



Joint Committee on the Draft Online Safety Bill

Corrected oral evidence: Consideration of government's draft Online Safety Bill

Monday 25 October 2021

5.15 pm

Watch the meeting: <https://parliamentlive.tv/event/index/cf4e4a95-fbc3-4779-8e78-12f46acbadfd>

Members present: Damian Collins MP (The Chair); Debbie Abrahams MP; Lord Clement-Jones; Lord Gilbert of Panteg; Baroness Kidron; Darren Jones MP; Lord Knight of Weymouth; John Nicolson MP; Dean Russell MP; Lord Stevenson of Balmacara; Suzanne Webb MP.

Evidence Session No. 14

Heard in Public

Questions 223 - 232

Witnesses

I: Leslie Miller, Vice President, Government Affairs and Public Policy, YouTube; Markham C. Erickson, Vice President, Government Affairs and Public Policy, Google.

USE OF THE TRANSCRIPT

1. This is an uncorrected transcript of evidence taken in public and webcast on www.parliamentlive.tv.
2. Any public use of, or reference to, the contents should make clear that neither Members nor witnesses have had the opportunity to correct the record. If in doubt as to the propriety of using the transcript, please contact the Clerk of the Committee.
3. Members and witnesses are asked to send corrections to the Clerk of the Committee within 14 days of receipt.

Examination of witnesses

Leslie Miller and Markham C. Erickson.

Q223 **The Chair:** Good afternoon and welcome to the second panel for our evidence session this afternoon. We are delighted to welcome Leslie Miller and Markham Erickson from Google.

I wonder if I could start off by asking about your advertising policies in terms of complying with local laws. Do you have effective policies in place to make sure that advertising that is in breach of UK consumer protection law is not served through Google's various platforms?

Markham C. Erickson: Good afternoon and thank you for inviting us to appear before the committee. We appreciate the opportunity to talk about our services and the proposed legislation.

When it comes to online advertising, yes, we have robust publisher guidelines and advertising protocols that would preclude advertisements that perpetuate frauds or scams, or would be harmful or promote illegal content, as defined under local laws.

The Chair: From your point of view, if you have those policies in place, you will have no objection to that being required of companies in the Online Safety Bill, to make sure that they do not allow illegal content—such ads would be considered a form of illegal content—to be shared through advertising.

Markham C. Erickson: Thank you for the question. As the committee recommends how to approach the scope of the legislation, we would encourage it to exclude the online advertising ecosystem. I say this for a couple of reasons. One is that it is already a regulated space. In the UK, the ASA regulates online advertisements, and there are a number of stakeholders who are responsible for ensuring that advertisements are appropriate. It includes the advertisers themselves, businesses, the technology companies and, in appropriate cases, separate regulators, for instance, in the financial services space. By including it in the proposed legislation, one would want to be sure that you are not creating confusion about who would need to comply with the Bill. We need to ensure for Google's products and services that the online advertising ecosystem is safe and trustworthy. It is the only way that we will ensure that advertisers use our products and services, and that users feel they will have a trusted experience when they click on advertisements.

I would make one further point, if I may. As I understand it, the Government have anticipated that the current scope of the legislation, which does not include the online advertising ecosystem, would regulate roughly 24,000 entities, and by including the online advertising ecosystem, potentially, the legislation would be regulating millions and millions of businesses that advertise on our platforms every day. We think there are safeguards in place both through separate regulation and through our natural incentives to ensure that this is a trusted and safe

experience.

The Chair: The advertising regulatory system in the UK is a self-regulatory model overseen by the Advertising Standards Authority, and financial advertisers comply with the FCA—Financial Conduct Authority—regulations, but that self-regulatory code relies very heavily on brand jeopardy on behalf of the advertiser; that is, that it would be shaming for an advertiser to have an advert that was considered to be in breach of the code, so they would withdraw it, and indeed traditional media companies, broadcasters and news organisations would not run those ads because they would be in danger themselves.

The issue here with online advertising is that we do not have a problem defining what we think an illegal ad is or a problem with interpreting the Advertising Standards Authority's guidelines. What we have a huge problem with here is enforcement online. Are you saying that, if you had guidance from the ASA that an ad was in breach of the code, Google would have systems in place to make sure that that ad, or a copy of that ad, did not appear anywhere on its systems?

Markham C. Erickson: What I meant to say is that it is an ecosystem that involves many different parties to ensure that the advertisements that one sees on Google are safe. However, we do not rely on others to ensure that it is a safe and secure experience. We do not wait for legislation to be passed to do so.

You mentioned the financial services space. We have partnered with the Financial Conduct Authority to come up with a system to ensure that financial services companies that want to advertise on Google platforms and services have to demonstrate that they have been certified by the FCA to sell their services and promote their services on Google products and platforms.

The Chair: In that case, you have effective enforcement of that. In theory, if an organisation is registered by the Financial Conduct Authority, action can be taken against it by the authority for offering poor services and products to customers. You are saying that you have policies in place that mean that a financial ad that is not FCA approved, or from an organisation that is not FCA approved, would be removed from Google's platforms and services.

Markham C. Erickson: We take it one step further. Before a financial services company in the UK can propose an ad for Google platforms and services, it first has to certify that it has been approved by the FCA as a legitimate financial services company that can market its services to the public on Google platforms and services. It is an ex ante approach rather than an ex post enforcement approach.

