PARLIAMENTARY DEBATES

HOUSE OF COMMONS
OFFICIAL REPORT
GENERAL COMMITTEES

Public Bill Committee

ONLINE SAFETY BILL

Second Sitting

Tuesday 24 May 2022

(Afternoon)

CONTENTS

No proofs can be supplied. Corrections that Members suggest for the final version of the report should be clearly marked in a copy of the report—not telephoned—and must be received in the Editor's Room, House of Commons,

not later than

Saturday 28 May 2022

The Committee consisted of the following Members:

Chairs: † Sir Roger Gale, Christina Rees

† Ansell, Caroline *(Eastbourne)* (Con)
† Bailey, Shaun *(West Bromwich West)* (Con)
† Blackman, Kirsty *(Aberdeen North)* (SNP)
Carden, Dan *(Liverpool, Walton)* (Lab)
† Davies-Jones, Alex *(Pontypridd)* (Lab)
† Double, Steve *(St Austell and Newquay)* (Con)
† Fletcher, Nick *(Don Valley)* (Con)
† Holden, Mr Richard *(North West Durham)* (Con)
† Keeley, Barbara *(Worsley and Eccles South)* (Lab)
† Leadbeater, Kim *(Batley and Spen)* (Lab)
† Miller, Mrs Maria *(Basingstoke)* (Con)
† Mishra, Navendu *(Stockport)* (Lab)
† Moore, Damien *(Southport)* (Con)
Nicolson, John *(Ochil and South Perthshire)* (SNP)
† Philp, Chris *(Parliamentary Under-Secretary of State for Digital, Culture, Media and Sport)*
† Russell, Dean *(Watford)* (Con)
Stevenson, Jane *(Wolverhampton North East)* (Con)

Katya Cassidy, Kevin Maddison, Seb Newman, *Committee Clerks*

† attended the Committee

Witnesses

Richard Earley, UK Public Policy Manager, Meta

Becky Foreman, UK Corporate Affairs Director, Microsoft

Katie O'Donovan, Director of Government Affairs and Public Policy, Google UK

Professor Clare McGlynn, Professor of Law, Durham University

Jessica Eagelton, Policy and Public Affairs Manager, Refuge

Janaya Walker, Public Affairs Manager, End Violence Against Women

Lulu Freemont, Head of Digital Regulation, techUK

Ian Stevenson, Chair, OSTIA

Adam Hildreth, CEO, Crisp

Jared Sine, Chief Business Affairs and Legal Officer, Match Group

Nima Elmi, Head of Public Policy in Europe, Bumble

Dr Rachel O'Connell, CEO TrustElevate

Rhiannon-Faye McDonald, Victim and Survivor Advocate, Marie Collins Foundation

Susie Hargreaves OBE, Chief Executive, Internet Watch Foundation

Ellen Judson, Lead Researcher at the Centre for the Analysis of Social Media, Demos

Kyle Taylor, Founder and Director, Fair Vote UK

Public Bill Committee

Tuesday 24 May 2022

(Afternoon)

[Sir Roger Gale *in the Chair*]

Online Safety Bill

2 pm

The Committee deliberated in private.

Examination of Witnesses

Richard Earley, Becky Foreman and Katie O'Donovan gave evidence.

2.2 pm

Q68 The Chair: Good afternoon. We are now sitting in public and these proceedings are being broadcast. This afternoon, we will first hear oral evidence from Richard Earley, the UK public policy manager of Meta, Becky Foreman, the UK corporate affairs director at Microsoft, and Katie O'Donovan, the director of Government affairs and public policy at Google and YouTube. Ladies and gentlemen, thank you very much indeed for joining us. For the sake of the record, could I just ask you to identify yourselves?

Richard Earley: Good afternoon. My name is Richard Earley, and I work in the public policy team at Meta, leading on content issues including the Online Safety Bill.

Becky Foreman: I am Becky Foreman; I am the corporate affairs director for Microsoft UK.

Katie O'Donovan: I am Katie O'Donovan; I am director of Government affairs and public policy for Google in the UK.

The Chair: May I just ask you, for the benefit of *Hansard*, to try to speak up a little? The sound system is not all that it might be in this room, and the acoustics certainly are not.

Q69 Alex Davies-Jones (Pontypridd) (Lab): Thank you to our witnesses for joining us this afternoon. Quite bluntly, I will get into it, because what is frustrating for us, as Parliamentarians, and for our constituents, is the fact that we need this legislation in the first place. Why are you, as platforms, allowing harmful and illegal content to perpetuate on your platforms? Why do we need this legislation for you to take action? It is within your gift to give, and despite all the things I am sure you are about to tell me that you are doing to prevent this issue from happening, it is happening and we are needing to legislate, so why?

The Chair: Mr Earley, I will go left to right to start with, if that is all right with you, so you have drawn the short straw.

Richard Earley: No worries, and thank you very much for giving us the opportunity to speak to you all today; I know that we do not have very much time. In short, we think this legislation is necessary because we believe that it is really important that democratically elected Members of Parliament and Government can provide input into the sorts of decisions that companies such as ours are making, every day, about how people use the internet. We do not believe that it is right for companies such as ours to be taking so many important decisions every single day.

Now, unfortunately, it is the case that social media reflects the society that we live in, so all of the problems that we see in our society also have a reflection on our services. Our priority, speaking for Meta and the services we provide—Facebook, Instagram and WhatsApp—is to do everything we can to make sure our users have as positive an experience as possible on our platform. That is why we have invested more than $13 billion over the past five years in safety and security, and have more than 40,000 people working at our company on safety and security every day.

That said, I fully recognise that we have a lot more areas to work on, and we are not waiting for this Bill to come into effect to do that. We recently launched a whole range of updated tools and technologies on Instagram, for example, to protect young people, including preventing anyone under the age of 18 from being messaged by a person they are not directly connected to. We are also using new technology to identify potentially suspicious accounts to prevent young people from appearing in any search results that those people carry out. We are trying to take steps to address these problems, but I accept there is a lot more to do.

Q70 Alex Davies-Jones: Before I bring in Becky and Katie to answer that, I just want to bring you back to something you said about social media and your platforms reflecting society like a mirror. That analogy is used time and again, but actually they are not a mirror. The platforms and the algorithms they use amplify, encourage and magnify certain types of content, so they are not a mirror of what we see in society. You do not see a balanced view of two points of an issue, for example.

You say that work is already being done to remove this content, but on Instagram, for example, which is a platform predominantly used by women, the Centre for Countering Digital Hate has exposed what they term an "epidemic of misogynistic abuse", with 90% of misogynistic abuse being sent via direct messaging. It is being ignored by the platform even when it is being reported to the moderators. Why is that happening?

Richard Earley: First, your point about algorithms is really important, but I do not agree that they are being used to promote harmful content. In fact, in our company, we use algorithms to do the reverse of that. We try to identify content that might break our policies—the ones we write with our global network of safety experts—and then remove those posts, or if we find images or posts that we think might be close to breaking those rules, we show them lower in people's feeds so that they have a lower likelihood of being seen. That is why, over the past two years, we have reduced the prevalence of harmful posts such as hate speech on Facebook so that now only 0.03% of views of posts on Facebook contain that kind of hate speech—we have almost halved the number. That is one type of action that we take in the public parts of social media.

When it comes to direct messages, including on Instagram, there are a range of steps that we take, including giving users additional tools to turn off any words they do not want to see in direct messages from anyone. We have recently rolled out a new feature called "restrict" which enables you to turn off any messages or comments from people who have just recently started to follow you, for example, and have just created their accounts. Those are some of the tools that we are trying to use to address that.

Q71 Alex Davies-Jones: So the responsibility is on the user, rather than the platform, to take action against abuse?

Richard Earley: No, the responsibility is absolutely shared by those of us who offer platforms, by those who are engaged in abuse in society, and by civil society and users more widely. We want to ensure we are doing everything we can to use the latest technology to stop abuse happening where we can and give people who use our services the power to control their experience and prevent themselves from encountering it.

The Chair: We must allow the other witnesses to participate.

Becky Foreman: Thank you for inviting me to give evidence to you today. Online safety is extremely important to Microsoft and sits right at the heart of everything we do. We have a "safety by design" policy, and responsibility for safety within our organisation sits right across the board, from engineers to operations and policy people. However, it is a complicated, difficult issue. We welcome and support the regulation that is being brought forward.

We have made a lot of investments in this area. For example, we introduced PhotoDNA more than 10 years ago, which is a tool that is used right across the sector and by non-governmental organisations to scan for child sexual abuse material and remove it from their platforms. More recently, we have introduced a grooming tool that automates the process of trying to establish whether there is a conversation for grooming taking place between an adult and a child. That can then be flagged for human review. We have made that available at no charge to the industry, and it has been licensed by a US NGO called Thorn. We take this really seriously, but it is a complicated issue and we really welcome the regulation and the opportunity to work with the Government and Ofcom on this.

Katie O'Donovan: Thank you so much for having me here today and asking us to give evidence. Thank you for your question. I have worked at Google and YouTube for about seven years and I am really proud of our progress on safety in those years. We think about it in three different ways. First, what products can we design and build to keep our users safer? Similar to Microsoft, we have developed technology that identifies new child sex abuse material and we have made that available across the industry. We have developed new policies and new ways of detecting content on YouTube, which means we have really strict community guidelines, we identify that content and we take it down. Those policies that underlie our products are really important. Finally, we work across education, both in secondary and primary schools, to help inform and educate children through our "Be Internet Legends" programme, which has reached about 4 million people.

There is definitely much more that we can do and I think the context of a regulatory environment is really important. We also welcome the Bill and I think it is really going to be meaningful when Ofcom audits how we are meeting the requirements in the legislation—not just how platforms like ours are meeting the requirements in the Bill, but a wide spectrum of platforms that young people and adults use. That could have a really positive additive effect to the impact.

It is worth pausing and reflecting on legislation that has passed recently, as well. The age-appropriate design code or the children's code that the Information Commissioner's Office now manages has also helped us determine new ways to keep our users safe. For example, where we have long had a product called SafeSearch, which you can use on search and parents can keep a lock on, we now also put that on by default where we use signals to identify people who we think are under 18.

We think that is getting the right balance between providing a safer environment but also enabling people to access information. We have not waited for this regulation. This regulation can help us do more, and it can also level the playing field and really make sure that everyone in the industry steps up and meets the best practice that can exist.

Q72 Alex Davies-Jones: Thank you, both, for adding context to that. If I can bring you back to what is not being done and why we need to legislate, Richard, I come back to you. You mentioned some of the tools and systems that you have put in place so users can stop abuse from happening. Why is it that that 90% of abuse on Instagram in direct messages is being ignored by your moderators?

