you define borderline content as misinformation and conspiracy theories, as not directly contravening your rules but a bit dodgy? Is that roughly how you would define borderline?

Marco Pancini: If I may I would like to challenge this assumption. We are investing in products, and you are going to hear more from us on this by the end of the year, in order to make sure that when the user comes on our platform they find content that is fulfilling them and is making them feel good, not content that is low quality, like, for example, conspiracy theories or content that is shocking. We do not have any interest, because it does not drive any positive experience, in borderline content.

Q902 **Brendan O'Hara:** The question I wanted to ask though was: how do you police borderline content? You cannot sit and watch every single video that is uploaded, can you, and then make a value judgment on the content of it?

Marco Pancini: Indeed. What we do, and maybe Rich can say more about it, is find ways when the user is searching for something, so in a general search, to provide authoritative sources, or content that is coming from qualitative sources. If a user wants for whatever reason to make a specific search for content that can be controversial, we cannot cross the line too much. We need to provide to the user the content that they are specifically looking for.

Rich Waterworth: To add to that specifically on the policing of borderline content, the model is very much the same as policing of content that goes beyond our community guidelines. That means that we

work with experts internally and externally to understand and create a shared and agreed definition of what borderline content should be. We want that to be a collaborative and collective view, not just our view. Then we enforce that view with a combination of humans and machines, using machines to pick up where there might be content that comes into one of these categories, and using human reviewers to confirm or deny and to make the machines more and more effective over time, which is what is happening there.

Q903 **Brendan O'Hara:** Moving on just a bit, Child Wise recently in its under-age social-platform use report said that YouTube was considered the worst offender for advertising, with the majority of children mentioning the video platform and its reputation organically without prompting. Why would Child Wise come to the conclusion that it has?

Claire Lilley: From research that I did when I was at NSPCC, children do not like advertising, and that is fair enough. None of us do, but it is the way in which the platform is funded. That is kind of a reality. One thing I would say is that we do have very restrictive advertising policies with regard to children on our YouTube Kids platform, which is specifically designed for nought to 12s. We have on there a much-restricted range of products that can be advertised, so no political or religious advertising, no beauty advertising, no food and drink advertising. The advertising that is on there must not have any incitement to purchase, for example, so the way in which the ads are presented is very conservative.

On that platform children will only be presented with an ad once every 15 minutes. Also, in recognition of the idea that children's digital literacy is not necessarily as developed as an adult's might be, what we do is serve them what we call a three-second bumper. It is basically a warning that what they are about to see is an ad, so they can understand when what they are watching is entertainment content and when what they are watching is commercial content. We really try on the YouTube Kids platform to reinforce positive, very conservative, very responsible advertising practices.

Q904 **Brendan O'Hara:** Why, in your opinion, would Child Wise have come to the conclusion that it did?

Claire Lilley: That children do not like advertising?

Q905 **Brendan O'Hara:** No, that YouTube is the worst offender for advertising to young people?

Claire Lilley: I do not think we treat young people aged 13 to 18, or aged 13 to 24, any differently than we treat other users in terms of advertising.

Marco Pancini: Take into consideration that, for example, we want to provide the user with full choice. If a user wants to have an experience of YouTube with no advertising, there is always the possibility to subscribe

to YouTube Premium, which is the YouTube platform with no advertising, plus more content.

Claire Lilley: I just want to correct my previous statement; we do treat children differently, because we do not serve interest-based ads to under 16s. It might be their frustration is that what they are seeing is ads that they do not feel are relevant to them because we do not profile them in the same way we do with adults.

Q906 **Brendan O'Hara:** Just finally, what is YouTube's policy about informing viewers if a video is sponsored or not?

Rich Waterworth: The ASA has very clear and good guidelines and we invest, as part of our resources and support for creators, in making sure that we reinforce and make sure that they are aware of that. That includes at the moment upload. If you are uploading a YouTube video there are reminders that if there is product being promoted it needs to be very transparent and clear. We think that that is a very good and important guideline, and we take steps to make sure that creators are aware, and if they do not adhere to that then they will be dealt with.

Q907 **Brendan O'Hara:** Every video that has a sponsor on YouTube is transparent and it can be seen as such?

Marco Pancini: Correct.

Q908 **Chair:** On the subject of children, I want to ask about some research that was conducted earlier this year by the Cyprus University of Technology, which analysed 130,000 videos targeting toddlers between the age of one and five.

Claire Lilley: What was the age range again?

Q909 **Chair:** It was children between the ages of one and five. It was based on people watching content aimed at children between the ages of one and five, so children's cartoons, children's programmes and that sort of thing. They found, based on this sample set of 130,000 videos, there was a 45% chance of a child coming across inappropriate footage within 10 clicks of a child-friendly video. They found that children stand a one in 20 chance of being recommended problematic videos after watching any video. I wonder if you recognise that research and those figures. Certainly, the idea that there is a 45% chance of coming across inappropriate footage within 10 clicks of watching a child-friendly video would seem to be concerning.

Claire Lilley: It certainly would be concerning, yes. I completely agree. Just for me to clarify, is that on YouTube?

Q910 **Chair:** Yes, on YouTube.

Claire Lilley: Okay. What I would say in response to that is this is exactly the reason why we created YouTube Kids, which is a platform designed for nought to 12 year-olds. We recognise that YouTube main is

for 13-plus, and there is a lot of content on there that is just not suitable for younger children. There is co-viewing on there, as I said, with parents, but in the main this is why we put in a lot of investment. YouTube Kids has a much-restricted corpus of content and a much higher threshold to be on there. There are lots of systems in place. Parents can choose, for example, to turn search off, they can choose to have only parent-approved mode, where they literally can handpick every video and every channel that their child might want to see. The videos that are served on the home screen have been verified by humans as 100% safe. It is an app designed with built-in features to protect those children at that very, very vulnerable life stage, because we absolutely do not want any child to see content of that nature.

Q911 **Chair:** I understand YouTube Kids, and this research, I believe, was on the YouTube main. However, again, it is just concerning that the way the algorithm is driving viewer behaviour is someone is watching a series of films that are designed to appeal to very young children, and then within a relatively short period of time they end up being directed towards more problematic content. That is a concern.

Rich Waterworth: One of the Up Next algorithm changes that Claire mentioned earlier—and this may be something that has been introduced since the time that you are quoting, I do not know—is that if you start a session and you are watching family-friendly content, even on YouTube, it will only now recommend family-friendly content. We are very conscious and aware, and this is why we have tried to emphasise that there are changes that have been made and are being made to make sure that that algorithm is delivering a safe and a high-quality experience for all of our users.

Q912 **Chair:** When was that change made?

Rich Waterworth: I would have to come back to you on exactly when it was made, but in the last six months, I think. It is quite recent.

Chair: Okay. If you could I would be grateful, because that research is dated for this year. I would be interested to see whether it was done before. We can check with the institution as well. Thank you.

Q913 **Rebecca Pow:** Apologies for being late, but I just have a couple of really quick questions. Just on that note, and listening to everything that is being said, I get the impression that you are just reacting as this whole thing grows like topsy with no concept of the direction it is all going in. It sounds like you have some good measures in place but do you not feel like you are firefighting every day to keep up with this thing that you have created, that it is almost impossible to control what people are seeing?

Rich Waterworth: To step back for a minute, we have just gone past the 50/50 moment, where more than 50% of the world is online. In the UK in many markets it is well beyond that, so very, very high levels of penetration. What happens when everyone is online is you get bad actors

as part of that. Digital consumption and penetration is an evolving story and continues to be. Our responsibility is and will continue to be to get ahead of what is happening. What is happening is changing over time, and partly it is changing over time because of our measures through the intelligence desk, through more robust enforcement procedures, and through much, much better algorithms that catch content at scale. As those measures get more effective, bad actors—who, let us face it, are around in any situation where all of humanity is there—have to try new things.

Q914 **Rebecca Pow:** I understand that, but it does sound like a lot of the time you are playing catch up rather than being ahead. Reacting to the Christchurch incident, for example, is catch up rather than being ahead.