The Chair: That is the ex ante approach, and that is fine, but what if they slip through the system? Are they removed retrospectively? If there is an ad out there that for whatever reason has been able to run, if it is brought to your attention, will you always take it down?

Markham C. Erickson: Yes. However robust we try to make the system by working with the FCA—we are proud that the chair of the FCA has said that we have an industry-leading approach to this space—no system is perfect. If we were to be made aware of an ad that was appearing that should not have been, we would take immediate action and block the ad from appearing. We blocked nearly 3.1 billion ads in 2020, including 123 million ads from financial services companies. We certainly do not want any ad that is harmful or perpetuating frauds to appear on our advertising platform.

The Chair: That is clear on financial advertising, but that is not the only type of scam ad. We had a problem here for years of ticket touts selling tickets through secondary ticketing sites such as Viagogo that were in breach of consumer protection legislation, and Google ran them for some time. I think the policy has changed now. In a case like that, where someone comes forward and says, “This product is being offered, is being sold in breach of consumer protection legislation and is illegal, and other advertisers are not accepting these adverts for that reason”, if it was a clear breach of local law, would you not accept the ad either?

Markham C. Erickson: We have over 40 different policies that are made clear to advertisers and publishers around what kind of content we will not accept, including advertisements that perpetuate frauds. When we see that there may be a higher risk to consumers of fraud or scams, we will take it a step further and create more robust certification programmes. A secondary ticket ecosystem is one such example of doing that. We have policies where, before advertisers can show their ads for tickets, they have to comply and certify that they are complying with the policies that we have in place. If we find that one is not doing so, we will of course remove it.

The Chair: So it would be fair to say your policies on advertising require compliance with domestic law.

Markham C. Erickson: That is right.

The Chair: There is a final point on the Advertising Standards Authority. If it notified you of an advert that it said should be withdrawn because it was in breach of the advertising code in the UK—as you said earlier, that is the regulatory body—again, would Google remove that ad, or copies of that ad, from running on its services?

Markham C. Erickson: We work with the ASA and we also work with other regulators, advertisers and businesses in the UK to ensure that in any such example like that, yes, we would take down that ad and take action.

The Chair: Going back to my earlier question, if you say you comply with ASA guidance, and you have to comply with local laws, and you comply with the FCA in terms of financial ads, and you apply all those rigorously to any ads that appear on your system, presumably, if the Online Safety Bill said that that is what we want people to do and what we want all

companies that serve ads to do, then you would have no objection to that because you do it anyway.

Markham C. Erickson: Certainly we have no objection to the objective that this committee is trying to achieve. The question is whether including potentially millions and millions of businesses in this advertising ecosystem is proportionate to the benefit that would accrue. We do think this is a safe and secure ecosystem. We are highly motivated to ensure that online advertisements are safe and secure. Advertisements that are legitimate and that are on our ecosystem do not want to appear or compete with advertisements that are fraudulent, or deceptive, or perpetrating fraud. I would hesitate to recommend that you include the ecosystem in the scope of the Bill. Certainly, we will continue conversations with you to explain how we approach various different aspects of the ad ecosystem and who we work with.

The Chair: With regards to proportionality, we have heard heart-rending evidence during this inquiry so far about the impact on people's lives of financial scams and frauds, including from the Competition and Markets Authority, the Financial Conduct Authority and the City of London Police, which investigate financial crimes in this country. They all support the inclusion of these provisions in the legislation. Even the Advertising Standards Authority said that it has no objection.

I do think it is important because I do not think it requires the regulation of the entire advertising ecosystem. What it requires, though, is that, when evidence of fraud and criminality is exposed, the media owners and the platforms respond by not allowing those ads to run, or any copies of them. I am just asking whether you would have any objection to that. I am not sure what your grounds of objection would be, given you say you largely try to do that anyway.

Markham C. Erickson: Certainly we will continue the conversations with you. We would want to see how that would be approached, recognising the caveats I mentioned of ensuring proportionality. We do think it is not necessary in terms of our incentives and how we operate. We will work as quickly as we can, as soon as we are made aware, to take down advertisements that violate local law and perpetuate frauds and scams. When we see that, I think our record shows that we take swift action to protect consumers from those kinds of scams.

The Chair: It may not be proportionate to your incentives, but it is proportionate in protecting the public interest. The suffering that many people have had as a consequence of scam ads that run on different social media platforms and online in different spaces almost unchecked, in a way they do in no other media, is a problem that people are concerned about.

Finally, I wanted to ask about advertising transparency. We have had evidence during this inquiry that a lot of platforms, websites and supposed news organisations profit from sharing and spreading hate speech. One way they do that is through programmatic advertising. Do

you think brands should be able to see a schedule of the sites where all their adverts appear? In the old days, everyone always knew, but you do not know when you are buying ads through an online auction. Do you think there should be more transparency so that brands do not see their ads ending up in places where they would never have wanted to advertise?

Markham C. Erickson: Thank you for the question. We are very transparent with our advertisers about the kinds of content and the publishers against which their advertisements will appear. If we find that an advertisement is mistakenly appearing on a site that is promoting hate speech, racist speech or other speech that violates our publisher guidelines—our publisher guidelines are explicit and clear on this point—we will remove that ad swiftly.

The Chair: If you found an ad served in a publication like that that was in breach of your guidelines, would you withdraw any ad tech support from Google for that website?