Richard Earley: I do not accept that figure. I believe that if you look at our quarterly transparency report, which we just released last week, you can see that we find more than 90% of all the content that we remove for breaking our policies ourselves. Whenever somebody reports something on any of our platforms, they get a response from us. I think it is really important, as we are focusing on the Bill, to understand or make the point that, for private messaging, yes, there are different harms and different risks of harm that can apply, which is why the steps that we take differ from the steps that we take in public social media.

One of the things that we have noticed in the final draft of the Bill is that the original distinction between public social media and private messaging, which was contained in the online harms White Paper and in earlier drafts of the Bill, has been lost here. Acknowledging that distinction, and helping companies recognise that there is different risk and then different steps that can be taken in private messaging to what is taken on public social media, would be a really important thing for the Committee to consider.

Q73 Alex Davies-Jones: Quite briefly, because I know we are short on time, exactly how many human moderators do you have working to take down disinformation and harmful illegal content on your platforms?

Richard Earley: We have around 40,000 people in total working on safety and security globally and, of those, around half directly review posts and content.

Q74 Alex Davies-Jones: How many of those are directly employed by you and how many are third party?

Richard Earley: I do not have that figure myself but I know it is predominantly the case that, in terms of the safety functions that we perform, it is not just looking at the pieces of content; it is also designing the technology that finds and surfaces content itself. As I said, more than 90% of the time—more than 95% in most cases—it is our technology that finds and removes content before anyone has to look at it or report it.

Q75 Alex Davies-Jones: On that technology, we have been told that you are not doing enough to remove harmful and illegal content in minority languages. This is a massive gap. In London alone, more than 250 languages are spoken on a regular basis. How do you explain your inaction on this? Can you really claim that your platform is safe if you are not building and investing in AI systems in a range of languages? What proactive steps are you taking to address this extreme content that is not in English?

Richard Earley: That group of 40,000 people that I mentioned, they operate 24 hours, 7 days a week. They cover more than 70 languages between them, which includes the vast majority of the world's major spoken languages. I should say that people working at Meta, working on these classifiers and reviewing content, include people with native proficiency in these languages and people who can build the technology to find and remove things too. It is not just what happens within Meta that makes a difference here, but the work we do with our external partners. We have over 850 safety partners that we work with globally, who help us understand how different terms can be used and how different issues can affect the spread of harm on our platforms. All of that goes into informing both the policies we use to protect people on our platform and the technology we build to ensure those policies are followed.

Q76 Alex Davies-Jones: Finally, which UK organisations that you use have quality assured any of their moderator training materials?

Richard Earley: I am sorry, could you repeat the question?

Alex Davies-Jones: The vast majority of people are third party. They are not employed directly by Meta to moderate content, so how many of the UK organisations you use have been quality assured to ensure that the training they provide in order to spot this illegal and harmful content is taken on board?

Richard Earley: I do not believe it is correct that for our company, the majority of moderators are employed by—

Alex Davies-Jones: You do not have the figures, so you cannot tell me.

Richard Earley: I haven't, no, but I will be happy to let you know afterwards in our written submission. Everyone who is involved in reviewing content at Meta goes through an extremely lengthy training process that lasts multiple weeks, covering not just our community standards in total but also the specific area they are focusing on, such as violence and incitement. If it is hate speech, of course, there is a very important language component to that training, but in other areas—nudity or graphic violence—the language component is less important. We have published quite a lot about the work we do to make sure our moderators are as effective as possible and to continue auditing and training them. I would be really happy to share some of that information, if you want.

Q77 Alex Davies-Jones: But that is only for those employed directly by Meta.

Richard Earley: I will have to get back to you to confirm that, but I think it applies to everyone who reviews content for Meta, whether they are directly employed by Meta or through one of our outsourced-in persistent partners.

The Chair: Thank you very much. Don't worry, ladies; I am sure other colleagues will have questions that they wish to pursue. Dean Russell, please.

Q78 Dean Russell (Watford) (Con): Thank you, Chair. I guess this is for all three of you, but it is actually directed primarily at Richard—apologies. I do not mean to be rude—well, I am probably about to be rude.

One of the reasons why we are bringing in this Bill is that platforms such as Facebook—Meta, sorry—just have not fulfilled their moral obligations to protect children from harm. What commitment are you making within your organisation to align yourself to deliver on the requirements of the Bill?

To be frank, the track record up until now is appalling, and all I hear when in these witness sessions, including before Christmas on the Joint Committee, is that it is as though the big platforms think they are doing a good job—that they are all fine. They have spent billions of pounds and it is not going anywhere, so I want to know what practical measures you are going to be putting into place following this Bill coming into law.

Richard Earley: Of course, I do not accept that we have failed in our moral obligation to our users, particularly our younger users. That is the most important obligation that we have. I work with hundreds of people, and there are thousands of people at our company who spend every single day talking to individuals who have experienced abuse online, people who have lived experience of working with victims of abuse, and human rights defenders—including people in public life such as yourself—to understand the impact that the use of our platform can have, and work every day to make it better.

Q79 Dean Russell: But do you accept that there is a massive gap between those who you perhaps have been protecting and those who are not protected, hence the need for us to put this law in place?

Richard Earley: Again, we publish this transparency report every quarter, which is our attempt to show how we are doing at enforcing our rules. We publish how many of the posts that break our rules we take down ourselves, and also our estimates of how likely you are to find a piece of harmful content on the platform—as I mentioned, it is around three in every 10,000 for hate speech right now—but we fully recognise that you will not take our word for it. We expect confidence in that work to be earned, not just assumed.

That is why last year, we commissioned EY to carry out a fully independent audit of these systems. It published that report last week when we published our most recent transparency report and, again, I am very happy to share it with you here. The reason we have been calling for many years for pieces of legislation like this Bill to come into effect is that we think having Ofcom, the regulator—as my colleagues just said—able to look in more detail at the work we are doing, assess the work we are doing, and identify areas where we could do more is a really important part of what this Bill can do.

Q80 Dean Russell: I am conscious of the time, sorry. I know colleagues want to come in, but what are the practical measures? What will you be doing differently moving forward following this Bill?

Richard Earley: To start with, as I said, we are not waiting for the Bill. We are introducing new products and new changes all the time.

Q81 Dean Russell: Which will do what, sorry? I do not mean to be rude, but what will they be?

Richard Earley: Well, I just spoke about some of the changes we made regarding young people, including defaulting them into private accounts. We have launched additional tools making it possible for people to put in lists of words they do not want to see. Many of those changes are aligned with the core objectives of the Bill, which are about assessing early the risks of any new tools that we launch and looking all the time at how the use of technology changes and what new risks that might bring. It is then about taking proactive steps to try to reduce the risk of those harms.

Q82 Dean Russell: May I ask you a specific question? Will that include enabling bereaved parents to see their children's Facebook posts and profile?

Richard Earley: This is an issue we have discussed at length with DCMS, and we have consulted a number of people. It is, of course, one of the most sensitive, delicate and difficult issues we have to deal with, and we deal with those cases very regularly. In the process that exists at present, there are, of course, coronial powers. There is a process in the UK and other countries for coroners to request information.

When it comes to access for parents to individuals' accounts, at present we have a system for legacy contacts on some of our services, where you can nominate somebody to have access to your account after you pass away. We are looking at how that can be expanded. Unfortunately, there are an awful lot of different obligations we have to consider, not least the obligations to a person who used our services and then passed away, because their privacy rights continue after they have passed away too.

Dean Russell: Okay, so there is a compassion element. I am conscious of time, so I will stop there.

The Chair: One moment, please. I am conscious of the fact that we are going to run out of time. I am not prepared to allow witnesses to leave without feeling they have had a chance to say anything. Ms Foreman, Ms O'Donovan, is there anything you want to comment on from what you have heard so far? If you are happy, that is fine, I just want to make sure you are not being short-changed.

Becky Foreman: No.

Katie O'Donovan: No, I look forward to the next question.

Q83 Kirsty Blackman (Aberdeen North) (SNP): Given the size of Facebook, a lot of our questions will be focused towards it—not that you guys do not have very large platforms, but the risks with social media are larger. You mentioned, Richard, that three in every 10,000 views are hate speech. If three in every 10,000 things I said were hate speech, I would be arrested. Do you not think that, given the incredibly high number of views there are on Facebook, there is much more you need to do to reduce the amount of hate speech?

Richard Earley: So, reducing that number—the prevalence figure, as we call it—is the goal that we set our engineers and policy teams, and it is what we are devoted to doing. On whether it is a high number, I think we are quite rare among companies of our size in providing that level of transparency about how effective our systems are, and so to compare whether the amount is high or low, you would require additional transparency from other companies. That is why we really welcome the part of the Bill that allows Ofcom to set standards for what kinds of transparency actually are meaningful for people.

We have alighted on the figure of prevalence, because we think it is the best way for you and the public to hold us to account for how we are doing. As I said, that figure of three in every 10,000 has declined from six in every 10,000 about 12 months ago. I hope the figure continues to go down, but it is not just a matter of what we do on our platform. It is about how all of us in society function and the regulations you will all be creating to help support the work we do.

Q84 Kirsty Blackman: I would like to follow up with a question about responding to complaints. The complaints process is incredibly important. Reports need to be made and Facebook needs to respond to those reports. The Centre for Countering Digital Hate said that it put in 100 complaints and that 51 did not get any response from Facebook. It seems as though there is a systemic issue with a lack of response to complaints.

Richard Earley: I do not know the details of that methodological study. What I can tell you is that every time anyone reports something on Facebook or Instagram, they get a response into their support inbox. We do not put the response directly into your Messenger inbox or IG Direct inbox, because very often when people report something, they do not want to be reminded of what they have seen among messages from their friends and family. Unfortunately, sometimes people do not know about the support inbox and so they miss the response. That could be what happened there, but every time somebody reports something on one of our platforms, they get a response.

Q85 Kirsty Blackman: Does the response just say, "Thanks for your report"?

Richard Earley: No. I want to be very constructive here. I should say that some of the concerns that are raised around this date from several years ago. I will

accept that five or 10 years ago, the experience on our platforms was not this comprehensive, but in the last few years, we have really increased the transparency we give to people. When you submit something and report it for a particular violation, we give you a response that explains the action we took. If we removed it, we would explain what piece of our community standards it broke. It also gives you a link to see that section of our policy so you can understand it.