Claire Lilley: In relation to kids, I do think we are ahead. I think we are industry-leading in the creation of tools like Family Link and tools like YouTube Kids. We saw that, and that has been a product of a decision five years ago to build those products and to recognise that future need. It is true that we are continually iterating the products, and that is right and fair, because what we see is people changing the way they are using the products all the time as new technology comes online. Particularly for me, the way offenders are manipulating and using systems is changing all the time, and we are always horizon scanning; the intelligence desk is always trying to get ahead of those. One of the key areas for me that is really important in getting ahead of this is about the education of young people and their parents and about how they can use technology in a safe, responsible and enjoyable way, and get the most out of it. That is why we are investing heavily in education programmes, both in primary and in secondary schools.

Q915 **Rebecca Pow:** At the moment I am glad I do not have small children, because it sounds like a minefield of trying to control what your children can and cannot see with all the different methods, switches, names and all of that. To me it sounds complicated for the parent.

Rich Waterworth: I am a parent and it is challenging. This is why, as Claire said, we are investing in education. Our Be Internet Legends programme has been experienced in 50% of UK primary schools and growing, and we are going to get to 100%. By the end of this year that programme will have been experienced by 1.9 million young people in the UK. These things are challenging and we need to continue to invest across a range of measures, education, intelligence, enforcement, that is what we are doing and will continue to do.

Q916 **Rebecca Pow:** One final question, just picking up from Mr O'Hara on the influencers, I see you have your set advertising structures and how you control adverts, but surely the influencers themselves are advertisers, aren't they? They are really endorsing products, they are valuable to advertisers. How much control do you have over that?

Rich Waterworth: I think this relates to what we were talking about earlier. There is an ASA code that states that if an influencer, if a creator, is promoting a product they have an obligation to be fully transparent about that. That is a code that we fully support and we educate our creators around it, including at the point of upload, and there are penalties if they do not adhere to that.

Marco Pancini: There are also laws in place. There is a recently approved law called AVMS, an audiovisual media services directive, that will be applicable also in the UK. There will be clear rules for promotion. Going to your question, if I may for one second, our CEO described this approach as responsible growth. It is really encompassing all the different product development and it needs a clear vision of making sure that every new product, every new policy that we are implementing, is responding to the need of the platform to grow, but in a responsible way. Indeed, we have a plan. One example I would like to make, and maybe Claire can say more, is Family Link. Because we know how parents are struggling across all the different platforms, we, as Google, wanted to give them one place, it is called Family Link, where they can go, they can log in and make sure that all of the different accounts of their kids are linked there. They can control from one simple page the behaviour of the kids and how the kids are interacting with technology.

Q917 **Ian C. Lucas:** I will come back to our conversation about liability earlier on. Just one short question: in the Online Harms White Paper it says, "The Government proposes the establishment of a new statutory duty of care to make companies take more responsibility for the safety of their users and tackle harm caused by content or activity on their services". Do you as YouTube accept that as a legal basis for operation or will you be opposing it?

Marco Pancini: As I said, we are absolutely engaged in the discussion, with our chief legal officer meeting with the Government to discuss that and to better define what the duty of care means. Because for us what it means is to act as a responsible platform, to have this approach of responsible growth, and to make sure that when we design a product this product is trying to maximise the experience of the user and keep the user safe.

Q918 **Ian C. Lucas:** I am sorry to interrupt you. I am just going to ask you a straight question here, and I would like a straight answer. Do you agree that you should take responsibility for harm caused by content or activity on your platform?

Marco Pancini: I can appreciate that is a very thin line but the responsibility, yes, the liability, no.

Q919 **Ian C. Lucas:** I do not understand this difference. I am a lawyer. Liability is legal responsibility for something. The only difference I see between the two words is moral responsibility but not having legal liability. What I want you to confirm is that you will have legal liability for the content of

your platform.

Marco Pancini: As a lawyer, I am also a lawyer, it has to be well defined.

Q920 **Ian C. Lucas:** That is why I want a straight answer to the question.

Marco Pancini: That is why I cannot give you a straight answer, because we are still in the process of defining what exactly the duty of care means in terms of knowing in advance—

Q921 **Ian C. Lucas:** That is why I have quoted you the Government's words. I would like you to tell me whether you will be opposing that proposal by the Government or will you be supporting it?

Marco Pancini: We will not oppose the proposal. We will be part of the process in discussing this proposal. However, I have also seen the hearing from your Secretary of State when he came here to discuss that, and indeed there are some areas that still need to be defined, such as the role of the regulators and the code of conduct. It is very important to define exactly what this duty of care will mean for us so we know when we have crossed the line and we become liable. In very practical terms, when we have this conversation with our colleagues from the US, they ask us, "Okay, Marco, in concrete terms what does it mean? In what situation will we, in a sense, break these rules? When will we not have fulfilled our duty of care?" Indeed, we are struggling to provide them with a clear example, because still we lack the code of conduct. We know, of course, the law. If we break the law, or if, as in the example made by Jeremy Wright, a website that you say is illegal should be IP blocked. Giving access to illegal content online would indeed raise a liability or responsibility and would be in breach of the duty of care. I think that is the right direction in designing how this will play out.

Q922 **Chair:** I think that is a good summary of why it is important that Parliament decides, rather than leaving it to companies. Thank you all very much for your evidence this afternoon. I appreciate it, thank you. We will now invite the representatives from Instagram to come and take their places on the witness table. Thank you.

Examination of witnesses

Karina Newton, Head of Public Policy, Instagram; Vishal Shah, Head of Product, Instagram.

Q923 **Chair:** Good afternoon, and thank you for joining us this afternoon for our evidence session on our inquiry into addictive and immersive technologies. Before we start the questions themselves, I just want to state for the record that obviously the Committee is bound by the House's rules on sub judice. For the avoidance of doubt for everyone in the room, the sub judice resolution applies to proceedings in the

Committee and, therefore, witnesses and members should not make reference to any individual case that is currently before the UK courts or cases that have been listed for a date to appear before the UK courts either. I just wanted to make that clear.

Vishal Shah, you are Head of Product at Instagram. You will probably, I am sure, be aware of the tragic news that has been reported today of the death of teenage girl in Malaysia. The incident seems to be that she posted on Instagram asking people whether she should live or die, and 69% of responders voted that she should die. She then took her own life. What can Instagram do to make sure there are no further tragic cases like this?

Vishal Shah: Thank you, Chair. I want to start by saying the news is certainly very shocking and deeply saddening, and our thoughts and prayers go out to the family of the young woman in Malaysia. The tools that we build, in this case the polling sticker in stories, are designed for helping people express themselves and to connect with the friends, family and followers that they have. By and large we believe that is how the platform is mostly used. There are these cases, like the one we have talked about today, where our responsibility for keeping our community safe and supported is tested, and we are constantly looking at our policies. I will pass over to Karina in a second to talk about our policies.

We are deeply looking at whether the products on balance are matching the expectations that we created them with. In cases like the polling sticker, we are finding more evidence of cases where it is not matching the expectations or the creative guidelines under which we created them. We are looking at whether we need to make some of those policy changes like we have done with Live more recently. Maybe Karina can talk a little bit more about how we think of suicide specifically from a policy perspective, and some of the product changes we have made as a result.

Karina Newton: Yes, it was absolutely heart breaking, and I have not talked to anyone at Instagram that did not feel the same. The way we have looked at our policies around suicide and self-harm has been formed by expert discussion. We work closely with academics, charities, and people who have lived experience to inform where we draw the line and what we allow and do not allow on our platform. We work with the Samaritans here in the UK. We do not allow the promotion or endorsement of suicide or self-injury content. That polling sticker as a story would violate our guidelines. It promotes the act of suicide.

We have allowed—and we think that there is an important role to play especially for Instagram to be a place that can help destigmatise mental health—that you can come on to Instagram under our policies and admit that you are struggling and that you need help. This offers people who are in your life, your friends and family, the opportunity to say, "Can I help you?" We also provide information; so if you post content and it is removed for suicide or self-injury content, we will provide the person who

posted it with resources on how to get help and connect them with local organisations, such as the Samaritans.