Markham C. Erickson: We would remove the ad altogether, yes. How our ad policy works is that there is the organic content that may appear on a news site and then there may be comments that appear relative to or next to or adjacent to the news. Sometimes it is not the publisher's content that has violated our policies, but comments appear from users that are racist or otherwise objectionable. In those cases, we will also remove the ad, but if that has been taken care of and we are working with the publisher, it is not a permanent ban against that publisher for its ability to monetise its publication.

The Chair: But would there be a cooling-off period? Would you say, "You are clearly operating outside our guidelines, you have no intention of changing what you are doing, and, therefore, until we see consistent evidence that you are running safe content, we will remove our ad support from your site"?

Markham C. Erickson: The way we approach this is that we would take action swiftly, and until we felt the publication was in compliance and we felt comfortable about that, we would not reinstate the publication for the appearance of ads.

The Chair: Thank you. Beeban Kidron.

Q224 **Baroness Kidron:** I was interested in your written evidence. It said this: "To best protect children while retaining their access to online services, it is vital that the Committee carefully considers how the Bill's child protection duties can be simplified and aligned with the Age Appropriate Design Code and existing international standards like the EU's AVMS Directive." I would like to ask you, first, what is the danger that you see, and, secondly, what would you see this approach adding above and beyond those two pieces of existing regulation?

Markham C. Erickson: I am certainly happy to begin and talk about this from a search perspective. When it comes to search, first, we would

encourage the committee to distinguish and recognise the distinctions between a search service and a user-to-user platform. The legislation acknowledges that there are important distinctions between the two ecosystems, but, unfortunately, as currently written, the duties that apply to both are nearly identical. We think this would potentially create an over-censorship of legitimate content—content that is important relative to freedom of expression—or result in a kind of age-gating that would substantially change the way search operates.

Our first point would be to distinguish between search services on the one hand and user-to-user generated platforms on the other. Search services, as you know, are not about engagement but rather about discoverability. A user comes to a search service looking for information, and we want to provide timely, relevant and helpful information in that moment. If we are successful in providing appropriate content that is relative to the search query, that is our first line of protecting the users who come to our search services.

Baroness Kidron: I will come to you, Leslie, but on that point, I am just a little bit confused because I would imagine that the aim of the Bill is to find appropriate search for under-18s. Could you give me an example of how censorship could happen because of how the Bill is currently drafted? I would just like to understand it.

Markham C. Erickson: Thank you for the question. As you know, search services do not host user-generated content. They are an index of trillions of web pages. We provide users with the ability to find relevant information relative to those index pages. The broad definition of online harms would require us potentially to have to make contextual decisions about the nature of content that we do not host, which is on another service, and require us to monitor those billions of web pages in doing so.

To give you an example, the legislation as currently drafted for search would require us to determine whether content could have a material negative “psychological impact on a child of ordinary sensibilities”. An example of how difficult that would be, in addition to trying to look for that across trillions of web pages, is something like the “Falling Man” image after 9/11. When it was first published, there was considerable outrage that it was harmful to show this image. It was horrific; it was damaging to the psyche. In retrospect, observers today would say that that image was important for the historical record, and that it captured the moment in an important way for the public good. There are other examples.

Baroness Kidron: I am sorry to cut across you; I apologise. I want to go the other way and say that we had some evidence from a couple of people who came in and said that you can put into Google, “How do I kill myself?”, and there were 139,000 sites or whatever the figure was—an endless figure of places—offering an answer. That is the other end of it, is it not? I am very sympathetic to what you are saying and I am trying to understand what good looks like, but I am interested what your position is when we are confronted with that evidence.

Markham C. Erickson: Thank you for that clarification of the question. I appreciate the way you have approached this whole ecosystem in the manner you have.

First, let me step back and say that we want to provide users with helpful and relevant information in times of trouble for that user. For things such as health or suicide ideation or self-harm, we work with experts to, first, promote content at the top of a search result that gives information that can be most helpful and appropriate for the user who is searching for that.

In the UK, we work with the Samaritans and with Shout. We work with them not only on providing ways for the user to connect with someone who can help but by being very careful about how we display information that is appropriate. There are important considerations here, to ensure that we are not unintentionally glorifying self-harm or suicide. We have a terrific partnership with these NGOs and we have similar ones around the world on those kinds of things. We feel responsible for providing helpful information at all times, but in those acute situations around health and safety, we try to promote information that is going to be responsible and helpful to a person who might be in crisis.

Baroness Kidron: I do want to go to Leslie so that she has a chance to come in on this, but, in your view, is it not a problem? Are we worrying about the wrong thing here?

Markham C. Erickson: We agree with the objective here, which is to ensure that UK citizens' security and safety are enhanced.

In terms of balancing proportionality and freedom of speech and privacy, I would encourage the committee to look at search services and acknowledge the unique distinctions that they have relative to freedom of speech and how they operate, and be proportionate in its response.

Baroness Kidron: Okay. Thank you.

Markham C. Erickson: I will leave it there, but I think we are aware of those kinds of situations. I would not say there is no problem in society, but when we find ways in which we can be helpful, we feel a responsibility to do stuff.

Baroness Kidron: Leslie, do you have anything to add?

Leslie Miller: It is nice to see you, albeit virtually. Thank you so much for inviting us and having a conversation about what I think we all realise are shared goals in making sure that users, particularly kids, have a safe experience online.