That is one way we have tried to increase the transparency we give to users. I think there is a lot more we could be doing. I could talk about some of the additional transparency steps we are taking around the way that our algorithms recommend content to people. Those are, again, all welcome parts of the Bill that we look forward to discussing further with Ofcom.

Q86 Kirsty Blackman: One of the things that has been recommended by a number of charities is increasing cross-platform and cross-company work to identify and take action on emerging threats. Do you think there would be the level of trust necessary for cross-platform co-operation with your competitors in the light of reports that, for example, Facebook employed somebody to put out negative things about TikTok in the US? Do you think that cross-platform working will work in that environment?

Richard Earley: Yes; I think it is already working, in fact. Others on the panel mentioned a few areas in which we have been collaborating in terms of open-sourcing some of the technology we have produced. A few years ago, we produced a grooming classifier—a technology that allows people to spot potentially inappropriate interactions between adults and children—and we open-sourced that and enabled it to be used and improved on by anyone else who is building a social media network.

A number of other areas, such as PhotoDNA, have already been mentioned. An obvious one is the Global Internet Forum to Counter Terrorism, which is a forum for sharing examples of known terrorist content so that those examples can be removed from across the internet. All those areas have been priorities for us in the past. A valuable piece of the Bill is that Ofcom—from what I can see from the reports that it has been asked to make—will do a lot of work to understand where there are further opportunities for collaboration among companies. We will be very keen to continue being involved in that.

Q87 Kirsty Blackman: I have a question for Katie on the algorithms that produce suggestions when you begin to type. It has been raised with me and in the evidence that we have received that when you begin to type, you might get a negative suggestion. If somebody types in, "Jews are", the algorithm might come up with some negative suggestions. What has Google done about that?

Katie O'Donovan: We are very clear that we want the auto-suggestion, as we call it, to be a helpful tool that helps you find the information that you are looking for quicker—that is the core rationale behind the search—but we really do not want it to perpetuate hate speech or harm for protected individuals or wider groups in society. We have changed the way that we use that auto-complete, and it will not auto-complete to harmful suggestions. That is a live process that we review and keep updated. Sometimes terminology, vernacular or slang change, or there is a topical focus on a particular group of people, so we keep it under review. But by our policy and implementation, those auto-suggestions should very much not be happening on Google search.

Q88 Kirsty Blackman: Would it be technically possible for all of the protected characteristics, for example, to have no auto-complete prompts come up?

Katie O'Donovan: That is an excellent question on where you do not want protections and safety to minimise user or individual impact. If you wanted a protected characteristic for Jewish people, for example, we see that as really important, and we should remove the ability for hate speech. If you wanted to do that for a Jewish cookbook, Jewish culture or Jewish history, and we removed everything, you would really minimise the amount of content that people could access.

The Bill is totally vital and will be incredibly significant on UK internet access, but that is where it is really important to get the balance and nuance right. Asking an algorithm to do something quite bluntly might look at first glance like it will really improve safety, but when you dig into it, you end up with the available information being much less sophisticated, less impactful and less full, which I think nobody really wants—either for the user or for those protected groups.

Q89 Kirsty Blackman: Would it not be easier to define all the protected characteristics and have a list of associated words than to define every possible instance of hate speech in relation to each?

Katie O'Donovan: The way we do it at the moment is through algorithmic learning. That is the most efficient way to do it because we have millions of different search terms, a large number of which we see for the very first time every single day on Google search. We rarely define things with static terms. We use our search rater guidelines—a guide of about 200 pages—to determine how those algorithms work and make sure that we have a dynamic ability to restrict them.

That means that you do not achieve perfection, and there will be changes and new topical uses that we perhaps did not anticipate—we make sure that we have enough data incoming to adjust to that. That is the most efficient way of doing it, and making sure that it has the nuance to stop the bad autocomplete but give access to the great content that we want people to get to.

The Chair: Thank you very much. Ms Foreman, do you want to add anything to that? You do not have to.

Becky Foreman: I do not have anything to add.

Q90 Barbara Keeley (Worsley and Eccles South) (Lab): I want to come back to transparency, which we touched on with my colleague Alex Davies-Jones earlier. Clearly, it is very important, and I think we could take a big step forward with the Bill. I want to ask you about child risk assessments, and whether they should be available publicly. I also want to ask about reports on the measures that you will have to take, as platforms, to manage the risks and mitigate the impact of harm. Harm is occurring at the moment—for example, content that causes harm is being left up. We heard earlier from the NSPCC that Facebook would not take down birthday groups for eight, nine and 10-year-old children, when it is known what purpose those birthday groups were serving for

those young children. I guess my question on transparency is, "Can't you do much better, and should there be public access to reports on the level of harm?"

Richard Earley: There are quite a few different questions there, and I will try to address them as briefly as I can. On the point about harmful Facebook groups, if a Facebook group is dedicated to breaking any of our rules, we can remove that group, even if no harmful content has been posted in it. I understand that was raised in the context of breadcrumbing, so trying to infer harmful intent from innocuous content. We have teams trying to understand how bad actors circumvent our rules, and to prevent them from doing that. That is a core part of our work, and a core part of what the Bill needs to incentivise us to do. That is why we have rules in place to remove groups that are dedicated to breaking our rules, even if no harmful content is actually posted in them.

On the question you asked about transparency, the Bill does an admirable job of trying to balance different types of transparency. There are some kinds of transparency that we believe are meaningful and valid to give to users. I gave the example a moment ago of explaining why a piece of content was removed and which of our community standards it broke. There is other transparency that we think is best given in a more general sense. We have our transparency report, as I said, where we give the figures for how much content we remove, how much of it we find ourselves—

Q91 Barbara Keeley: I am not talking here about general figures for what you have removed. I am talking about giving real access to the data on the risks of harm and the measures to mitigate harm. You could make those reports available to academics—we could find a way of doing that—and that would be very valuable. Surely what we want to do is to generate communities, including academics and people who have the aim of improving things, but you need to give them access to the data. You are the only ones who have access to the data, so it will just be you and Ofcom. A greater community out there who can help to improve things will not have that access.

Richard Earley: I completely agree. Apologies for hogging more time, but I think you have hit on an important point there, which is about sharing information with researchers. Last year, we gave data to support the publishing of more than 400 independent research projects, carried out along the lines you have described here. Just yesterday, we announced an expansion of what is called our Facebook open research tool, which expands academics' ability to access data about advertising.

Q92 Barbara Keeley: My question is, will you publish the risk assessment and the measures you are taking to mitigate?

Richard Earley: Going back to how the Bill works, when it comes to—

Barbara Keeley: No, I am not just asking about the Bill. Will you do that?

Richard Earley: We have not seen the Ofcom guidance on what those risk assessments should contain yet, so it is not possible to say. I think more transparency should always be the goal. If we can publish more information, we will do so.

Q93 Barbara Keeley: It would be good to have that goal. Can I come to you, Katie O'Donovan?

Katie O'Donovan: To begin with, I would pick up on the importance of transparency. We at Google and YouTube publish many reports on a quarterly or annual basis to help understand the actions we are taking. That ranges from everything on YouTube, where we publish by country the content we have taken down, why we have taken it down, how it was detected and the number of appeals. That is incredibly important information. It is good for researchers and others to have access to that.

We also do things around ads that we have removed and legal requests from different foreign Governments, which again has real validity. I think it is really important that Ofcom will have access to how we work through this—

Q94 Barbara Keeley: I was not just asking about Ofcom; I was wanting to go further than that and have wider access.

Katie O'Donovan: I do not want to gloss over the Ofcom point; I want to dwell on it for a second. In anticipation of this Bill, we were able to have conversations with Ofcom about how we work, the risks that we see and how our systems detect that. Hopefully, that is very helpful for Ofcom to understand how it will audit and regulate us, but it also informs how we need to think and improve our systems. I do think that is important.

We make a huge amount of training data available at Google. We publish a lot of shared APIs to help people understand what our data is doing. We are very open to publishing and working with academics.

It is difficult to give a broad statement without knowing the detail of what that data is. One thing I would say—it always sound a bit glib when people in my position say this—is that, in some cases, we do need to be limited in explaining exactly how our systems work to detect bad content. On YouTube, you have very clear community guidelines, which we know we have to publish, because people have a right to know what content is allowed and what is not, but we will find people who go right up to the line of that content very deliberately and carefully—they understand that, almost from a legal perspective. When it comes to fraudulent services and our ads, we have also seen people pivot the way that they attempt to defraud us. There needs to be some safe spaces to share that information. Ofcom is helpful for that too.

The Chair: Okay. Kim Leadbetter, one very quick question. We must move on—I am sorry.

Q95 Kim Leadbeater (Batley and Spen) (Lab): Okay, I will try to be very quick. The draft Bill contained a proposed new media literacy duty. That seems to have now disappeared. What are your digital media literacy strategies?

Becky Foreman: We have a range of strategies. One thing I would point to is research that we conduct every year and have done for a number of years called the digital civility index. It is a set of research that speaks to teens and adults in a number of countries around the world to understand what harms they are concerned about online and to ascertain whether those harms are increasing or decreasing and how they vary between different geographies. That is one way in which we are

trying to make more data and information available to the general public about the type of harms they might come across online and whether they are increasing or decreasing.

Richard Earley: We have a range of different organisations that we work with in the UK and internationally. One that I would like to draw attention to is the Economist Educational Foundation's Burnet News Club. We have supported them to increase their funding to be able to aim to reach 10% of all state schools with a really incredibly immersive and impressive programme that enables young people to understand digital literacy and digital numeracy and the media. We are also members of the media literacy taskforce of the Department for Digital, Culture, Media and Sport at the moment, which has been working to build on the strategy that the Government published.

Overall, there is a really important role for us as platforms to play here. We regularly commission and start new programmes in this space. What is also really important is to have more guidance from Government and civil society organisations that we work with on what is effective, so that we can know where we can put our resources and boost the greatest work.

Katie O'Donovan: Thank you for the question. It is really important. We were disappointed to see the literacy focus lost in the Bill.

We really take the issue seriously. We know there is an absolute responsibility for us when it comes to product, and an absolute responsibility when it comes to policy. Even within the safest products and with the most impressive and on-it parents, people can be exposed in content in ways that are surprising and shocking. That is why you need this holistic approach. We have long invested in a programme that we run with the non-governmental organisation Parent Zone called "Be internet legends". When we developed that, we did it with the PSHE Association to make sure it was totally compliant with the national curriculum. We regularly review that to check that it is actually making a difference. We did some recent research with MORI and got some really good results back.