We will also, when we review that content, if we are worried about immediate harm work with local law enforcement to get someone there. We have heard anecdotally that we have saved lives through this programme. We are looking very much holistically at our policies to ensure that they are in the right place. The case in this tragedy is no different. We will be looking at it and making sure that we are doing everything we can to ensure the safety of the people who come to Instagram.

Vishal Shah: If I could add one thing to answer your question specifically, one of the things that we see in cases like this where there is some amount of media attention is that then results in some mimicry on the platform, where other people try to copy this behaviour. We are always looking at trends and new ways in which people are expressing these kinds of mental-health issues or suicidal tendencies, and it is a constantly evolving space. In this case we are paying specific attention to the usage of the polling sticker and cases where there might be other troubled people who are using the same methodology.

Q924 **Chair:** I have seen—of course this was yesterday—the poll was live for about five hours or so. What more can your systems do to identify people using the polling sticker for something that could lead to harm, as this did, sooner? I assume that no one reported it to Instagram while it was live. If they did report it and no action was taken, I think that would be a very serious matter, I am sure you would agree. I would be grateful if, if possible, after the event you could write to confirm that was the case; there is no record of this poll being reported to Instagram. What can the company do to identify people who are using these tools in ways that could cause harm?

Vishal Shah: For this case specifically it is an active area of investigation. Our team on the ground in Malaysia is engaged with the local authorities to make sure we can understand more information about it. Generally, this is a bit of what I was referring to earlier, where sometimes the ways in which these tools are used and how people are expressing these thoughts can evolve. Without talking about the specifics in this case, because it is still active, there is speech that maybe we would recognise as potentially coming from people who are suffering from mental illness and expressing these self-harm type thoughts. It is my understanding that sometimes it is the use of emojis, or the use of things that are not part of standard parlance, around why and how people are expressing these things, especially with young people.

Sometimes it is about making sure that when cases like this occur, and we learn of a new way in which people are expressing these thoughts, we can add those to the proactive detection of these things and put up warnings offering the producer of that content, the person expressing those thoughts, a way to get help. Also, for the people who have come

across this content, especially again with young people, to make sure they have the opportunity to get help if they come across content that they are discovering on Instagram.

Karina Newton: I would like to add just one thing. We as an industry are really looking at how we can be very proactive in this space. You may be aware that Secretary of State Matt Hancock recently brought together several members of tech companies, including us, and we are going to be working with the Samaritans to invest in research in this area and use this research to inform how we think about things, from policies, to enforcement, to support resources. I do think it is a really positive development to work with the Government and work with charities such as the Samaritans to better understand how we can invest support and prevent instances like this.

Q925 **Chair:** Given that Instagram holds a huge amount of data on its users, is there not more that you could do to identify people who are at risk, maybe because of the content they are engaging with and the profile that is being created by the data based on their engagement on the platform? Can more be done to identify these people and to monitor their activity, or even to reach out to them?

Vishal Shah: This is a constantly evolving space. From a technology perspective, we are trying to get better at predicting and practically removing or identifying this content. Sometimes it is about past behaviour; sometimes it is about the act that they are currently in the process of sharing or expressing. At the same time, we want to be conscious of ensuring that the content either violates our policies, which Karina talked about, or is a case where we want to be offering help more proactively. There is a spectrum there. In general, I agree with you that our responsibility to our community of keeping people safe and supportive means we should be continuing to invest, as we have been, to help keep people safe.

Q926 **Chair:** Instagram uses the same advertising tools as Facebook. That is correct, is it not?

Vishal Shah: We use the same underlying advertising platform as Facebook. That is correct.

Q927 **Chair:** Is it not possible to use some of the tools that are used to profile people to sell them stuff to identify people that may be suffering and, therefore, may be at risk?

Vishal Shah: To be clear on the advertising side, there are many categories of targeting that are not possible through our advertising tools. Things like mental health and depression are in that category of things that are not targetable. Maybe not to jump into jargon, but it is important to think about how we can use all of the signals that we have on the platform, all the ways that people interact with content, and the accounts that they follow, to understand if there is additional risk, and from a product perspective, see if we can get ahead of some of these

issues. I do think it is important to make sure that we have the right balance of very clearly defining when something violates our policies and making sure we act on that explicitly, versus cases where we can do a better job of detecting something in advance and working with law enforcement if we believe there is imminent harm for a young person using our platform.

Q928 **Chair:** Do the same sort of Lookalike Audience tools exist on Instagram that exist on Facebook?

Vishal Shah: We support that, because we sit on top of the same advertising platform, yes.

Q929 **Chair:** Yes. So, would it not be possible, where there are cases of people that are known to have been engaged in harmful content and may be at risk, that analysis can be done to see what other users share similar characteristics based on an interest in the same sorts of subjects?

Karina Newton: These are some of the things that we are trying to evaluate and understand. Obviously, we have some privacy concerns of users, and also some GDPR obligations as well. But one of the things that we are trying to see, if you search right now in Instagram we have a couple of different surfaces; there is a home tab, which is where most people spend most of their time. That is where you see what you have followed. Then we have a search and explore tab, which is where you find content. If you search for a term such as "suicide", we have a pretty good indication that you are interested in that subject, and so we use this as an opportunity to help. As soon as you search for that term you will receive a pop-up message that says, "Can we help? Based on what you are searching for, this is some information".

That is localised in 40 different countries with resources there. We are using the signals that we have from what people are searching for, and we are looking for additional opportunities to do that, perhaps based upon accounts you follow, or perhaps based upon other search terms. We really do want to use the time that people are on Instagram and our platform as a way to provide resources in as many ways as possible.

Q930 **Chair:** Yes. I have seen those prompts that you describe. You mentioned that there are GDPR privacy issues on using data to try to identify people that might be at risk. I suppose the concern we would have is that those privacy and GDPR issues do not exist in terms of using these advertising tools, but inferred data is used all the time in advertising tools to sell people stuff. Instead, could that data be used to identify people who could be at risk of self-harm?

Vishal Shah: Just to be clear, you cannot target people based on their potential likelihood to have mental illness or to have depression. From a technology perspective, there are perhaps more granular signals that can be used, as opposed to ones that are about specific profiles or interests that you might have, can potentially help identify these things in a compliant way from a data perspective. Right now we have that in the

search capability where it is a very clear intent. I think the steps are in the right direction, which is offering people help when we think they need it. Now the question is: where else can we highlight that and show it across the product?

Q931 **Chair:** Let's say I was the Samaritans and I wanted to run a campaign on Instagram, and I said, "I have a load of data on people that are already in our database that have given us their data, can you expand that audience through lookalike audiences to find more people who are like that?" Is that something you would offer?

Vishal Shah: Lookalike audiences are available. You have to give proof that you have permission to use the custom audience in the first place. The way the algorithm would find potentially other people that are interested in the same topic would not necessarily be based on the specific interests that they might have. It might be other accounts that they might follow. Let's say this happens to be young people; it might index on young people and things that they care about.

Q932 **Chair:** I understand that, but if you are a charity organisation that works with vulnerable people and says, "We have a dataset of vulnerable people", the purpose of using a tool like lookalike audiences would be to find other people that share similar characteristics. The tools that are available would do that. You have confirmed that would be possible to do. I just wonder whether that is something the platform could look at doing, working with charities, to try to identify proactively people who may be at risk and maybe someone should reach out to them.

Vishal Shah: I do think that the engagement we are in with the Samaritans and other charitable organisations helps us in a couple of ways. One is it allows us to leverage expertise on these topics that we may not have inside the company, and the second is it allows us to engage in research. I know we have been working with some of these organisations on grants to do some of this research to then incorporate best practice into the product.

In general, I think this is an important area for us to be investing in, given that young people do use our platform. It is Mental Health Awareness Week here in the UK, and it is something we certainly deeply care about in terms of destigmatising and then being able to play our role in helping combat it.