On your question about what we have real concerns about, on kids' safety and where it could be improved as it relates to the Bill, I certainly agree with the goals, but it would be helpful to have a little more clarity on what we mean by definitions. I think it is about priority and primary content. There are definitions around the need to prevent material being

made available versus having systems in place to mitigate the risk of that material. It is just a matter of getting almost in the weeds of the Bill itself to make sure that, as we talk about these terms, there is a shared understanding of what the definition is, as well as the expectation for the companies as to what our obligations will be, which balance the need not necessarily overly to monitor proactively but to make sure that we are capturing this content and doing it efficiently at speed, and making sure that people are not being confronted with it.

Baroness Kidron: The committee has talked a lot about safety by design, which is taking the age-appropriate design system as a model and saying, "Let's have a look at the features that seem to be problematic. Let's look at the spread, the groups, the 'recommend' and so on". Are you sympathetic to that view, which is a little more upstream than perhaps some of those definitions seem to suggest?

Leslie Miller: I am certainly open to having more transparency around the community guidelines for every company, including us. For example, when YouTube talks about its community guidelines, it has a quarterly transparency report. I do not know how often people read them. We have rolled them out since 2018, but we realise that people say, "Well, that is table stakes now. We need additional information as it relates to the percentage of violative content that people are viewing on the platform". That was why we rolled out the violative view rate earlier this year. I am open to the idea of staring at the systems and practices, making sure that we are appropriately enforcing our terms of service and community guidelines for content that is legal but has been determined to be harmful.

Q225 **Baroness Kidron:** I have one last brief question; I am not sure who should answer it. You also talk about your digital literacy programme for seven to 11 year-olds. It is creative and very good as far as it goes, but it does not deal with compulsive technologies and systems; it does not deal with data extraction. One of the things I noticed—I have told your colleagues in London about this—is that it talks about a lot of features that are available on services that are really only for 13 year-olds, but it does not say you should not be on them before you are 13. I just wonder whether you would be sympathetic to having official mandatory guidance on what digital literacy and digital safety should be taught in schools so that kids get the full picture of what is difficult about the digital world, rather than just bad actors and bad behaviour.

Leslie Miller: I can speak about this on behalf of YouTube. Please interrupt me if I have misunderstood your question. We have done several partnerships in promoting digital literacy across the board. This includes during the time of Covid when we have rolled out campaigns, for example, with the NHS to encourage younger people to get vaccinated. As I think you know, we rolled out YouTube Kids in 2015 in large part because YouTube main is not a platform meant for those under 13. There is always more that we can be doing on YouTube Kids—for example, setting Autoplay to default off and encouraging healthy behaviour.

Baroness Kidron: I take your invitation and am very grateful for it. Personally, I was absolutely delighted that you turned off Autoplay. Thank you very much. However, what I am talking about is what goes on in schools. If you are doing assemblies and not saying that there are age restrictions, and not giving the full scope of the problems of interacting with digital services, you are not really educating kids. We are concerned—at least I am concerned—that there may be a situation where you are teaching children half the truth and no one else is picking up the rest. Would you welcome official guidance that says, “It’s great that you are doing this work, but you must cover these subjects, even if it is a little self-critical”?

Markham C. Erickson: We want parents and minors to understand how they can best protect themselves online, including with our services. We have developed our educational programmes with experts. We have developed a curriculum with the Personal, Social, Health & Economic Education Association. I am happy to think further about this and work with you on how we can expand this relative to your question. I am happy to consider this.

Baroness Kidron: Thank you for that answer.

Q226 **The Chair:** I have a quick question for you, Markham Erickson. I want to follow up the earlier question about advertising. In its study looking at Covid misinformation sites that ran advertising, NewsGuard found that 67% of them had Google advertising tags on them. Many people see this as Google, through the back door, funding Covid disinformation. How does that comply with your platform policies?

Markham C. Erickson: It is absolutely against our platform policies. We take a tremendous amount of action to raise authoritative information about Covid. We work with local health authorities, including those in the UK—the National Health Service—on the location of vaccine information. As to that kind of content relative to our advertisements, we would not want our ads to be shown. We take down a tremendous amount of ads that are violative, but even one ad is not appropriate. I take your point.

The Chair: It is 67% of those sites. Why do you think the number is so high?

Markham C. Erickson: I am not familiar with that study or that figure, so I cannot comment on that.

The Chair: Perhaps we can follow it up in writing and get a response after this session.

Markham C. Erickson: I would be happy to do so.

Q227 **Darren Jones:** To begin with, I should declare that Google is a paying member of the Parliamentary Internet, Communications and Technology Forum, which is an all-party parliamentary group that I co-chair and has previously sponsored an event hosted by Labour Digital, which I chair.

My first question is for Leslie Miller. Ms Miller, will you be the person responsible for submitting the risk assessment on behalf of YouTube?

Leslie Miller: That is a good question. I do not know that I will be personally responsible, because it will be based on a cross-functional group of people participating in the follow-up of the various codes of practice that are going to be developed and the requirements that I know Ofcom will be responsible for. It will be a cross-functional team of people dedicated to assessing risk and providing information. I will certainly be involved, but I do not want to say that I will be the only executive who will be responsible for this.

Darren Jones: Where in the corporate governance in YouTube will that risk assessment get signed off?

Leslie Miller: It is still to be determined as we make our way through this process. I cannot give you any specifics, but it will certainly have a review by executives.

Darren Jones: Thank you. Mr Erickson, I put the same question to you. Will you be responsible for Google's submission of its risk assessment to Ofcom?

Markham C. Erickson: Thank you for the question. We have not made a determination about who will sign the risk assessment. We know that we will work with Ofcom in providing the risk assessment and it will be reviewed at an appropriate level.