We used to deliver that programme face to face in schools up and down the country. Obviously, the pandemic stopped that. We went online and while we did not enjoy it quite as much, we were able to reach real scale and it was really effective. Along with doing the assemblies, which are now back in person, we deliver a pack for teachers so they can also take that up at scale. We run similar programmes through YouTube with teenagers. It is absolutely incumbent on us to do more, but it must be part of the debate, because if you rely just on technological solutions, you will end up reducing access to lawful information, with some of the harms still being prevalent and people not having the skills to navigate them.

The Chair: I am sorry, but I must move on. Minister, I am afraid you only have five minutes.

Q96 The Parliamentary Under-Secretary of State for Digital, Culture, Media and Sport (Chris Philp): Welcome to the Committee's proceedings and thank you for joining us this afternoon. I would like to start on the question of the algorithmic promotion of content. Last week, I met with the Facebook whistleblower, Frances Haugen, who spoke in detail about she had found when working for Facebook, so I will start with you, Richard. On the question of transparency, which other Members of the Committee have touched on, would you have any objection to sharing all the information you hold internally with trusted researchers?

Richard Earley: What information are you referring to?

Chris Philp: Data, in particular on the operation of algorithmic promotion of particular kinds of content.

Richard Earley: We already do things like that through the direct opportunity that anyone has to see why a single post has been chosen for them in their feed. You can click on the three dots next to any post and see that. For researcher access and support, as I mentioned, we have contributed to the publishing of more than 400 reports over the last year, and we want to do more of that. In fact, the Bill requires Ofcom to conduct a report on how to unlock those sorts of barriers, which we think should be done as soon as possible. Yes, in general we support that sort of research.

I would like say one thing, though. I have worked at Facebook—now Meta—for almost five years, and nobody at Facebook has any obligation, any moral incentive, to do anything other than provide people with the best, most positive experience on our platform, because we know that if we do not give people a positive experience, through algorithms or anything else, they will leave our platform and will not use it. They tell us that and they do it, and the advertisers who pay for our services do not want to see that harmful content on our platforms either. All of our incentives are aligned with yours, which are to ensure that our users have a safe and positive experience on our platforms.

Q97 Chris Philp: Yet the algorithms that select particular content for promotion are optimised for user engagement —views, likes and shares—because that increases user stickiness and keeps them on the site for longer. The evidence seems to suggest that, despite what people say in response to the surveys you have just referenced, what they actually interact with the most—or what a particular proportion of the population chooses to interact with the most—is content that would be considered in some way extreme, divisive, or so on, and that the algorithms, which are optimised for user engagement, notice that and therefore uprank that content. Do you accept that your algorithms are optimised for user engagement?

Richard Earley: I am afraid to say that that is not correct. We have multiple algorithms on our services. Many of them, in fact, do the opposite of what you have just described: they identify posts that might be violent, misleading or harmful and reduce the prevalence of them within our feed products, our recommendation services and other parts of the service.

We optimise the algorithm that shows people things for something called meaningful social interaction. That is not just pure engagement; in fact, its focus—we made a large change to our algorithms in 2018 to focus on this—is on the kinds of activities online that research shows are correlated with positive wellbeing outcomes. Joining a group in your local area or deciding to go to

an event that was started by one of your friends—that is what our algorithms are designed to promote. In fact, when we made that switch in 2018, we saw a decrease in more than 50 million hours of Facebook use every day as a result of that change. That is not the action of a company that is just focused on maximising engagement; it is a company that is focused on giving our users a positive experience on our platform.

Q98 Chris Philp: You have alluded to some elements of the algorithmic landscape, but do you accept that the dominant feature of the algorithm that determines which content is most promoted is based on user engagement, and that the things you have described are essentially second-order modifications to that?

Richard Earley: No, because as I just said, when we sent the algorithm this instruction to focus on social interaction it actually decreased the amount of time people spent on our platform.

Q99 Chris Philp: It might have decreased it, but the meaningful social interaction score is, not exclusively, as you said, but principally based on user engagement, isn't it?

Richard Earley: As I said, it is about ensuring that people who spend time on our platform come away feeling that they have had a positive experience.

Q100 Chris Philp: That does not quite answer the question.

Richard Earley: I think that a really valuable part of the Bill that we are here to discuss is the fact that Ofcom will be required, and we in our risk assessments will be required, to consider the impact on the experience of our users of multiple different algorithms, of which we have hundreds. We build those algorithms to ensure that we reduce the prevalence of harmful content and give people the power to connect with those around them and build community. That is what we look forward to demonstrating to Ofcom when this legislation is in place.

Q101 Chris Philp: Yes, but in her testimony to, I think, the Joint Committee and the US Senate, in a document that she released to *The Wall Street Journal*, and in our conversation last week, Frances Haugen suggested that the culture inside Facebook, now Meta, is that measures that tend to reduce user engagement do not get a very sympathetic hearing internally. However, I think we are about to run out of time. I have one other question, which I will direct, again, to Richard. Forgive me, Katie and Becky, but it is probably most relevant for Meta.

Q102 The Chair: Just one moment, please. Is there anything that the other witnesses need to say about this before we move on? It will have to be very brief.

Katie O'Donovan: I welcome the opportunity to address the Committee. It is so important that this Bill has parliamentary scrutiny. It is a Bill that the DCMS has spent a lot of time on, getting it right and looking at the systems and the frameworks. However, it will lead to a fundamentally different internet for UK users versus the rest of the world. It is one of the most complicated Bills we are seeing anywhere in the world. I realise that it is very important to have scrutiny of us as platforms to determine what we are doing, but I think it is really important to also look at the substance of the Bill. If we have time, I would welcome the chance to give a little feedback on the substance of the Bill too.

Becky Foreman: I would add that the Committee spent a lot of time talking to Meta, who are obviously a big focus for the Bill, but it is important to remember that there are numerous other networks and services that potentially will be caught by the Bill and that are very different from Meta. It is important to remember that.

Chris Philp: While the Bill is proportionate in its measures, it is not designed to impose undue burdens on companies that are not high risk. I have one more question for Richard. I think Katie was saying that she wanted to make a statement?

The Chair: We are out of time. I am sorry about this; I regard it as woefully unsatisfactory. We have got three witnesses here, a lot of questions that need to be answered, and not enough time to do it. However, we have a raft of witnesses coming in for the rest of the day, so I am going to have to draw a line under this now. I am very grateful to you for taking the trouble to come—the Committee is indebted to you. You must have the opportunity to make your case. Would you be kind enough to put any comments that you wish to make in writing so that the Committee can have them. Feel free to go as broad as you would like because I feel very strongly that you have been short-changed this afternoon. We are indebted to you. Thank you very much indeed.

Richard Earley: We will certainly do that and look forward to providing comments in writing.

Examination of Witnesses

Professor Clare McGlynn, Jessica Eagelton and Janaya Walker gave evidence.

2.48 pm

The Chair: Good afternoon. We now hear oral evidence from Professor Clare McGlynn, professor of law at Durham University, Jessica Eagleton, policy and public affairs manager at Refuge, and Janaya Walker, public affairs manager at End Violence Against Women. Ladies, thank you very much for taking the trouble to join us this afternoon. We look forward to hearing from you.

Q103 Alex Davies-Jones: Thank you, Sir Roger, and thank you to the witnesses for joining us. We hear a lot about the negative experiences online of women, particularly women of colour. If violence against women and girls is not mentioned directly in the Bill, if misogyny is not made a priority harm, and if the violence against women and girls code of practice is not adopted in the Bill, what will that mean for the experience of women and girls?

Janaya Walker: Thank you for the opportunity to speak today. As you have addressed there, the real consensus among violence against women and girls organisations is for VAWG to be named in the Bill. The concern is that without that, the requirements that are placed on providers of regulated services will be very narrowly tied to the priority illegal content in schedule 7, as well as other illegal content.

We are very clear that violence against women and girls is part of a continuum in which there is a really broad manifestation of behaviour; some reaches a criminal threshold, but there is other behaviour that is important to be understood as part of the wider context. Much of the abuse that women and girls face cannot be understood by only looking through a criminal lens. We have to think about the relationship between the sender and the recipient—if it is an ex-partner, for example—the severity of the abuse they have experienced, the previous history and also the reach of the content. The worry is that the outcome of the Bill will be a missed opportunity in terms of addressing something that the Government have repeatedly committed to as a priority.

As you mentioned, we have worked with Refuge, Clare McGlynn, the NSPCC and 5Rights, bringing together our expertise to produce this full code of practice, which we think the Bill should be amended to include. The code of practice would introduce a cross-cutting duty that tries to mitigate this kind of pocketing of violence against women and girls into those three categories, to ensure that it is addressed really comprehensively.

Q104 Alex Davies-Jones: To what extent do you think that the provisions on anonymity will assist in reducing online violence against women and girls? Will the provisions currently in the Bill make a difference?

Janaya Walker: I think it will be limited. For the End Violence Against Women Coalition, our priority above all else is having a systems-based approach. Prevention really needs to be at the heart of the Bill. We need to think about the choices that platforms make in the design and operation of their services in order to prevent violence against women and girls in the first instance.

Anonymity has a place in the sense of providing users with agency, particularly in a context where a person is in danger and they could take that step in order to mitigate harm. There is a worry, though, when we look at things through an intersectional lens—thinking about how violence against women and girls intersects with other forms of harm, such as racism and homophobia. Lots of marginalised and minoritised people rely very heavy on being able to participate online anonymously, so we do not want to create a two-tier system whereby some people's safety is contingent on them being a verified user, which is one option available. We would like the focus to be much more on prevention in the first instance.

The Chair: Professor McGlynn and Ms Eagelton, you must feel free to come in if you wish to.

Q105 Alex Davies-Jones: My final question is probably directed at you, Professor McGlynn. Although we welcome the new communications offence of cyber-flashing, one of the criticisms is that it will not actually make a difference because of the onus on proving intent to cause harm, rather than the sender providing consent to receive the material. How do you respond to that?

Professor Clare McGlynn: I think it is great that the Government have recognised the harms of cyber-flashing and put that into the Bill. In the last couple of weeks we have had the case of Gaia Pope, a teenager who went missing and died—an inquest is currently taking place in Dorset. The case has raised the issue of the harms of cyber-flashing, because in the days before she went missing she was sent indecent images that triggered post-traumatic stress disorder from a previous rape. On the day she went missing, her aunt was trying to report that to the police, and one of the police officers was reported as saying that she was "taking the piss".

What I think that case highlights, interestingly, is that this girl was triggered by receiving these images, and it triggered a lot of adverse consequences. We do not know why that man sent her those images, and I guess my question would be: does it actually matter why he sent them? Unfortunately, the Bill says that why he sent them does matter, despite the harm it caused, because it would only be a criminal offence if it could be proved that he sent them with the intention of causing distress or for sexual gratification and being reckless about causing distress.