Q933 **Chair:** Could I just ask a final question on this? Will Instagram be taking action against the accounts that voted in this poll, particularly those that encouraged her to take her life?

Karina Newton: We are currently investigating that right now.

Q934 **Chair:** Would you potentially suspend those accounts or cancel those accounts?

Karina Newton: Literally, as I am sitting here, a team is investigating. We are happy to circle back after that investigation is over.

Q935 **Chair:** You say you are investigating, but it is a matter of fact there was a poll. People voted in it. You have records as to who voted. Is this a point of principle? Will the company be taking action against those accounts?

Karina Newton: I hope you can understand. It is just so soon. Our team is looking into it to see what the content violations are. Again, that poll in and of itself goes against our policies for promotion of suicide.

Q936 **Ian C. Lucas:** Can I just ask some questions about the governance setup within Instagram and the relationship with Facebook? Do you have a separate board that runs Instagram? Are you a separate business to Facebook? How does that work?

Vishal Shah: We are a product group inside Facebook, to use our internal terminology, which means that product decisions that we make around Instagram specifically, how content is displayed, how our algorithms work, how people express themselves and the tools we create are all governed by Instagram. We were part of the same company, and therefore the Facebook Board of Directors is also the Instagram Board of Directors. From a product perspective, we have considerable autonomy in terms of the decisions that we make.

Q937 **Ian C. Lucas:** The buck stops with Mark Zuckerberg with Instagram?

Vishal Shah: Instagram is run by Adam Mosseri. He is the head of Instagram and the majority of decisions will land with Adam. We are a company and Mark Zuckerberg is CEO, and there are some decisions that go to the CEO.

Q938 **Ian C. Lucas:** In terms of the data arrangements between Instagram and Facebook, are Instagram data shared with Facebook?

Vishal Shah: Because Instagram and Facebook are part of the same company, it is clear as part of the terms when you sign up for this service that the data policies are shared between the two platforms.

Q939 **Ian C. Lucas:** Does that include WhatsApp as well?

Vishal Shah: WhatsApp, to the best of my understanding—and Karina, please correct me, I have more familiarity with Instagram—is on a separate data policy, but Instagram and Facebook are together.

Q940 **Ian C. Lucas:** Are the data shared with WhatsApp?

Karina Newton: We have separate data policies that govern them. Facebook and Instagram share a data policy and WhatsApp has separate terms and data policy. The data are not shared between the two companies.

Q941 **Ian C. Lucas:** They are shared between Instagram and Facebook but

they are not shared with WhatsApp. Is that right?

Karina Newton: That is correct. There are some exceptional circumstances or some other circumstances. Let's say we find an instance of terrorism content. We create a digital fingerprint of that content, we store that and we share that information through organisations with other tech companies as well. WhatsApp could use that digital fingerprint as well. There are a number of ways in which we have industry partnerships across technology where information is shared. In terms of the data policy that Facebook and Instagram share, we operate as one company.

Q942 **Ian C. Lucas:** How old is the average user of Instagram?

Vishal Shah: Unfortunately I do not have that data to share with you today. That is something we can follow up with if possible, but I do not have that.

Q943 **Ian C. Lucas:** I cannot believe you do not know that. That is a pretty important thing for you, isn't it? Is it 65? Is it lower? Is it mainly people under 18?

Vishal Shah: In the UK and in accordance with GDPR, when a user signs up for Instagram, if they indicate that they are less than the age of 18, then we ask them for their exact date of birth to verify that they are over 13. If they are over the age of 18, then they do not provide any additional information in terms of their age when they sign up for Instagram.

Q944 **Ian C. Lucas:** If they are over 13, what do you they have to do then?

Vishal Shah: If they say that they are under 18, then they are required to put in their date of birth. If they are over the age of 18, they are only required to state that they are over 18.

Q945 **Ian C. Lucas:** If you are between 13 and 18, do you then sign up? How does it work?

Vishal Shah: Yes. Then you can input your date of birth and you can sign up for Instagram.

Q946 **Ian C. Lucas:** A 13 year-old signs up with Instagram. Does that person give his or her consent to sharing data with Facebook?

Vishal Shah: When someone signs up for Instagram, the terms that they sign up for indicate that Instagram and Facebook are part of the same company and share a data policy.

Q947 **Ian C. Lucas:** How does a 13 year-old give legal consent to share data with Facebook?

Vishal Shah: Just to be very clear, they are not creating a Facebook account. They are creating an Instagram account. The company—

Q948 **Ian C. Lucas:** How do they legally contract with Instagram as a minor, as someone under the age of 18?

Karina Newton: In our registration process we have the terms, which express how we do and do not use information, and the data policy clearly states what we do and do not share. We also have very plain language in those terms, and we have additional information in our help centre to help provide people with information. I believe the age of consent here in the UK is set at 13.

Q949 **Ian C. Lucas:** I am a bit rusty on this sort of stuff but I am not sure about that. What concerns me here is that we have been looking into this subject for approaching two years now. It is very complex. Issues like the sharing of data are incredibly complex, which is why I have been asking you some dull questions about company structures and sharing of data. I do not think that someone who is 13, 14 or 15—or in fact a lot of adults—really gives informed consent on these issues and understands the importance of these issues, particularly when they are under 18. Don't you think that is a valid objection?

Karina Newton: We really do try our best to make it as clear as possible what we do, what we use data for, and try to inform young people as well as our entire community what we use data for. We want to be transparent with our users so they understand, when they are using Instagram, what that means for us and them.

Q950 **Ian C. Lucas:** I am still slightly unclear about WhatsApp. You are saying that information is not passed to WhatsApp from Instagram?

Karina Newton: We are here today as Instagram experts. We are not experts in WhatsApp. If you have further questions on this, if we can come back to you on those—

Q951 **Ian C. Lucas:** Yes, but I am asking you about Instagram information and whether it is shared with WhatsApp, so I would have thought that people from Instagram would know that.

Karina Newton: We work on the products that we build and on the policies that affect what content is and is not allowed and how we work on the safety and security of our community, but we would really be happy to come back on additional questions. We just do not want to unintentionally mislead you because we are not experts in that space.

Q952 **Ian C. Lucas:** I understand that, but do you work on what data you share, with data-profiling of customers that you have? Is that part of your responsibility?

Karina Newton: I do not work on the ads side but I do work on ensuring that the privacy of our users is respected when we build products. I am happy to speak to any number of privacy matters and how we think about privacy on Instagram.

Q953 **Ian C. Lucas:** An important privacy matter for me is whether, when I give information to a platform, it is shared with another platform. Do you think that is a reasonable point of view? Do you agree with that?

Karina Newton: Absolutely.

Q954 **Ian C. Lucas:** What I am asking you is whether information that users of Instagram give is shared with WhatsApp. That is to do with privacy. Does that happen?

Karina Newton: I, again, just do not want to be—

Ian C. Lucas: You do not know.

Karina Newton: Thank you.

Ian C. Lucas: If you could come back to me on that.

Karina Newton: Absolutely.

Q955 **Chair:** I would just like to ask one question. Vishal, you said that from a product point of view, Instagram and Facebook are separate. You may know that Christopher Wylie gave evidence to the Committee last year. He was a Facebook whistleblower in the Cambridge Analytica affair. As a consequence of being a whistleblower, he had his Facebook accounts cancelled. He was also blocked from Instagram as well. Given that he did not say anything about Instagram, I wondered why that would have been the case.

Vishal Shah: We do share community guidelines between the two platforms for both content and accounts. If an account was to violate community guidelines, they would be removed from the platforms that they violated.

Q956 **Chair:** I think in this case he did not violate Instagram community guidelines because he did not say anything about Instagram. His whistleblowing only related to Facebook.

Karina Newton: I am not familiar with that case but we would be more than happy to look into it and circle back.

Q957 **Clive Efford:** Just going back to your answer about the age profile of people who use Instagram, how much is Instagram worth? What is its revenue? What is its turnover?