Darren Jones: Have either of you reported to your Audit and Compliance Committee about online harms in your product areas?

Leslie Miller: If I understand your question as to whether I have reported in, no, I have not.

Darren Jones: Mr Erickson?

Markham C. Erickson: I am not sure I totally understand the question. If the question is whether we work with colleagues in trust and safety and legal and product to review how our products and services are being used and where we think they could be improved, yes, I and my team do weigh in in that regard.

Darren Jones: I understand that. To clarify my question, we are interested in board-level responsibility and accountability for the measures under this Bill. Obviously, Alphabet has its series of committees; I understand your products are in different lines under that. Alphabet has an Audit and Compliance Committee that is tasked with understanding risk to Google, including compliance with regulations and presumably, therefore, the safety of your customers using your products. My question is whether either of you, even before this Bill becomes law, has had to report in to that committee about online safety issues in your product areas.

Leslie Miller: Thus far I have not, but what I can tell you is that the programmes, tools and systems—everything that we put in place as it relates to the safety of our users—are not segmented into a non-executive part of the company. It is, quite honestly, comprehensive and holistic across the company.

Darren Jones: Mr Erickson, is that the same for Google?

Markham C. Erickson: Cross-functional teams report up so that our compliance committees are adequately informed. I appreciate the question.

Q228 **Darren Jones:** My next question is about research because, under a duty of care, there is an obligation to deal with harms that are reasonably foreseeable. Ms Miller, how much of YouTube’s marketing and research budget is spent on trying to identify online harm issues for your users?

Leslie Miller: I am not aware of a specific line item in our budget that is dedicated to this. We work with experts in fields such as child development and mental health. We rely on their insights and research because we want to make sure that our product and policy decisions are up to date, based on their latest research and insights.

Darren Jones: How do you identify safety issues on your product?

Leslie Miller: I think you can glean a lot as it relates to the transparency report, for example, on our community guidelines. We release that quarterly. In the last quarter, we removed over 6 million videos. For example, we look at how fast our machines are capturing that content relative to user flags. Over 90% of the content is flagged by our machines. It is also about how many views that content received before we took action on it and making sure we are reducing them—for example, that it is fewer than 10 views. In the last quarter, approximately two-thirds of violative content were removed with fewer than 10 views.

Q229 **Darren Jones:** The reason for my question is that there is a distinction between transparency reporting and operational standards that comply with your guidelines and terms and conditions, and funding research that identifies online safety issues being faced by your customers. So my next question is this: under this Bill, or in the debate surrounding it, Ofcom has the right to request information. There is a debate about whether external researchers should be given access to data that comes from the use of your products to help you understand the online harms that you might need to tackle. Would YouTube be willing to do that?

Leslie Miller: It depends on what Ofcom ends up defining as the type of material that needs to be made available—for example, material that could be used in nefarious ways to manipulate the platform. We would certainly be happy to work with Ofcom as it develops the requirements, and obviously we will comply, but this is where we would like to be involved over the process in the definitions and application of the codes.

Darren Jones: Mr Erickson, from a search perspective, how much of your research and marketing budget do you spend on researching potential safety issues?

Markham C. Erickson: Thank you for the question. If I may approach the question from this way, we have a number of ways to ensure that our search product is appropriately contemplating the safety of users and that when they are querying something they will get relevant information, not harmful information or spam that is inappropriate relative to what they are looking for.

In addition to internal trust and safety teams and mechanisms, we have a set of external search evaluators who constantly review our search tool. We prioritise things like health and safety and ask them to look at the results being returned. We make changes to our search tool when we see that we are not quite meeting our expectations to ensure that users are getting the most relevant information they are looking for. That is our primary means of ensuring that users are protected.

Darren Jones: Would Google be willing to share data with third-party researchers to help inform its understanding of harms from its products?

Markham C. Erickson: Google is very transparent about how its search product works. We publish material that explains how our search tool works and what its goals are, but I should also say that, unlike perhaps certain social media platforms, we are an open ecosystem. We are able to have educators, NGOs and researchers running search queries to determine whether there is bias in the results being returned. Therefore, we see a tremendous amount of third-party research being done because we are an open platform.

Darren Jones: You may disagree with my characterisation of your answers but, to play back what I have heard, neither of you knows which executive would be responsible for submitting the risk assessment; neither of you has reported to Alphabet's Audit and Compliance Committee on safety issues; and neither of you has been able to comment on research that you pay for to understand harms, or future harms, from the use of your products. Do either of you have concerns about Google's ability to comply with the provisions in this draft Bill?

Markham C. Erickson: As currently written, our concern with the draft Bill is that it is overly complicated; it is too complex. We would encourage the committee to tighten the definitions, particularly on things like online harm. We think that will ensure that we are able appropriately to comply with the law. The penalties under the proposed law are, as you know, quite severe. Not only do we have natural incentives to ensure that our products and services are being used as intended and are providing a safe and secure experience; of course, when the legislation is enacted, we will also have incentives to comply with the law based on civil penalty structures that are being proposed.

Darren Jones: Ms Miller, do you have anything to add?

Leslie Miller: I appreciate that is your interpretation, but I also think that part of the reason why we are having the conversation today is to see whether there are areas to modify the proposal to get a better understanding of the empowerment of the chief executive of Ofcom. If you suggest here that we know who exactly will oversee the risk assessment and the details of our compliance, the process is still ongoing, but I hope what you take away is that we are a committed partner in making sure that the Bill accomplishes its goals and is workable for us.