That has two main consequences. First, it is not comprehensive, so it does not cover all cases of cyber-flashing. The real risk is that a woman, having seen the headlines and heard the rhetoric about cyber-flashing being criminalised, might go to report it to the police but will then be told, "Actually, your case of cyber-flashing isn't criminal. Sorry." That might just undermine women's confidence in the criminal justice system even further.

Secondly, this threshold of having to prove the intention to cause distress is an evidential threshold, so even if you think, as might well be the case, that he sent the image to cause distress, you need the evidence to prove it. We know from the offence of non-consensual sending of sexual images that it is that threshold that limits prosecutions, but we are repeating that mistake here with this offence. So I think a consent-based, comprehensive, straightforward offence would send a stronger message and be a better message from which education could then take place.

The Chair: You are nodding, Ms Eagelton.

Jessica Eagelton: I agree with Professor McGlynn. Thinking about the broader landscape and intimate image abuse as well, I think there are some significant gaps. There is quite a piecemeal approach at the moment and issues that we are seeing in terms of enforcing measures on domestic abuse as well.

Q106 Mrs Maria Miller (Basingstoke) (Con): Thank you to all the panellists; it is incredibly helpful to have you here. The strength of the Bill will really be underpinned by the strength of the criminal law that underpins it, and schedule 7 lists offences that relate to sexual images, including revenge pornography, as priority offences. Can the witnesses say whether they think the law is sufficient to protect women from having their intimate pictures shared without their consent, or indeed whether the Bill will do anything to prevent the making and sharing of deepfake images? What would you like to see?

Professor Clare McGlynn: You make a very good point about how, in essence, criminal offences are now going to play a key part in the obligations of platforms under this Bill. In general, historically, the criminal law has not been a friend to women and girls. The criminal law was not written, designed or interpreted with women's harms in mind. That means that you have a very piecemeal,

confusing, out-of-date criminal law, particularly as regards online abuse, yet that is the basis on which we have to go forward. That is an unfortunate place for us to be, but I think we can strengthen it.

We could strengthen schedule 7 by, for example, including trafficking offences. There are tens of thousands of cases of trafficking, as we know from yourselves and whistleblowers, that platforms could be doing so much more about, but that is not a priority offence. The Obscene Publications Act distribution of unlawful images offence is not included. That means that incest porn, for example, is not a priority offence; it could be if we put obscene publications in that provision. Cyber-flashing, which again companies could take a lot of steps to act against, is not listed as a priority offence. Blackmail—sexual extortion, which has risen exponentially during the pandemic—again is not listed as a priority offence.

Deepfake pornography is a rising phenomenon. It is not an offence in English law to distribute deepfake pornography at the moment. That could be a very straightforward, simple change in the Bill. Only a few words are needed. It is very straightforward to make that a criminal offence, thanks to Scots law, where it is actually an offence to distribute altered images. The way the Bill is structured means the platforms will have to go by the highest standard, so in relation to deepfake porn, it would be interpreted as a priority harm—assuming that schedule 7 is actually altered to include all the Scottish offences, and the Northern Irish ones, which are absent at the moment.

The deepfake example points to a wider problem with the criminal law on online abuse: the laws vary considerably across the jurisdictions. There are very different laws on down-blousing, deepfake porn, intimate image abuse, extreme pornography, across all the different jurisdictions, so among the hundreds of lawyers the platforms are appointing, I hope they are appointing some Scots criminal lawyers, because that is where the highest standard tends to be.

Q107 Mrs Miller: Would the other panellists like to comment on this?

Jessica Eagelton: I think something that will particularly help in this instance is having that broad code of practice; that is a really important addition that must be made to the Bill. Refuge is the largest specialist provider of gender-based violence services in the country. We have a specialist tech abuse team who specialise in technology-facilitated domestic abuse, and what they have seen is that, pretty consistently, survivors are being let down by the platforms. They wait weeks and weeks for responses—months sometimes—if they get a response at all, and the reporting systems are just not up to scratch.

I think it will help to have the broad code of practice that Janaya mentioned. We collaborated with others to produce a workable example of what that could look like, for Ofcom to hopefully take as a starting point if it is mandated in the Bill. That sets out steps to improve the victim journey through content reporting, for example. Hopefully, via the code of practice, a victim of deepfakes and other forms of intimate image abuse would be able to have a more streamlined, better response from platforms.

I would also like to say, just touching on the point about schedule 7, that from the point of view of domestic abuse, there is another significant gap in that: controlling and coercive behaviour is not listed, but it should be. Controlling and coercive behaviour is one of the most common forms of domestic abuse. It carries serious risk; it is one of the key aggravating factors for domestic homicide, and we are seeing countless examples of that online, so we think that is another gap in schedule 7.

The Chair: Ms Walker?

Janaya Walker: Some of these discussions almost reiterate what I was saying earlier about the problematic nature of this, in that so much of what companies are going to be directed to do will be tied only to the specific schedule 7 offences. There have been lots of discussions about how you respond to some harms that reach a threshold of criminality and others that do not, but that really contrasts with the best practice approach to addressing violence against women and girls, which is really trying to understand the context and all of the ways that it manifests. There is a real worry among violence against women and girls organisations about the minimal response to content that is harmful to adults and children, but will not require taking such a rigorous approach.

Having the definition of violence against women and girls on the face of the Bill allows us to retain those expectations on providers as technology changes and new forms of abuse emerge, because the definition is there. It is VAWG as a whole that we are expecting the companies to address, rather than a changing list of offences that may or may not be captured in criminal law.

Q108 Kirsty Blackman: Why is it important that we have this? Is this a big thing? What are you guys actually seeing here?

Jessica Eagelton: I can respond to that in terms of what we are seeing as a provider. Technology-facilitated domestic abuse is an increasing form of domestic abuse: technology is providing perpetrators with increasing ways to abuse and harass survivors. What we are seeing on social media is constant abuse, harassment, intimate image abuse, monitoring and hacking of accounts, but when it comes to the responses we are getting from platforms at the moment, while I acknowledge that there is some good practice, the majority experience of survivors is that platforms are not responding sufficiently to the tech abuse they are experiencing.

Our concern is that the Bill could be a really good opportunity for survivors of domestic abuse to have greater protections online that would mean that they are not forced to come offline. At the moment, some of the options being given to survivors are to block the perpetrator—which in some cases has a minimal impact when they can easily set up new fake accounts—or to come offline completely. First, that is not a solution to that person being able to maintain contact, stay online and take part in public debate. But secondly, it can actually escalate risk in some cases, because a perpetrator could resort to in-person forms of abuse. If we do not make some of these changes—I am thinking in particular about mandating a VAWG code of practice, and looking at schedule 7 and including controlling and coercive behaviour—the Bill is going to be a missed opportunity. Women and survivors have been waiting long enough, and we need to take this opportunity.

Janaya Walker: If I could add to that, as Jessica has highlighted, there is the direct harm to survivors in terms of the really distressing experience of being exposed to these forms of harm, or the harm they experience offline being exacerbated online, but this is also about indirect harm. We need to think about the ways in which the choices that companies are making are having an impact on the extent to which violence against women and girls is allowed to flourish.

As Jessica said, it impacts our ability to participate in online discourse, because we often see a mirroring online of what happens offline, in the sense that the onus is often on women to take responsibility for keeping themselves safe. That is the status quo we see offline, in terms of the decisions we make about what we are told to wear or where we should go as a response to violence against women and girls. Similarly, online, the onus is often on us to come offline or put our profiles on private, to take all those actions, or to follow up with complaints to various different companies that are not taking action. There is also something about the wider impact on society as a whole by not addressing this within the Bill.

Q109 Kirsty Blackman: How does the proposed code of practice—or, I suppose, how could the Bill—tackle intersectionality of harms?

Janaya Walker: This is a really important question. We often highlight the fact that, as I have said, violence against women and girls often intersects with other forms of discrimination. For example, we know from research that EVAW conducted with Glitch during the pandemic that black and minoritised women and non-binary people experience a higher proportion of abuse. Similarly, research done by Amnesty International shows that black women experience harassment at a rate 84% higher than that experienced by their white counterparts. It is a real focal point. When we think about the abuse experienced, we see the ways that people's identities are impacted and how structural discrimination emerges online.

What we have done with the code of practice is try to introduce requirements for the companies to think about things through that lens, so having an overarching human rights and equalities framework and having the Equality Act protected characteristics named as a minimum. We see in the Bill quite vague language when it comes to intersectionality; it talks about people being members of a certain group. We do not have confidence that these companies, which are not famed for their diversity, will interpret that in a way that we regard as robust—thinking very clearly about protected characteristics, human rights and equalities legislation. The vagueness in the Bill is quite concerning. The code of practice is an attempt to be more directive on what we want to see and how to think through issues in a way that considers all survivors, all women and girls.

Professor Clare McGlynn: I wholly agree. The code of practice is one way by which we can explain in detail those sorts of intersecting harms and what companies and platforms should do, but I think it is vital that we also write it into the Bill. For example, on the definitions around certain characteristics and certain groups, in previous iterations reference was made to protected characteristics. I know certain groups can go wider than that, but naming those protected characteristics is really important, so that they are front and centre and the platforms know that that is exactly what they have to cover. That will cover all the bases and ensure that that happens.

Kirsty Blackman: I have a quite specific question on something that is a bit tangential.

The Chair: Last one, please.

Q110 Kirsty Blackman: If someone has consented to take part in pornography and they later change their mind and would like it to be taken down, do you think they should have the right to ask a porn website, for example, to take it down?

Professor Clare McGlynn: That is quite challenging not only for pornography platforms but for sex workers, in that if you could participate in pornography but at any time thereafter withdraw your consent, it is difficult to understand how a pornography company and the sex worker would be able to make a significant amount of money. The company would be reluctant to invest because it might have to withdraw the material at any time. In my view, that is a quite a challenge. I would not go down that route, because what it highlights is that the industry can be exploitative and that is where the concern comes from. I think there are other ways to deal with an exploitative porn industry and other ways to ensure that the material online has the full consent of participants. You could put some of those provisions into the Bill—for example, making the porn companies verify the age and consent of those who are participating in the videos for them to be uploaded. I think that is a better way to deal with that, and it would ensure that sex workers themselves can still contract to perform in porn and sustain their way of life.

Q111 Kim Leadbeater: Thank you very much—this is extremely interesting and helpful. You have covered a lot of ground already, but I wonder whether there is anything specific you think the Bill should be doing more about, to protect girls—under-18s or under-16s—in particular?