Vishal Shah: We report revenue at Facebook as a company by region, and we do not break it down by individual products. It is broken down by the regions in which we operate as a business.

Q958 **Clive Efford:** What would be our nearest region and what would it be worth?

Vishal Shah: I do not have off the top of my head the most recent European or EMEA revenue, but I am happy to get that number. It would be joint with Facebook numbers.

Q959 **Clive Efford:** It is fair to say it is quite a sizeable company and you do not know who your primary customers are? You do not have a profile of them?

Vishal Shah: Just to make sure I am answering your question accurately, when a user signs up for Instagram, they do not give their specific date of birth unless they are under the age of 18.

Q960 **Clive Efford:** If you are directing content at your customers, surely you have some idea of who they are and what they might be interested in, and you are saying that you have no idea of the age profile of people who use Instagram?

Vishal Shah: From an advertising perspective, from a business perspective, because we use the same underlying ads system as Facebook, if a Facebook account has provided their age, that can be used as part of our advertising platform. Your specific question was around the age of people on Instagram, which we do not necessarily have full coverage on because we do not ask for age on Instagram directly.

Q961 **Clive Efford:** You do not assess the age profile of the people that Instagram is targeted at?

Karina Newton: We do have some information. If you have, for instance, added your date of birth to your Facebook profile and connected that to your Instagram account, we would know how old you are. We also are able to tell within some range what an age might be based upon similar accounts that you follow. There are some large buckets of areas of age groups that we have some familiarity with.

Q962 **Clive Efford:** Are you aware of the study by the Royal Society for Public Health? It says that Instagram was ranked worst for its effects on young people's mental health. In particular, it scored lowest for its effect on body image and was also linked with young people fearing that they were missing out on situations, FOMO, and sleep deprivation. Do you recognise that?

Karina Newton: We have seen that study, yes.

Q963 **Clive Efford:** What is your reaction to it?

Karina Newton: These are really tough issues that Governments and academics and technology companies like ours, as well as society, are thinking about. It is a really novel time. Social media is rather new, so we are all trying and invested in figuring out what our impact is. At Instagram we are deeply invested in researching that in terms of our wellbeing research—Vishal can talk a little bit about our wellbeing team at Instagram, as well as working with outside academics to better understand this. We have a grant programme that we are investing in to help charities and others research what our impact is. We are deeply invested in that.

I think there are a number of areas that we know we really want to invest in based on some of the initial information coming in. For instance, young people and bullying is a serious issue. Bullying is obviously something

that happens on and off Instagram, but as a platform for young people we think that we have an opportunity to really do something about it.

Vishal Shah: We have a dedicated team at Instagram named Wellbeing, and it is responsible for a variety of issues on the platform. Some of them are around content and making sure that we remove content that violates our policies, but a large amount of it is thinking about wellbeing of people broadly on the platform. It is the team that is responsible for our comment filters and our ability to manage time, and some of the issues that you are mentioning and that Karina mentioned around bullying. We are trying to take a leading stance on some of those issues where we believe we can, and where the opportunity of having young people on the platform and the ability to shift behaviour does exist and could be there for good.

I want to also clarify that there are conflicting studies on some of these issues that you mentioned. I think it is helpful to have a variety of engagements with policymakers and also academic institutions—like a report that came out from Oxford just last week—to make sure that we have a balance and we are looking at the apples-to-apples comparison.

Q964 **Clive Efford:** Yes, but if I were a drug company and I were testing a drug that had not been proved to be safe on young people, I would be prevented from doing that. Why do we accept this from online companies, from social media companies? Should we have some degree of safety before we go out there and feel our way because it is all new?

Vishal Shah: The safety and responsibility of our users is incredibly important. To be clear, I think the vast majority of usage of Instagram is beneficial to people. It allows them to connect with their friends and their family. It allows young people to connect with their colleagues and their schoolmates where they are not next to one another but they can still connect with one another. That being said, these are novel issues and there are content challenges that we are dealing with and we are facing. We are here today to make sure that you all on the Committee hear from us that we are committed to investing in that from both a product and policy perspective.

Q965 **Clive Efford:** If I can just go back on this issue about age, how does Instagram ensure that its users are not below the age limit set by the platform?

Vishal Shah: That is an important question and one that we are, as an industry, really trying to understand and figure out how best to implement. Certainly there are things we have in place today that are around asking for age upfront, but I will admit that a young person can get around that. We have abilities to report content after someone has joined Instagram. We will review those accounts. If the account is under the age of consent, they will be deactivated.

Broadly speaking, we think there is a role for technology to play to help us detect these things in a privacy-safe way. That is the opportunity that we are excited to pursue. I know there have been other proposals around age verification that have privacy challenges that I am happy to chat through, but it is an important industry-wide problem.

Q966 **Clive Efford:** What is your comment on the Child Wise report that says few accounts are shut down, and they are too easily replaced, for not reaching the age requirement? What is your comment on that?

Karina Newton: We have really easy reporting, and in fact you do not even have to be an Instagram user to report. We made a change last year as well. Now, if an account is being reviewed for any reason—it could be bullying, it could be any number of things—reviewers will look to see if the account appears to be of someone who is underage. We will then what we call "checkpoint" the user, which basically blocks them from using Instagram until they are able to prove with a photo ID that they are in fact over the age of 13. If they are not able to do that, we delete the account.

Q967 **Clive Efford:** Yes, but the point is that this is being done in too few cases. That is the point they are making.

Karina Newton: I wish we had the numbers to share with you, but I would dispute that.

Q968 **Clive Efford:** Why would you come here without them? Can I ask?

Karina Newton: We unfortunately do not have them. I do not want to sit here and claim we cannot be doing more or there is not more technology that we can invest in because we absolutely are, and age verification is something that people as an industry, lawmakers and Government are all trying to figure out a solution that works. It is not easy. We probably would have something in the duty of care already if it was easy. This is something that we are definitely working on. In terms of our ability to action on accounts that are reported to us and in terms of our ability to take down accounts, we are taking down accounts on a very regular basis.

Q969 **Clive Efford:** Perhaps you can supply us with those figures. You say you are working on it. What sorts of things are you working on to improve performance in this area?

Vishal Shah: Some of it is around proactive detection of people under the age of consent, so that might be the types of comments that they use or the types of captions that they use. This is a new and burgeoning area of technology and one where we can come up with a privacy-safe way of doing this. We do not have the clear answer on that yet. If we did, we would have implemented it already. The goal is to be able to detect these things more proactively in addition to the reactive takedowns that Karina mentioned.

Q970 **Clive Efford:** Are parental controls part of that?

Vishal Shah: Because we do not allow children under the age of 13, the controls that we have are for all users on Instagram.

Q971 **Clive Efford:** There is no way to apply any form of parental control on the site that would help you?

Vishal Shah: Maybe you can talk about the clarity around getting parental consent.

Karina Newton: We do not have parental controls on Instagram accounts. We do provide a parents' guide to help parents understand and have conversations with their young teenagers about how to use Instagram and what is an appropriate way to use it. We encourage people who are under 18 to have a private account, for instance, and not share their location. These are conversations that we are trying to empower parents to have as to how to have a safe experience on Instagram as a young person.

Vishal Shah: Just to be clear, from a product perspective, we are actively looking at if there are things we should be changing from a product perspective to make Instagram safer for young people and whether there is age-appropriate design. We are actively engaged in that conversation with the UK Government and others around things that maybe would change from a product perspective based on someone's age.

Q972 **Clive Efford:** What do you say to parents who want to ensure that their children do not access harmful content, either adult content or pornography? What would you say to them? What action can they take?

Karina Newton: My advice to parents—and I am a parent as well—is to make your account private, to only follow people that you know, and only allow people who you know to follow you. You could remove people who follow you on Instagram. If you have a private account, only those who you approve are able to see your photos.

In addition, we do not allow any messages to go into your direct inbox unless you follow that person. Everything goes into a pending inbox. We urge parents and young people as well to not accept any messages from people that they do not follow or know in real life.