Darren Jones: More broadly, my point is that you should already have a board-level executive responsible for these issues, not just when the Bill becomes law. Ms Miller, on that point, I understand that you are going to become the exec chair next year at Google. Could you explain what that means to us, please?

Leslie Miller: I am not becoming the exec. Do you mean of GIFCT?

Darren Jones: No. We have been told that you are going to be the exec chair of Google.

Leslie Miller: That is news to me. That is quite a promotion. I will be the executive chair of the Global Internet Forum, which we co-founded several years ago to address violent extremism. I will be stepping in on that, which is equally important but not specific to Google.

Darren Jones: There may have been an error in our briefing notes. I will not therefore offer you congratulations.

Leslie Miller: That is okay.

Q230 **Lord Clement-Jones:** Hello to both of you. I want to pursue Darren Jones's point, in part about the issue of proactivity but also the question of responsibility. How proactive are you in preventing the impact and spread of harmful content on both Google and YouTube? In a sense, if you then adopt technology to mitigate that, who takes the decision—not only who develops it, but who takes the decision to introduce technology or, as you did with turning off Autoplay, who makes the decision not to deploy the technology? For instance, Leslie, who took the decision to turn off the Autoplay function, and what was the basis for that?

Leslie Miller: Ultimately, Susan Wojcicki is the CEO of YouTube and is responsible for the platform and business, but in the process of assessing Autoplay and whether it should be default on or default off, for example on YouTube Kids or supervised experiences, this is where we benefit from outside expertise as well as having experts inside the company, for example, on child development, digital literacy and mental health. We make determinations on the healthiest experience that a user will have on the platform, particularly in the area of kids.

In the 12 years I have worked at the company, I have found that nothing is static in the ways in which we commit to make sure that we adapt to the threats around us and that we put tools in place for parents and users

more broadly. For any of these things, I do not think we have made one decision and our work is done. It is ongoing, adaptive and happening all the time.

Lord Clement-Jones: For instance, have you evaluated what the impact of turning off Autoplay on YouTube Kids has been?

Leslie Miller: I am not directly aware whether we have done any assessment of impact in having Autoplay default off, but I would be happy to follow up with you on this.

Lord Clement-Jones: Turning to Google, Markham, we have the issue of searches producing self-harm material. We know that on Google Chrome there is a particular app with pop-up content that gives support to those looking for that kind of material, but it has not been mainstreamed. Is that the kind of thing that you are continuously evaluating? Is there a barrier to doing that sort of thing—tweaking the search algorithm, for instance?

Markham C. Erickson: Thank you for the question. I also thought that I was going to have to congratulate Leslie on her promotion, but she will be an excellent leader at GIFCT; it is well deserved in that regard.

When it comes to self-harm and threats of suicide ideation, we want to meet people where they are in those moments of distress. You referenced that there is a third-party plug-in to our Chrome browser that provides information for users who are exploring self-harm, but, to be clear, we do not wait for that, nor did we not take action ourselves. We work very closely with the Samaritans in the UK so that, when someone is searching for self-harm terms or suicide ideations, our service pops up a box with information to be helpful to that person in that moment, working with the experts at the Samaritans and Shout, so that they have information that can help them in that moment, including ways to reach someone quickly for help in that time of distress.

The plug-in that you referenced is in addition to that. We were happy to work with the party that developed that plug-in, but the Google service itself also responds to those types of threats with its own content that has been developed with expert input.

Lord Clement-Jones: Where would that evaluation and decision have been made to do what you have done?

Markham C. Erickson: When we develop our products and services, including improvements to our products and services, we think about how we can ensure that users have a safe and secure experience. In things like health and safety, we ask our search evaluators to pay particular attention to what can be done to ensure that we are meeting people in that moment with the information that will be most helpful to them, depending on what they are searching for.

Lord Clement-Jones: Coming back to Leslie, in the safety policy that you very comprehensively set out in your evidence, one of the key issues

is when a single piece of policy-violating content gets taken down. You probably have to hand the average time that it is allowed to stay up before it is removed. Who makes the decisions about that sort of area, if you like, of the way that you moderate the platform?

Leslie Miller: This is an area where we feel an obligation to be transparent, which is something I referred to earlier. You say that, by the time we have taken action on a video, probably too many people have been exposed to it and the harm has been done, but that is why in our transparency report we reference two things: first, that our machines can catch this content at scale around the globe; and, secondly, that we take action with very few views. In the last quarter alone, over two-thirds of the videos that violated our community guidelines were removed with fewer than 10 views, because we understand that just removing a video does not necessarily explain the totality of whether we have mitigated the harm. That is why we include those two other additional components.

Lord Knight of Weymouth: To pick up the last answer, what about the other third? When we are talking to these large platforms, it is easy for us to be quoted, "Oh well, we deal with 90% using our algorithm and only 10% are flagged by humans", but how many video views do you get every day on YouTube?

Leslie Miller: I am not sure about the number of video views on YouTube. Obviously, it depends on the user. We have 500 hours of video uploaded every minute, to give a bit of scale. On your question, this was why earlier this year we rolled out—it is a mouthful—the violative view rate because it goes to exactly what you are asking, which is this: that is okay, but relative to what people are viewing on the platform, how much of it is in fact violative? What we announced in the past quarter was that, for every 10,000 views on the platform, 19 to 21 of them were of content that violated our community guidelines.