Janaya Walker: A lot of what we have discussed in terms of naming violence against women and girls on the face of the Bill includes children. We know that four in five offences of sexual communications with a child involved girls, and a lot of child abuse material is targeted at girls specifically. The Bill as a whole takes a very gender-neutral approach, which we do not think is helpful; in fact, we think it is quite harmful to trying to reduce the harm that girls face online.

This goes against the approach taken in the Home Office violence against women and girls strategy and its domestic abuse plan, as well as the gold-standard treaties the UK has signed up to, such as the Istanbul convention, which we signed and have recently committed to ratifying. The convention states explicitly that domestic laws, including on violence against women and girls online, need to take a very gendered approach. Currently, it is almost implied, with references to specific characteristics. We think that in addressing the abuse that girls, specifically, experience, we need to name girls. To clarify, the words "women", "girls", "gender" and "sex" do not appear in the Bill, and that is a problem.

Jessica Eagelton: May I add a point that is slightly broader than your question? Another thing that the Bill does not do at the moment is provide for specialist victim support for girls who are experiencing online abuse. There has been some discussion about taking a "polluter pays" approach; where platforms are not compliant with the duties, for example, a percentage of the funds that go to the regulator could go towards victim support services, such as the revenge porn helpline and Refuge's tech abuse team, that provide support to victims of abuse later on.

Professor Clare McGlynn: I can speak to pornography. Do you want to cover that separately, or shall I do that now?

Kim Leadbeater: That is fine.

Professor Clare McGlynn: I know that there was a discussion this morning about age assurance, which obviously targets children's access to pornography. I would emphasise that age assurance is not a panacea for the problems with pornography. We are so worried about age assurance only because of the content that is available online. The pornography industry is quite happy with age verification measures. It is a win-win for them: they get public credibility by saying they will adopt it; they can monetise it, because they are going to get more data—especially if they are encouraged to develop age verification measures, which of course they have been; that really is putting the fox in charge of the henhouse—and they know that it will be easily evaded.

One of the most recent surveys of young people in the UK was of 16 and 17-year-olds: 50% of them had used a VPN, which avoids age verification controls, and 25% more knew about that, so 75% of those older children knew how to evade age assurance. This is why the companies are quite happy—they are going to make money. It will stop some people stumbling across it, but it will not stop most older children accessing pornography. We need to focus on the content, and when we do that, we have to go beyond age assurance.

You have just heard Google talking about how it takes safety very seriously. Rape porn and incest porn are one click away on Google. They are freely and easily accessible. There are swathes of that material on Google. Twitter is hiding in plain sight, too. I know that you had a discussion about Twitter this morning. I, like many, thought, "Yes, I know there is porn on Twitter," but I must confess that until doing some prep over the last few weeks, I did not know the nature of that porn. For example, "Kidnapped in the wood"; "Daddy's little girl comes home from school; let's now cheer her up"; "Raped behind the bin"—this is the material that is on Twitter. We know there is a problem with Pornhub, but this is what is on Twitter as well.

As the Minister mentioned this morning, Twitter says you have to be 13, and you have to be 18 to try to access much of this content, but you just put in whatever date of birth is necessary—it is that easy—and you can get all this material. It is freely and easily accessible. Those companies are hiding in plain sight in that sense. The age verification and age assurance provisions, and the safety duties, need to be toughened up.

To an extent, I think this will come down to the regulator. Is the regulator going to accept Google's SafeSearch as satisfying the safety duties? I am not convinced, because of the easy accessibility of the rape and incest porn I have just talked about. I emphasise that incest porn is not classed as extreme pornography, so it is not a priority offence, but there are swathes of that material on Pornhub as well. In one of the studies that I did, we found that one in eight titles on the mainstream pornography sites described sexually violent material, and the incest material was the highest category in that. There is a lot of that around.

Q112 Barbara Keeley: We are talking here about pornography when it is hosted on mainstream websites, as opposed to pornographic websites. Could I ask you to confirm what more, specifically, you think the Bill should do to tackle pornography on mainstream websites, as you have just been describing with Twitter? What should the Bill be doing here?

Professor Clare McGlynn: In many ways, it is going to be up to the regulator. Is the regulator going to deem that things such as SafeSearch, or Twitter's current rules about sensitive information—which rely on the host to identify their material as sensitive—satisfy their obligations to minimise and mitigate the risk? That is, in essence, what it will all come down to.

Are they going to take the terms and conditions of Twitter, for example, at face value? Twitter's terms and conditions do say that they do not want sexually violent material on there, and they even say that it is because they know it glorifies violence against women and girls, but this material is there and does not appear to get swiftly and easily taken down. Even when you try to block it—I tried to block some cartoon child sexual abuse images, which are easily available on there; you do not have to search for them very hard, it literally comes up when you search for porn—it brings you up five or six other options in case you want to report them as well, so you are viewing them as well. Just on the cartoon child sexual abuse images, before anyone asks, they are very clever, because they are just under the radar of what is actually a prohibited offence.

It is not necessarily that there is more that the Bill itself could do, although the code of practice would ensure that they have to think about these things more. They have to report on their transparency and their risk assessments: for example, what type of content are they taking down? Who is making the reports, and how many are they upholding? But it is then on the regulator as to what they are going to accept as acceptable, frankly.

Barbara Keeley: Do any other panellists want to add to that?

Janaya Walker: Just to draw together the questions about pornography and the question you asked about children, I wanted to highlight one of the things that came up earlier, which was the importance of media literacy. We share the view that that has been rolled back from earlier versions of the draft Bill.

There has also been a shift, in that the emphasis of the draft Bill was also talking about the impact of harm. That is really important when we are talking about violence against women and girls, and what is happening in the context of schools and relationship and sex education. Where some of these things like non-consensual image sharing take place, the Bill as

currently drafted talks about media literacy and safe use of the service, rather than the impact of such material and really trying to point to the collective responsibility that everyone has as good digital citizens—in the language of Glitch—in terms of talking about online violence against women and girls. That is an area in which the Bill could be strengthened from the way it is currently drafted.

Jessica Eagelton: I completely agree with the media literacy point. In general, we see very low awareness of what tech abuse is. We surveyed some survivors and did some research last year—a public survey—and almost half of survivors told no one about the abuse they experienced online at the hands of their partner or former partner, and many of the survivors we interviewed did not understand what it was until they had come to Refuge and we had provided them with support. There is an aspect of that to the broader media literacy point as well: increasing awareness of what is and is not unacceptable behaviour online, and encouraging members of the public to report that and call it out when they see it.

Q113 Barbara Keeley: Thank you. Can I ask for a bit more detail on a question that you touched on earlier with my colleague Kirsty Blackman? It is to Professor McGlynn, really. I think you included in your written evidence to the Committee a point about using age and consent verification for pornography sites for people featured in the content of the site—not the age verification assurance checks on the sites, but for the content. Could I just draw out from you whether that is feasible, and would it be retrospective for all videos, or just new ones? How would that work?

Professor Clare McGlynn: Inevitably, it would have to work from any time that that requirement was put in place, in reality. That measure is being discussed in the Canadian Parliament at the moment—you might know that Pornhub's parent company, MindGeek, is based in Canada, which is why they are doing a lot of work in that regard. The provision was also put forward by the European Parliament in its debates on the Digital Services Act. Of course, any of these measures are possible; we could put it into the Bill that that will be a requirement.

Another way of doing it, of course, would be for the regulator to say that one of the ways in which Pornhub, for example—or XVideos or xHamster—should ensure that they are fulfilling their safety duties is by ensuring the age and consent of those for whom videos are uploaded. The flipside of that is that we could also introduce an offence for uploading a video and falsely representing that the person in the video had given their consent to that. That would mirror offences in the Fraud Act 2006.

The idea is really about introducing some element of friction so that there is a break before images are uploaded. For example, with intimate image abuse, which we have already talked about, the revenge porn helpline reports that for over half of the cases of such abuse that it deals with, the images go on to porn websites. So those aspects are really important. It is not just about all porn videos; it is also about trying to reduce the distribution of non-consensual videos.

Q114 Nick Fletcher (Don Valley) (Con): I think that it would have been better to hear from you three before we heard from the platforms this morning. Unfortunately, you have opened my eyes to a few things that I wish I did not have to know about—I think we all feel the same.

I am concerned about VPNs. Will the Bill stop anyone accessing through VPNs? Is there anything we can do about that? I googled "VPNs" to find out what they were, and apparently there is a genuine need for them when using public networks, because it is safer. Costa Coffee suggests that people do so, for example. I do not know how we could work that.

You have obviously educated me, and probably some of my colleagues, about some of the sites that are available. I do not mix in circles where I would be exposed to that, but obviously children and young people do and there is no filter. If I did know about those things, I would probably not speak to my colleagues about it, because that would probably not be a good thing to do, but younger people might think it is quite funny to talk about. Do you think there is an education piece there for schools and parents? Should these platforms be saying to them, "Look, this is out there, even though you might not have heard of it—some MPs have not heard of it." We ought to be doing something to protect children by telling parents what to look out for. Could there be something in the Bill to force them to do that? Do you think that would be a good idea? There is an awful lot there to answer—sorry.

Professor Clare McGlynn: On VPNs, I guess it is like so much technology: obviously it can be used for good, but it can also be used to evade regulations. My understanding is that individuals will be able to use a VPN to avoid age verification. On that point, I emphasise that in recent years Pornhub, at the same time as it was talking to the Government about developing age verification, was developing its own VPN app. At the same time it was saying, "Of course we will comply with your age verification rules."

Don't get me wrong: the age assurance provisions are important, because they will stop people stumbling across material, which is particularly important for the very youngest. In reality, 75% know about VPNs now, but once it becomes more widely known that this is how to evade it, I expect that all younger people will know how to do so. I do not think there is anything else you can do in the Bill, because you are not going to outlaw VPNs, for the reasons you identified—they are actually really important in some ways.

That is why the focus needs to be on content, because that is what we are actually concerned about. When you talk about media literacy and understanding, you are absolutely right, because we need to do more to educate all people, including young people—it does not just stop at age 18—about the nature of the pornography and the impact it can have. I guess that goes to the point about media literacy as well. It does also go to the point about fully and expertly resourcing sex and relationships education in school. Pornhub has its own sex education arm, but it is not the sex education arm that I think many of us would want to be encouraging. We need to be doing more in that regard.

Q115 Nick Fletcher: This might sound like a silly question. Can we not just put age verification on VPN sites, so that you can only have VPN access if you have gone through age verification? Do you understand what I am saying?