We do also advise parents that you should be aware and have open dialogue with your young teenagers as to what kinds of experiences they are having. One of the things that I hear the most from parents if they are not an Instagram user is, "How do I talk to my kid about this? What questions should I be asking?" That is one of the things we worked on in the parents' guide to help facilitate that dialogue, because frequently there is a generational difference where they are not speaking the same language about what a safe experience looks like on this online platform.

Q973 **Clive Efford:** Final question. Do you age-gate any content or hashtags from minors?

Karina Newton: We do age-gate content.

Q974 **Brendan O'Hara:** Similarly to Mr Lucas, I am amazed that you are unable to tell us the average age of Instagram users. Who then at the company would have that information?

Vishal Shah: We have ages for people who are on Facebook. If there is a match or link between a Facebook account and an Instagram account that a user has linked to, then we would have an age for that user on Instagram. That is only the case where those accounts are explicitly linked.

Q975 **Brendan O'Hara:** No one at Instagram knows the age of the people that use the platform?

Vishal Shah: Unfortunately, the specific information that I have that I can share with you today is that we do not have ages for people who use Instagram directly because we do not collect age. Therefore, I do not have that data to share.

Q976 **Brendan O'Hara:** Would you be surprised if the age of the typical Instagram user was 65?

Vishal Shah: We can certainly all theorise.

Q977 **Brendan O'Hara:** Would you be surprised if it was 48? Would you be surprised if it was 22? Do you have no intelligence at all about your own customer base?

Vishal Shah: We certainly do research and we talk to people who use Instagram, and we know that there are young people who are on the platform.

Q978 **Brendan O'Hara:** Predominantly young people?

Vishal Shah: I would not be able to say that with any definitiveness without the data.

Q979 **Brendan O'Hara:** Who has that data, or does it not exist?

Vishal Shah: We do not have age data for people who use Instagram directly.

Karina Newton: Just to clarify, we do have ranges. We do not have data bursts for them. I would say that we do have some ranges. If you were to say, "Are the majority of people on Instagram 65?" my answer to that would be, "Likely not", but we are not able to provide that breakdown.

You might want to mention the time spent numbers that we have, and that is an age range that we have.

Vishal Shah: Using some of these rough ranges, we have talked about how people under the age of 25 spend about 30 minutes per day on Instagram, which gives you a sense of the usage of the platform, but we have not broken down the specific buckets or the ranges in terms of what the age and range of broad Instagram usage is.

Q980 **Brendan O'Hara:** On the point of usage, how many times a day does an average user open his or her Instagram account, and how long on average would that person stay on the platform?

Vishal Shah: It varies pretty widely. What we do is we provide controls from an aggregate basis in terms of time management. That is telling people how much time they do spend on the platform as well as being able to set limits against that time.

Q981 **Brendan O'Hara:** Yes, I know that. I will come on to that in a minute. From your own data, how many times a day does the average user open Instagram and how long does he or she spend on the platform in a day?

Vishal Shah: I understand the line of questioning and I would be happy to provide that privately to the Committee, the reason being that it is of commercial interest for us and for our competitors in the space. I am happy to provide it in terms of the breakdown that we do have if that is helpful for the Committee.

Q982 **Brendan O'Hara:** Thank you. What was the company's motivation that led it to introduce the usage monitoring tool?

Vishal Shah: There were a variety of motivations for the usage tool, primarily driven around people's wellbeing. What we have heard is that there is no clear answer on how much time someone should or should not be spending on Instagram or any digital platform or any medium at all, but it is important for people to have a sense of what their own usage is and to define sometimes for themselves what their usage should be and set limits, or in the case of young people to engage with their parents and to have a reasonable conversation about that in the same way they might have it for television, for example.

When we built that tool, we had an expectation that we would be limiting engagement for people who would sign up for these limits. Because wellbeing is an overarching priority for us, we were okay with the loss in engagement for people who decided to set a limit on their time.

Q983 **Brendan O'Hara:** Were there concerns that young people were spending too much time on Instagram?

Vishal Shah: There was a concern broadly that people wanted to have more control over the way that they were spending their time. That I think has also shown in other digital platforms, whether it is platforms like Instagram or even some of the mobile operating systems, where more visibility into how people are spending their time is something that a lot of platforms have introduced.

I think this is an important topic broadly speaking, not just for digital or for social media but one that we have been having for many years on entertainment with television or radio or magazines. This is a continued conversation in terms of how people spend their time.

Just to be clear with the Committee, I do think there is a limit for how long people should be spending on the platform. I just do not think it is the same number for everybody, and I do think that having these controls gives people the tools to make those decisions for themselves.

Q984 **Brendan O'Hara:** What proportion of Instagram users use this facility, the "Your Activity" page? How many of them have set it up?

Vishal Shah: We have only launched those tools about six months ago and we are still monitoring usage and making sure that we understand exactly how people are using it and what they are opting into. I don't have those specific numbers to share with you today because the—

Q985 **Brendan O'Hara:** Could you give us any indication? Have you been surprised by the uptake, disappointed by the uptake? What targets had you set for the first six months or the first year of uptake and have they been met or not met?

Vishal Shah: We have been surprised, but pleasantly so, that people have really responded to these tools and that was important. Instagram is a platform to be able to report on some of these numbers. Most of that research has been anecdotal through our research after we launched the tool. We do not have any specific goals against usage, nor do we have any specific numbers to share against them.

Q986 **Brendan O'Hara:** You must have specific numbers of people who have availed themselves of this, surely? That number must exist?

Vishal Shah: Unfortunately, I do not have that number to share with you today but we can follow up on that.

Karina Newton: May I also just say we are really invested in getting people to use the tools? We want them to use this. We want them to have meaningful experiences on Instagram and be very aware of how they are spending their time. We held a design hack, for instance, here in the UK with several charities and we wanted to also involve young people in that conversation as to, "What would convince you to turn this on? What kind of messages reach you?" That is something that we are also working on. Now that we have fully rolled out, I believe in December, to 100% and now that we have had four or five months of product testing and ensuring that it is working as people intend, we really want to get the word out to ensure that people avail themselves of this control.

Q987 **Brendan O'Hara:** Isn't there a contradiction, though, that you are introducing this usage tool at the same time as you are creating features such as IGTV in order to keep people hooked and keep people on the platform? There is an inherent contradiction, surely.

Vishal Shah: I do not think that is a fair characterisation of the motivation behind IGTV or many of the tools that we build. The primary goal of Instagram is to let people express themselves and for people to discover their interests and the people that they know and care about. IGTV specifically was built to help people connect with the interests that they have and to be able to consume longer form content than was currently available on Instagram through our existing products like Feed and Stories.

We think of wellbeing generally as important across the entire product, not just our primary uses around Feed and Stories but including IGTV. To my point earlier around knowing that there would be some reduction in engagement with these tools, I believe that overall time spent and management of that time was important to us that it applied across all Instagram.

Q988 **Brendan O'Hara:** What research did the company do before introducing the usage monitoring tool in order to establish why they were introducing it and what goals it set itself to test how useful it had been?

Vishal Shah: We use a variety of research techniques for any of the products that we build and specifically ones on wellbeing, and maybe Karina can chime in on what we have learned from experts and some of the research we have done there.

Karina Newton: On wellbeing and generally, there are a number of things that we get from both qualitative and quantitative research to understand what people are struggling with and things that they find unpleasant and so on. On time spent tools, this is actually a decision that was made from the top down. Kevin Systrom, who is our co-founder and former CEO, said, "I think we should have these time management tools. This is the right thing to do". There was not any research really that led into it saying, "We must do this because X% of users want this". That was not where it was coming from. It was coming from the place of, "We believe in controls. We believe in giving people control over their experiences".

I think that where the research and testing comes in is how easy they are to use. What options do people want? Do they like a scroll bar? Do they like an alarm buzz? A lot of the research that we do is about the product itself and making it work as people expect it to.

Q989 **Brendan O'Hara:** Finally, have you or Instagram ever been approached to share the user engagement data with academics or researchers looking into wellbeing and, if you have not and you were to be in the future, would you be prepared to?