Lord Knight of Weymouth: I go back to the scale issue. If you have billions of views every day, even if you manage to filter out all but 0.5%, that is still a lot of views we should be concerned about in terms of its violative content.

Leslie Miller: Yes, and this is why we make this data available and update it quarterly, because we want to keep doing better on it and want people to hold us accountable accordingly. It is also why we have applied machine learning because, to answer your point, being able to moderate content on the platform at scale requires a mix of machines and humans to make sure it is as safe as possible.

Q231 **Lord Knight of Weymouth:** Mr Erickson, I am interested in whether Google has evidence, through cookies or other sorts of technology for tracking, that people are finding conspiracy theories and misinformation elsewhere, perhaps on Facebook, which we have just been talking to, and then going to Google to search for further information based on what they might have seen on a social media platform elsewhere.

Markham C. Erickson: Thank you for the question. I want to be responsive to this and accurate but I am not sure I understand the question.

Lord Knight of Weymouth: What I am driving at is understanding the viral spread of misinformation and how easily you can end up down a rabbit hole. You start off on a platform; you get exposed to content that misinforms you; and you then go back to Google, because that is the search engine of choice for so many people, and search for that. Is there any way for you as Google to be able to flag the misinformation that is coming out and track those sorts of user journeys? Clearly, as a platform you are tracking user journeys all the time; that is your business.

Markham C. Erickson: Thank you for the question. As you may know, we actually announced our intention to eliminate cookies altogether from our platform, as well as any tracking device, because we understand that users do not want to be tracked across the internet. When users come to Google to search for something, we provide a result that is relevant and appropriate to the question that they are asking. The way we do that is by providing authoritative information relative to that search. That is the primary means of avoiding misinformation. People come to the search query. They want to know how tall a mountain is and we give them the answer quickly and efficiently.

Lord Knight of Weymouth: Does that extend to ads? I cannot remember the stat, but a fair proportion of users of your search engine struggle to differentiate between something that is a paid-for ranking using Google AdWords or whatever and something generated on the basis of your search algorithm. I appreciate that your search algorithm is looking for authoritative content, but how much do you worry about the ad taking people down that rabbit hole even further because it matches their search query and might be placed by someone who is not so authoritative as the people you are pointing to?

Markham C. Erickson: Thank you for the question. We do have very explicit advertising policies and guidelines with publishers that prevent ads that would perpetuate fraud or a scam. When it comes to sensitive information—sensitive health information or information about a sensitive event—we will not accept ads altogether with that but, rather, promote authoritative content in response to a query, for instance about Covid, to explain where one can find a vaccine or relevant information about the pandemic. So we have policies that are intended to prevent that kind of perpetuation of misinformation.

Lord Knight of Weymouth: In response to what you said about cookies, which is interesting, was the decision to try to move away from cookie because of worries about user journeys following rabbit holes of misinformation in the past?

Markham C. Erickson: Thank you for the question. The announcement to eliminate support of cookies that track a user across the internet was made because users do not want to be tracked across the internet.

Increasingly, that is the expectation, and we want to match the expectations of our users. We are transitioning altogether away from those kinds of tracking devices. We are doing it with some intentionality because there is an ecosystem of advertisers, publishers and small businesses that need to transition from cookie technology like that to alternative technology that is more privacy protective.

Lord Knight of Weymouth: Thank you. Finally, back to Leslie. The wonders of modern technology mean that I am told that 1.6 million hours of disinformation, hate speech and other banned content are viewed on YouTube every day. Is that a number you recognise?

Leslie Miller: Absolutely not. I am not sure where that number is from, and I am interested in the definition of disinformation. As for content that we prohibit, we move expeditiously to get it off the platform.

Lord Knight of Weymouth: Probably the first witness this committee heard from was a well-known footballer and commentator on football here called Rio Ferdinand. He raised with us how quickly IP content relating to, say, a football match that has a commercial arrangement attached to it is filtered out and taken down from a platform like YouTube, and yet racist comments and content are not. What is the answer to Rio Ferdinand as to why the technology works so well in respect of the commercial interests of intellectual property and yet does not work so well in respect of racism and hate-fuelled comment?

Leslie Miller: I respectfully see it differently. We prohibit hate and harassment content on the platform. In fact, we overhauled those policies in 2019 to make sure we were drawing the lines in a place where YouTube was inhospitable to this type of material. That is separate from our copyright tools where a rights holder can alert us to material for which that rights holder either decides he or she would like it to remain on the platform and a share in revenue from that content remaining on the platform, or submits a claim to remove the copyrighted material. Doing both is important for us in being a responsible platform, but we certainly do not prioritise IP over making sure that our platform is responsible.

Q232 **The Chair:** I have a few final questions for Leslie on YouTube. I believe YouTube says that it takes down 94% of content if it is in violation of its policies before it is reported. Is that correct?

Leslie Miller: It fluctuates a little bit quarter by quarter, but it tends to hover at around 90% of the content flagged by machines that we take action on versus user flags, for example.

The Chair: I want to be clear what that means. Does that mean that, of all the content you remove, 90% is identified?

Leslie Miller: It is significant.

The Chair: That does not mean your system is finding 90% of all problematic content on the platform.

Leslie Miller: Correct, yes. Of the content we have identified, over 90% of it—

The Chair: You have identified 90% of it and the other 10% is identified by users.

Leslie Miller: Yes.