Professor Clare McGlynn: I do. We are beginning to reach the limits of my technical knowledge.

Nick Fletcher: You have gone beyond mine anyway.

Professor Clare McGlynn: You might be able to do that through regulations on your phone. If you have a phone that is age-protected, you might not be able to download a particular VPN app, perhaps. Maybe you could do that, but people would find ways to evade that requirement as well. We have to tackle the content. That is why you need to tackle Google and Twitter as well as the likes of Pornhub.

Nick Fletcher: Can we have them back in, Sir Roger?

The Chair: Minister?

Q116 Chris Philp: Thank you, Sir Roger, and thank you to the witnesses for coming in and giving very clear, helpful and powerful evidence to the Committee this afternoon. On the question of age verification or age assurance that we have just spoken about, clause 11(14) of the Bill sets a standard in the legislation that will be translated into the codes of practice by Ofcom. It says that, for the purposes of the subsection before on whether or not children can access a particular set of content, a platform is

"only entitled to conclude that it is not possible for children to access a service…if there are systems or processes in place…that achieve the result that children are not normally able to access the service".

Ofcom will then interpret in codes of practice what that means practically. Professor McGlynn, do you think that standard set out there—

"the result that children are not normally able to access the service or that part of it"

—is sufficiently high to address the concerns we have been discussing in the last few minutes?

Professor Clare McGlynn: At the moment, the wording with regard to age assurance in part 5—the pornography providers—is slightly different, compared with the other safety duties. That is one technicality that could be amended. As for whether the provision you just talked about is sufficient, in truth I think it comes down, in the end, to exactly what is required, and of course we do not yet know what the nature of the age verification or age assurance requirements will actually be and what that will actually mean.

I do not know what that will actually mean for something like Twitter. What will they have to do to change it? In principle, that terminology is possibly sufficient, but it kind of depends in practice what it actually means in terms of those codes of practice. We do not yet know what it means, because all we have in the Bill is about age assurance or age verification.

Q117 Chris Philp: Yes, you are quite right that the Ofcom codes of practice will be important. As far as I can see, the difference between clauses 68 and 11(14) is that one uses the word "access" and the other uses the word "encounter". Is that your analysis of the difference as well?

Professor Clare McGlynn: My understanding as well is that those terms are, at the moment, being interpreted slightly differently in terms of the requirements that people will be under. I am just making a point about it probably being easier to harmonise those terms.

Q118 Chris Philp: Thank you very much. I wanted to ask you a different question—one that has not come up so far in this session but has been raised quite frequently in the media. It concerns freedom of speech. This is probably for Professor McGlynn again. I am asking you this in your capacity as a professor of law. Some commentators have suggested that the Bill will have an adverse impact on freedom of speech. I do not agree with that. I have written an article in *The Times* today making that case, but what is your expert legal analysis of that question?

Professor Clare McGlynn: I read your piece in *The Times* this morning, which was a robust defence of the legislation, in that it said that it is no threat to freedom of speech, but I hope you read my quote tweet, in which I emphasised that there is a strong case to be made for regulation to free the speech of many others, including women and girls and other marginalised people. For example, the current lack of regulation means that women's freedom of speech is restricted because we fear going online because of the abuse we might encounter. Regulation frees speech, while your Bill does not unduly limit freedom of speech.

Q119 Chris Philp: Okay, I take your second point, but did you agree with the point that the Bill as crafted does not restrict what you would ordinarily consider to be free speech?

Professor Clare McGlynn: There are many ways in which speech is regulated. The social media companies already make choices about what speech is online and offline. There are strengths in the Bill, such as the ability to challenge when material is taken offline, because that can impact on women and girls as well. They might want to put forward a story about their experiences of abuse, for example. If that gets taken down, they will want to raise a complaint and have it swiftly dealt with, not just left in an inbox.

There are lots of ways in which speech is regulated, and the idea of having a binary choice between free speech and no free speech is inappropriate. Free speech is always regulated, and it is about how we choose to regulate it. I would keep making the point that the speech of women and girls and other marginalised people is minimised at the moment, so we need regulation to free it. The House of Lords and various other reports about free speech and regulation, for example, around extreme pornography, talk about regulation as being human-rights-enhancing. That is the approach we need to take.

The Chair: Thank you very much indeed. Once again, I am afraid I have to draw the session to a close, and once again we have probably not covered all the ground we would have liked. Professor McGlynn, Ms Walker, Ms Eagleton, thank you very much indeed. As always, if you have further thoughts or comments, please put them in writing and let us know. We are indebted to you.

Examination of Witnesses

Lulu Freemont, Ian Stevenson and Adam Hildreth gave evidence.

3.32 pm

The Chair: We will now hear oral evidence from Lulu Freemont, head of digital regulation at techUK; Ian Stevenson, the chairman of OSTIA; and Adam Hildreth,

[The Chair]

chief executive officer of Crisp, who is appearing by Zoom—and it works. Thank you all for joining us. I will not waste further time by asking you to identify yourselves, because I have effectively done that for you. Without further ado, I call Alex Davies-Jones.

Q120 Alex Davies-Jones: Thank you, Sir Roger; thank you, witnesses. We want the UK to become a world leader in tech start-ups. We want those employment opportunities for the future. Does this legislation, as it currently stands, threaten that ability?

Lulu Freemont: Hi everybody. Thank you so much for inviting techUK to give evidence today. Just to give a small intro to techUK, so that you know the perspective I am coming from, we are the trade body for the tech sector. We have roughly 850 tech companies in our membership, the majority of which are small and medium-sized enterprises. We are really focused on how this regime will work for the 25,000 tech companies that are set to be in scope, and our approach is really on the implementation and how the Bill can deliver on the objectives.

Thank you so much for the question. There are some definite risks when we think about smaller businesses and the Online Safety Bill. Today, we have heard a lot of the names that come up with regard to tech companies; they are the larger companies. However, this will be a regime that impacts thousands of different tech companies, with different functionalities and different roles within the ecosystem, all of which contribute to the economy in their own way.

There are specific areas to be addressed in the Bill, where there are some threats to innovation and investment by smaller businesses. First, greater clarity is needed. In order for this regime to be workable for smaller businesses, they need clarity on guidelines and on definitions, and they also need to be confident that the systems and processes that they put in place will be sustainable—in other words, the right ones.

Certain parts of the regime risk not having enough clarity. The first thing that I will point to is around the definitions of harm. We would very much welcome having some definitions of harmful content, or even categories of harmful content, in primary legislation. It might then be for Ofcom to determine how those definitions are interpreted within the codes, but having things to work off and types of harmful content for smaller businesses to start thinking about would be useful; obviously, that will be towards children, given that they are likely to be category 2.

The second risk for smaller businesses is really around the powers of the Secretary of State. I think there is a real concern. The Secretary of State will have some technical powers, which are pretty much normal; they are what you would expect in any form of regulation. However, the Online Safety Bill goes a bit further than that, introducing some amendment powers. So, the Secretary of State can modify codes of practice to align with public policy. In addition to that, there are provisions to allow the Secretary of State to set thresholds between the categories of companies.

Smaller businesses want to start forming a strong relationship with Ofcom and putting systems and processes in place that they can feel confident in. If they do not have that level of confidence and if the regime could be changed at any point, they might not be able to progress with those systems and processes, and when it comes to kind of pushing them out of the market, they might not be able to keep up with some of the larger companies that have been very much referenced in every conversation.

So, we need to think about proportionality, and we need to think about Ofcom's independence and the kind of relationship that it can form with smaller businesses. We also need to think about balance. This regime is looking to strike a balance between safety, free speech and innovation in the UK's digital economy. Let us just ensure that we provide enough clarity for businesses so that they can get going and have confidence in what they are doing.

Q121 Alex Davies-Jones: Thank you, Lulu. Adam and Ian, if either of you want to come in at any point, please just indicate that and I will bring you in.

The Chair: May I just apologise before we go any further, because I got you both the wrong way round? I am sorry. It is Mr Stevenson who is online and it is Adam Hildreth who is here in body and person.

Adam Hildreth: I think we have evolved as a world actually, when it comes to online safety. I think that if you went back five or 10 years, safety would have come after your people had developed their app, their platform or whatever they were creating from a tech perspective. I think we are now in a world where safety, in various forms, has to be there by default. And moving on to your point, we have to understand what that means for different sizes of businesses. The risk assessment word or phrase for me is the critical part there, which is putting blocks in front of people who are innovating and creating entrepreneurial businesses that make the online world a better place. Putting those blocks in without them understanding whether they can compete or not in an open and fair market is where we do not want to be.

So, getting to the point where it is very easy to understand is important—a bit like where we got to in other areas, such as data protection and where we went with the GDPR. In the end, it became simplified; I will not use the word "simplified" ever again in relation to GDPR, but it did become simplified from where it started. It is really important for anyone developing any type of tech platform that the Online Safety Bill will affect that they understand exactly what they do and do not have to put in place; otherwise, they will be taken out just by not having a legal understanding of what is required.

The other point to add, though, is that there is a whole other side to online safety, which is the online safety tech industry. There are tons of companies in the UK and worldwide that are developing innovative technologies that solve these problems. So, there is a positive as well as an understanding of how the Bill needs to be created and publicised, so that people understand what the boundaries are, if you are a UK business.

The Chair: Mr Stevenson, you are nodding. Do you want to come in?

Ian Stevenson: I agree with the contributions from both Adam and Lulu. For me, one of the strengths of the Bill in terms of the opportunity for innovators is that so much is left to Ofcom to provide codes of practice and so on in the future, but simultaneously that is its weakness in the short term. In the absence of those codes of practice and definitions of exactly where the boundaries between merely undesirable and actually harmful and actionable might lie, the situation is very difficult. It is very difficult for companies like my own and the other members of the Online Safety Tech Industry Association, who are trying to produce technology to support safer experiences online, to know exactly what that technology should do until we know which harms are in scope and exactly what the thresholds are and what the definitions of those harms are. Similarly, it is very hard for anybody building a service to know what technologies, processes and procedures they will need until they have considerably more detailed information than they have at the moment.

I agree that there are certain benefits to having more of that in the Bill, especially when it comes to the harms, but in terms of the aspiration and of what I hear is the objective of the Bill—creating safer online experiences—we really need to understand when we are going to have much more clarity and detail from Ofcom and any other relevant party about exactly what is going to be seen as best practice and acceptable practice, so that people can put in place those measures on their sites and companies in the Online Safety Tech Industry Association can build the tools to help support putting those measures in place.