Karina Newton: We are actively looking at that right now. We are in active conversations with the Samaritans as a result of the conversation with the Secretary of State as to what information would be helpful. I think that there are also privacy concerns to ensure that we are giving

aggregate data, but this is definitely something that we would like to explore for the wellbeing of our entire community. We cannot do all the research ourselves. Academic institutions are incredibly important, so it is something that we are looking at.

Q990 **Chair:** Have you seen the Information Commissioner's Office draft age appropriate design code that has been published in the UK?

Vishal Shah: Yes, we are very engaged in that conversation and the conversation that we have been having internally on the product side in addition to that.

Q991 **Chair:** I suppose one of the central things that the Information Commissioner is recommending is that for under-18s using social media a lot of the prompt and reward tools should not apply, in particular things like likes, which might rank popularity of post and could lead to people feeling upset about the number of likes they either have or do not have. Are those recommendations that you are considering?

Vishal Shah: In the specific case of things like likes, we are looking at that broadly independent of someone's age. This goes back to the questions that we were asked earlier around what something was designed for originally versus how we saw it evolve and what it is being used for now.

Just to clarify why likes were created on Instagram, when you express and when you share something to the platform there was no indication, certainly back in 2010 when the product launched, that someone was connecting with you and seeing that content that you share, that you express. A like was a way of knowing that someone saw that content and appreciated it.

As the platform evolved, as it grew, as the amount of followers that people had grew, as the amount of public figures and professional content creators grew, the pressure around likes increased. That is something that we have certainly been aware of in the research but also something that we have seen ourselves. We are running a test now in Canada to remove like counts from broad availability and visibility. The idea is to reduce the pressure around feeling like you are not only competing with yourself on your previous posts but with the rest of Instagram, frankly, and feeling that every time you have to always be perfect and achieve a certain number of likes in order to feel like your expression was worthwhile. That is something that we are looking at more broadly, not just for young people.

Q992 **Chair:** But likes is an advertising tool as well, isn't it? It is a data-gathering tool that you use for advertising?

Vishal Shah: Our advertising tools are based in outcomes or things that advertisers are hoping to achieve, whether that is sales or downloading of an app. We specifically think of likes as a proxy to real outcomes for

advertisers. From the beginning, Instagram has always had outcome-oriented advertising products.

Q993 **Chair:** No, but the data about what people like is part of the data that exists about that user, which is then valuable to advertisers, helping to identify audiences. It is a very valuable data tool for you, likes, isn't it?

Vishal Shah: Maybe to take a step back, likes is a product or a feature that is part of Feed, a part of our Instagram Feed product, but not something that is part of Stories. Something we learned in Stories is that people view content and discover new topics and new interests, and because feedback is private—because you cannot provide a like and you can only respond through a private message—people are still able to understand the value that is being created. Advertising happens in Stories without likes. We are taking some of those learnings of how the platform has evolved both in Feed and with new expectations in products like Stories and seeing how we can apply that more broadly. It is very much a wellbeing-driven initiative knowing that there is going to be some engagement loss as a result.

Q994 **Chair:** Yes, but given you have the same advertising tools as Facebook, if I was to go to Facebook and say, "I am interested in advertising to people who have these interests", then likes data is something that helps to develop those audiences, isn't it?

Vishal Shah: Interests on Instagram are pretty broadly defined. It can be the accounts that you follow. It can be the content that you see.

Q995 **Chair:** The likes is part of it, isn't it? That is what I am getting at.

Vishal Shah: It absolutely is part of it.

Q996 **Chair:** The simple question I was just trying to get to was that likes is not just about showing what people like, it is a data-gathering tool for the company as well and that data is used to help advertising targeting.

Vishal Shah: Certainly, the data that is generated by the platform is used from an advertising perspective.

Chair: Yes, and that includes likes?

Vishal Shah: That does include likes, yes.

Q997 **Chair:** For younger users in particular, you talked about making likes private, which means that the user cannot see what people have liked but the company can because the company is still gathering the data. Do you think that for younger users that function should just be turned off altogether? Is it appropriate that data is being gathered about a 13 year-old in order to serve advertising to them?

Vishal Shah: If I could answer your question in two different ways, one is that we have given people control over the way that their content can be interacted with. In the same way that there was a decision to introduce the time spent tool, another decision like that was to give

people the ability to remove comments or turn off comments for an individual piece of content. It was something that was not very common in social platforms but we decided it was an important thing for people to not only be able to turn them off completely but be able to restrict the kinds of words or terms that people could put in their comments.

The first step in removing like counts is the idea of removing the pressure around comparison and just being extremely clear on what the test is. As someone who is producing a piece of content, you can see for yourself how many likes you received, but no one else can see how many likes you received. It is information you might know and, again, important to the purpose of expression but not something that you have a fear of being on a public competition or stage with the people around you.

Q998 **Chair:** Okay, but the question I was asking is about younger users. Given that the interests of the company in gathering certain data is about having rich datasets that facilitate advertising—because ultimately you are an advertising business and you make your money from selling advertising—is it appropriate to gather the same sort of data about a 13 year-old that you have gathered about a 30 year-old?

Karina Newton: We are really involved in this process now and we would love to invite the Committee as well to learn more about how we are thinking about age across the company. I think that many of these ideas are being thought about holistically and we are definitely reading up on what people are saying and taking that in. I think that this is going to continue to be a dialogue that we all can learn from each other on.

Q999 **Chair:** I just want to change the subject. I want to ask a question about filters and body image. There are filters on Instagram that can make your lips fuller or your hips narrower and some people have raised concerns that this is creating unfair pressure in terms of people's body image and self-esteem. Is that something that Instagram is considering reviewing? Is that a feature that you might consider discontinuing?

Vishal Shah: It is something that we are looking at closely. Just to talk about it from a product perspective—and then, Karina, please chime in from a policy perspective—there were a few important considerations for us when we were launching filters, and filters on Instagram are much more recent on our platform than on other platforms. One was that it had to be opt in, that people had to choose to apply a filter, and also that it was not restricted to filters that can alter appearance but things around putting on a hat or changing the world around you. That being said, I think that body image is a really important topic not just on Instagram but broadly speaking and this is something that we are taking seriously, especially in some of the research that we are doing.

Q1000 **Chair:** Sorry to interrupt but just to be clear, though, the filters do allow you to change your physical appearance, don't they, as well, not just add a hat or something else?

Vishal Shah: There are some filters that allow you to do that, yes.

Karina Newton: I was just going to say that one of the things that I think Instagram has done for the good on this is provide different images of bodies. For decades we have had what an ideal body is. Sometimes that image has changed but there has normally been just a few. I think that Instagram and social media in general has done a great deal in terms of I can now see people who look like me on Instagram. These are different voices, these are different ethnicities, these are different abilities, these are different interests.

If I may give a personal example, when I was 13 I was diagnosed with scoliosis, which is a curvature of the spine. I had the joy of wearing back braces for several years until my curve progressed so much that I had to have spinal fusion surgery. I have this moment seared into my brain when I was in the bottom of a department store basement trying on dresses for my prom and I was trying to figure out what I could wear over the brace. I just felt like the most awkward person ever and I didn't know anyone in my school who had scoliosis. I didn't know anyone who had gone through spinal fusion surgery. Teenage years are obviously awkward enough but then you add that on.

Now if I go to Instagram, I can find #scoliosisfashion. I can find women who are going through the same thing as I am. I am now a professional woman who is quite confident and body positive but I was trying on my wedding dress and I was trying to figure out how to hide my curve. There are ways in which body positivity impacts women and men differently over the course of their lives, and I think that social media is an incredible forum to be able to connect with people who might be feeling the same thing, people you might not have access to in your immediate community. I would say that Instagram on the whole is providing a new way for people to see and connect and have other ideas of what beauty and body standards are.