The Chair: In terms of your AI, what proportion of hate speech, if you like, on YouTube do you believe is identified and removed?

Leslie Miller: I would have to go into the details of our latest quarterly transparency report. The numbers are publicly available. I am sorry; of the 6 million I referred to, I do not remember what the hard number and percentage is as they relate to hate, but I can easily get that. I apologise that I do not remember it off the top of my head.

The Chair: For content that is removed, particularly hate speech, do you also have details of the view time for that content? You have published data about the number of views as a proportion of the total number of views on the platform, but do you gather data about the view time?

Leslie Miller: I am not really sure what you mean by view time. I do not know whether you mean how long it is up before it is removed, which goes to the figures and why we care about sharing. Moving in a direction of having fewer than 10 views is obviously very important to us. On a quarterly basis, we hover around two-thirds to three-quarters in making sure we remove the content, but I may be misunderstanding your question or point.

The Chair: You are saying that two-thirds to three-quarters of content that is removed has 10 or fewer views. That leaves another third that has more than that. Do you say that this number of hours of problematic content—hate speech or other content—that breaches your platform policies is how much of it that is watched?

Leslie Miller: Yes.

The Chair: That will be based on the number of films and the total amount of time they are watched.

Leslie Miller: Yes.

The Chair: That is on the basis that one film that is an hour long and is watched 10 times is watched for 10 hours.

Leslie Miller: I am not aware of analysis based on something like that—for example, the hour of watch time—but you raise a great point about which we realised we needed to be more transparent, and that is the violative view rate. We also update that quarterly now. When we think about what is the amount of content that viewers are seeing that violates our content relative to the other content that does not, in the past

quarter for every 10,000 views on the platform, 19 to 21 were of content that violated our community guidelines.

The Chair: What I am interested in is that, for people who are exposed to content that breaches your community guidelines, how much content like that do they see? What we have here is a very big platform and most of the content and most people's experience of that content will largely be quite positive, but for some people their experience will be very negative and they probably see a disproportionately large amount of content like that. Therefore, do you do analysis of how much people see of the problematic content to which they are exposed?

Leslie Miller: I would welcome following that up with you. I am not aware that this is something we do analysis on, but we also think about making sure, for example, that people are not being radicalised and, based on a video they are watching, are not going down that radicalisation route, which is why we have also made changes to recommendations, but I welcome following up with you on the specifics you are asking about.

The Chair: If you are concerned about radicalisation, this is the sort of thing that you would be tracking and researching.

Leslie Miller: It is a matter of, first, if the content violates our community guidelines around violent extremism, hate and harassment, we are moving to get it off the platform as quickly as possible. However, content may not violate our community guidelines but may be perceived to be borderline. We have put a number of changes in place in our recommendations systems over the past couple of years whereby content that is determined to be borderline is demoted out of recommendations.

The Chair: Do you track how many times a film that has been removed from the platform was recommended, using the recommendation tools?

Leslie Miller: If it has been removed from the platform, I do not believe it would have been in recommendations, but I can follow that up.

The Chair: If a third of the films that are removed have been viewed more than 10 times, it suggests they have probably been on the platform for a while and have been recommended because they were not flagged for removal before that point.

Leslie Miller: I think that is a fair assumption. It could also be that a video may be somewhat newsworthy and ultimately violative, but it already received more than 10 views in a matter of minutes, for example.

The Chair: But it would be fair to say that something like the conspiracy film "Plandemic" about Covid, which was removed in the end, was viewed so many times that it would almost certainly have been recommended, using your recommendation tools.

Leslie Miller: I am just not certain of that, but I want to give you an honest answer, so let me take that back and follow up with you on it.

The Chair: It would seem likely that, for Autoplay and tools like NextUp, films like that were recommended for people who had an interest in that sort of content. A lot of the research we have received about all social media platforms that use AI recommendation systems suggests that the model is based on engagement. We have discussed this a lot with Facebook, obviously but I think this would be pretty similar for YouTube as well.

Leslie Miller: Again, I am not convinced by what you are saying because, in addition to demoting borderline content out of recommendations, in areas like Covid we make sure that we raise authoritative information. If somebody is coming to YouTube to search for something relative to Covid, we will make sure that the search results are from authoritative sources.

The Chair: For search results, sure, but a lot of people's experience is not based on search. I believe that, three years ago, YouTube said that 70% of plays on YouTube were selected by the platform for the user; they were not films that people had searched out for themselves. Is that figure still the same or has it changed in the past three years?

Leslie Miller: I do not know, but I would be happy to get you the answer to that.

The Chair: I am surprised you do not know. It is quite a fundamental question about how the platform works. If 70% of what people watch on YouTube is played for them, that is a pretty fundamental part of the business.

Leslie Miller: Again, I just do not have the answer for you at this moment, but I would be happy, like I said—

The Chair: Are you familiar with that figure?

Leslie Miller: To be honest, if it is being said that 70% of the watch time is because of either Autoplay or recommendations, I do not want to speak to that specific number just because I do not recall it. It does not mean we may not have said it, but I just need to follow up.

The Chair: I would welcome it if you could write back to us on that point as well.

Leslie Miller: Of course.

The Chair: I think it would be fair to assume that people who engage with conspiracy theories or harmful content probably have quite a lot of it played to them and recommended to them; they are not necessarily searching for it.

Leslie Miller: Of course. I definitely want to follow up with you on that.

The Chair: Okay. Thank you. I think that concludes the committee's questions. Thank you for your evidence.