Q1001 **Chair:** That is a great story and I do not dispute that at all. That is certainly a clear benefit that social media brings. I suppose what you are describing is people sharing images of them as they are. What I was asking about was if you can use filters to change your physical appearance, to make yourself in some ways conform to someone else's view of what beauty looks like. I would be interested to know the stats on this as to who uses these filters. My suspicion would be that the sort of filters that are used to change your physical appearance to make yourself conform to an idea of beauty are probably predominantly used by younger users, particularly younger girls. I think that does feed back to the issues about self-image, self-confidence and maybe even self-harm that we discussed at the beginning. I would think a responsible attitude towards considering this should be maybe examining whether these filters are actually helping. It might make some people feel happy but it may be negatively affecting a lot of others.

Karina Newton: We are looking at how filters can help. We had a filter for kindness. We have a Ramadan filter this month celebrating the month

of kindness. It is another way to express yourself. We also have filters that are used in different ways, but I definitely take your point that we need to be vitally aware of how filters might impact people's wellbeing. We are taking a thoughtful approach here but it is definitely something where we cannot ignore the impact on society and part of the wellbeing research that we are continuously investing in.

Q1002 **Chair:** What we are talking about is not filters per se, because filters can be used for all sorts of different things and many of them are fine. We had the same conversation with Snapchat when it gave evidence to the Committee as well. Where you have filters that are designed to effectively provide a form of digital cosmetic surgery on your image, is that appropriate? Is it appropriate for a 13 year-old girl to be trying to make her lips look fuller in an Instagram post because she thinks that is conforming to an idea of what beautiful looks like?

Vishal Shah: I think that it is an important point considering the broad use case of whether we should be taking a harder look at some of those filters.

Q1003 **Jo Stevens:** Can I ask for clarification on one of the answers that you gave earlier? We were talking about data shared between Instagram and Facebook. Did you say that if you have an Instagram account that your Instagram account could be linked or connected to your Facebook account?

Vishal Shah: Yes, that is correct.

Q1004 **Jo Stevens:** What proportion of your Instagram customers or account holders are also on Facebook as well?

Vishal Shah: We do not have that data here today, but I am happy to follow up with that specifically if that is important to the Committee.

Q1005 **Jo Stevens:** The age verification system we have just been looking at. It does not work, does it?

Vishal Shah: As we talked about earlier, the current systems that we have in place are designed to ask people for their age and if they circumvent them—

Q1006 **Jo Stevens:** Well, "Are you over 18? Are you under 18?"

Vishal Shah: That is correct. If they bypass those, then there are ways in which we try to detect underage usage or people under the age of consent. There is more work to be done here. We can do more to improve the ability to catch this upfront.

Q1007 **Jo Stevens:** Say I am 10 years-old; I am obviously not but say I am 10 years-old. I can just go on to Instagram and I can set up an account because I can just say I am over 18. I can just do it, can't I? The system does not work.

Karina Newton: This is one of those things that as an industry age verification is what we are trying to figure out.

Q1008 **Jo Stevens:** But why didn't you work it out before you rolled out your product? We have talked about this in previous sessions. If you have a pharmaceutical product, you have to go through product safety testing. Why are you always reacting to things rather than proactively doing it?

Vishal Shah: I would not characterise that as reactive versus proactive but more that in general when we let people sign up for Instagram we are trying to collect as little information upfront from them as possible. It is why we simply ask for whether you are over the age of 18 or not. Verification generally is challenging for all the reasons that Karina mentioned. From a privacy perspective, there are more than a billion people who use Instagram around the world and not everyone has access to a form of identification. Not everyone has access to a passport. Sometimes what seems like a fairly straightforward and simple way to verify age actually eliminates usage of the platform for many of the world's population. We want to find a balance there where we can hopefully detect some of these things proactively. We certainly have policies in place now and we can remove folk reactively, but the goal is to get to a place where we can detect more of this stuff proactively and we can do a better job upfront.

Q1009 **Jo Stevens:** When was Instagram first launched?

Vishal Shah: Instagram was first launched in 2010.

Q1010 **Jo Stevens:** It has been going for nine years with an age verification system that does not work. When are you going to have a system that does work?

Vishal Shah: We are committed to building a system and working with the rest of the industry to figure out the best way to do this. I do think technology has a role in making this something that is applicable for everybody.

Karina Newton: I would also add that technology is improving every day and some of the ways in which you can use technology to identify things such as this are rather new. That is one of the things that we are looking closely at and I am sure other partners in industry are as well.

Q1011 **Jo Stevens:** I wanted to ask you about influencers on Instagram. Until very recently, I did not know anything about this and I find it really intriguing. We know from our previous inquiry with Facebook that there was a big risk of fake accounts on Facebook and that distorts the digital market. What steps do you take to make sure that you do not have fake accounts that are distorting or artificially inflating influencers' popularity? The figures for the number of people who follow and look at what the influencers do are enormous, aren't they?

Vishal Shah: In general, we think of all accounts on the platform, not just creators or those with large followings. The engagement has to be authentic and it is against our policies for engagement to be inauthentic. There is a variety of ways and means in which that happens today. There are two different ways to talk about it. Sometimes there are purely fake accounts, which we, of course, want to remove off the platform. But many times there are people who have given their real username and real password to a third party service in order to inflate their numbers, to create fake engagement on to the platform. We have taken very aggressive steps recently—and Karina can talk about this in more detail—to both detect and shut down those networks that create this fake engagement, both from a systems perspective as well as from a legal and criminal perspective.

Karina Newton: We have recently been taking more actions against these fake engagement farms that are essentially trying to bump up numbers for a paid service or other areas. We have filed a number of lawsuits to crack down on it as well as starting to remove inauthentic engagement from accounts when we can detect it. We are building detection systems internally to detect this fake engagement. We are also sending notices to users that we suspect of creating fake engagement or using or benefiting from some unknowingly because they have signed up for a third party service and logged in with their username and password, perhaps to help manage their Instagram account, and now that they have given over that username and password there may be other actions taken on their account that they are not aware of. Not everybody who has fake engagement on their account is aware of it, so part of what we are doing is helping inform them, encouraging them to change their password, which then cuts that link from the bots or computers. Inauthentic engagement is something that we are cracking down on from so many different levels, from the legal side, from the technology side. Of course, it is against our policies.

Q1012 **Jo Stevens:** Quite a lot of your most famous influencers are children. I have just been looking at the McClure twins on Instagram. What is your position on whether United States child labour laws apply to children influencers on your platform?

Karina Newton: I am not familiar with that law, unfortunately. I think that might be an issue that we need to circle back with you.

Q1013 **Jo Stevens:** You technically have a platform that does not allow under-13s on to it. That is your age limit. You have children on the platform who are influencers with millions of followers. They clearly cannot own the accounts because they are below your age limit. They are doing work on your platform. That work is being monetised, but the children are not being paid that money. Do you have any kind of system or controls about what is happening around that?

Vishal Shah: From a branded content or paid partnership perspective, we have seen from an influencer and creator perspective that there was

inconsistencies in the way people were disclosing these paid relationships, sometimes using hashtags like #ad or #sponsored but doing it in very inconsistent ways. From a transparency perspective, it was important to us to make sure that from a creator perspective as well as for a person who is seeing that content it is clear what their financial relationships are. We started testing a tool last year and we are in the process of rolling that out more broadly now to very clearly let any sort of branded content be clearly disclosed inside of the post so that it is clear why the post was there. I also want to make it clear that from an Instagram perspective we are not participating in any of those branded content deals and we do not monetise as a result of any of that branded content.

Q1014 **Jo Stevens:** But other people are using your platform to monetise the work, whether it is consented to or not, of children on your platform, aren't they?

Vishal Shah: We have been very committed from a transparency perspective to ensuring that these disclosures are clear. As Karina mentioned, there are some specifics that neither of us are as familiar with in terms of individual laws that we will be happy to follow up on.

Jo Stevens: I would be very interested to hear about that. Thank you.

Chair: Thank you very much. I think that concludes the questions from the Committee this afternoon. Thank you very much for your evidence. Thank you for travelling all this way to be with us.

Karina Newton: Thank you for having us.

Vishal Shah: Thank you.