Q122 Alex Davies-Jones: Thank you all. Lulu, you mentioned concerns about the Secretary of State's powers and Ofcom's independence. Other concerns expressed about Ofcom include its ability to carry out this regulation. It is being hailed as the saviour of the internet by some people. Twenty-five thousand tech companies in the UK will be under these Ofcom regulations, but questions have been asked about its technical and administrative capacity to do this. Just today, there is an online safety regulator funding policy adviser role being advertised by the Department for Digital, Culture, Media and Sport. Part of the key roles and responsibilities are:

"The successful post holder will play a key role in online safety as the policy advisor on Funding for the Online Safety Regulator."

Basically, their job is to raise money for Ofcom. Does that suggest concerns about the role of Ofcom going forward, its funding, and its resource and capacity to support those 25,000 platforms?

Lulu Freemont: It is a very interesting question. We really support Ofcom in this role. We think that it has a very good track record with other industries that are also in techUK's membership, such as broadcasters. It has done a very good job at implementing proportionate regulation. We know that it has been increasing its capacity for some time now, and we feel confident that it is working with us as the trade and with a range of other experts to try to understand some of the detail that it will have to understand to regulate.

One of the biggest challenges—we have had this conversation with Ofcom as well—is to understand the functionalities of tech services. The same functionality might be used in a different context, and that functionality could be branded as very high risk in one context but very low risk in another. We are having those conversations now. It is very important that they are being had now, and we would very much welcome Ofcom publishing drafts. We know that is its intention, but it should bring everything forward in terms of all the gaps in this regulation that are left to Ofcom's codes, guidance and various other documentation.

Adam Hildreth: One of the challenges that I hear a lot, and that we hear a lot at Crisp in our work, is that people think that the Bill will almost eradicate all harmful content everywhere. The challenge that we have with content is that every time we create a new technology or mechanism that defeats harmful or illegal content, the people who are creating it—they are referred to in lots of ways, but bad actors, ultimately—create another mechanism to do it. It is very unlikely that we will ever get to a situation in which it is eradicated from every platform forever—though I hope we do.

What is even harder for a regulator is to be investigating why a piece of content is on a platform. If we get to a position where people are saying, "I saw this bit of content; it was on a platform," that will be a really dangerous place to be, because the funding requirement for any regulator will go off the charts—think about how much content we consume. I would much prefer to be in a situation where we think about the processes and procedures that a platform puts in place and making them appropriate, ensuring that if features are aimed at children, they do a risk assessment so that they understand how those features are being used and how they could affect children in particular—or they might have a much more diverse user group, whereby harm is much less likely.

So, risk assessments and, as Ian mentioned, technologies, processes and procedures—that is the bit that a regulator can do well. If your risk assessment is good and your technology, process and procedures are as good as they can be based on a risk assessment, that almost should mean that you are doing the best job you possibly can to stop that content appearing, but you are not eradicating it. It really worries me that we are in a position whereby people are going to expect that they will never see content on a platform again, even though billions of pieces of potentially harmful content could have been removed from those platforms.

Q123 Alex Davies-Jones: On that point, you mentioned that it is hard to predict the future and to regulate on the basis of what is already there. We have waited a long time for the Bill, and in that time we have had new platforms and new emerging technology appear. How confident are you that the Bill allows for future-proofing, in order that we can react to anything new that might crop up on the internet?

Adam Hildreth: I helped personally in 2000 and 2001, when online grooming did not even exist as a law, so I have been involved in this an awful long time, waiting for laws to exist. I do not think we will ever be in a situation in which they are future-proofed if we keep putting every possibility into law. There needs to be some principles there. There are new features launched every day, and assessments need to be made about who they pose a risk to and the level of risk. In the same way as you would do in all kinds of industries, someone should do an assessment from a health and safety perspective. From that, you then say, "Can we even launch it at all? Is it feasible? Actually, we can, because

we can take this amount of risk." Once they understand those risk assessments, technology providers can go further and develop technology that can combat this.

If we can get to the point where it is more about process and the expectations around people who are creating any types of online environments, apps or technologies, it will be future-proofed. If we start trying to determine exact pieces of content, what will happen is that someone will work out a way around it tomorrow, and that content will not be included in the Bill, or it will take too long to get through and suddenly, the whole principle of why we are here and why we are having this discussion will go out the window. That is what we have faced every day since 1998: every time the technology works out how to combat a new risk—whether that is to children, adults, the economy or society—someone comes along and works out a way around the technology or around the rules and regulations. It needs to move quickly; that will future-proof it.

The Chair: I have four Members plus the Minister to get in, so please be brief. I call Dean Russell.

Q124 Dean Russell: Thank you, Sir Roger. My question builds on the future-proofing. Obviously, the big focus now is the metaverse and a virtual reality world. My question has two parts. First, is the Bill helping already by encouraging the new start-ups in that space to put safety first? Secondly, do you agree that a Joint Committee of the Houses of Parliament that continued to look at the Act and its evolution over the long term once it had been passed would be beneficial? I will come to you first, Lulu.

Lulu Freemont: On future-proofing, one of the real strengths of the Bill is the approach: it is striving to rely on systems and processes, to be flexible and to adapt to future technologies. If the Bill sticks to that approach, it will have the potential to be future-proof. Some points in the Bill raise a slight concern about the future-proofness of the regulation. There is a risk that mandating specific technologies—I know that is one of Ofcom's powers under the Bill—would put a bit of a timestamp on the regulation, because those technologies will likely become outdated at some point. Ensuring that the regulation remains flexible enough to build on the levels of risk that individual companies have, and on the technologies that work for the development and innovation of those individual companies, will be a really important feature, so we do have some concerns around the mandating of specific technologies in the Bill.

On the point about setting up a committee, one of the things for which techUK has called for a really long time is an independent committee that could think about the current definitions of harm and keep them under review. As companies put in place systems and processes that might mitigate levels of risk of harm, will those levels of harm still be harmful? We need to constantly evolve the regime so that it is true to the harms and risks that are present today, and to evaluate it against human rights implications. Having some sort of democratically led body to think about those definitional points and evaluate them as times change and harm reduces through this regime would be very welcome.

Adam Hildreth: To add to that, are people starting to think differently? Yes, they definitely are. That ultimately, for me, is the purpose of the Bill. It is to get people to start thinking about putting safety as a core principle of what they do as an overall business—not just in the development of their products, but as the overall business. I think that will change things.

A lot of the innovation that comes means that safety is not there as the principal guiding aspect, so businesses do need some help. Once they understand how a particular feature can be exploited, or how it impacts certain demographics or particular age groups—children being one of them—they will look for solutions. A lot of the time, they have no idea before they create this amazing new metaverse, or this new metaverse game, that it could actually be a container for harmful content or new types of harm. I think this is about getting people to think. The risk assessment side is critical, for me—making sure they go through that process or can bring on experts to do that.

Ian Stevenson: I would split the future-proofing question into two parts. There is a part where this Bill will provide Ofcom with a set of powers, and the question will be: does Ofcom have the capacity and agility to keep up with the rate of change in the tech world? Assuming it does, it will be able to act fairly quickly. There is always a risk, however, that once a code of conduct gets issued, it becomes very difficult to update that code of conduct in a responsive way.

There is then a second piece, which is: are the organisations that are in scope of regulation, and the powers that Ofcom has, sufficient as things change? That is where the idea of a long-term committee to keep an eye on this is extremely helpful. That would be most successful if it did not compromise Ofcom's independence by digging deeply into individual codes of conduct or recommendations, but rather focused on whether Ofcom has the powers and capacity that it needs to regulate as new types of company, platform and technology come along.

Dean Russell: Thank you.

Q125 Kirsty Blackman: My first question is for Lulu. Do small tech companies have enough staff with technical expertise to be able to fulfil their obligations under the Bill?

Lulu Freemont: It is a great question. One of the biggest challenges is capacity. We hear quite a lot from the smaller tech businesses within our membership that they will have to divert their staff away from existing work to comply with the regime. They do not have compliance teams, and they probably do not have legal counsel. Even at this stage, to try to understand the Bill as it is currently drafted—there are lots of gaps—they are coming to us and saying, "What does this mean in practice?" They do not have the answers, or the capability to identify that. Attendant regulatory costs—thinking about the staff that you have and the cost, and making sure the regulation is proportionate to the need to divert away from business development or whatever work you might be doing in your business—are really fundamental.

Another real risk, and something in the Bill that smaller businesses are quite concerned about, is the potential proposal to extend the senior management liability provisions. We can understand them being in there to enable the regulators to do their job—information requests—but if there is any extension into individual pieces of content, coupled with a real lack of definitions,

those businesses might find themselves in the position of restricting access to their services, removing too much content or feeling like they cannot comply with the regime in a proportionate way. That is obviously a very extreme case study. It will be Ofcom's role to make sure that those businesses are being proportionate and understand the provisions, but the senior management liability does have a real, chilling impact on the smaller businesses within our membership.

Adam Hildreth: One of the challenges that we have seen over the last few years is that you can have a business that is small in revenue but has a huge global user base, with millions of users, so it is not really a small business; it just has not got to the point where it is getting advertisers and getting users to pay for it. I have a challenge on the definition of a small to medium-sized business. Absolutely, for start-ups with four people in a room—or perhaps even still just two—that do not have legal counsel or anything else, we need to make it simple for those types of businesses to ingest and understand what the principles are and what is expected of them. Hopefully they will be able to do quite a lot early on.

The real challenge comes when someone labels themselves as a small business but they have millions of users across the globe—and sometimes actually quite a lot of people working for them. Some of the biggest tech businesses in the world that we all use had tens of people working for them at one point in time, when they had millions of users. That is the challenge, because there is an expectation for the big-tier providers to be spending an awful lot of money, when the small companies are actually directly competing with them. There is a challenge to understanding the definition a small business and whether that is revenue-focused, employee-focused or about how many users it has—there may be other metrics.

Ian Stevenson: One of the key questions is how much staffing this will actually take. Every business in the UK that processes data is subject to GDPR from day one. Few of them have a dedicated data protection officer from day one; it is a role or responsibility that gets taken on by somebody within the organisation, or maybe somebody on the board who has some knowledge. That is facilitated by the fact that there are a really clear set of requirements there, and there are a lot of services you can buy and consume that help you deliver compliance. If we can get to a point where we have codes of practice that make very clear recommendations, then even small organisations that perhaps do not have that many staff to divert should be able to achieve some of the basic requirements of online safety by buying in the services and expertise that they need. We have seen with GDPR that many of those services are affordable to small business.

If we can get the clarity of what is required right, then the staff burden does not have to be that great, but we should all remember that the purpose of the Bill is to stop some of the egregiously bad things that happen to people as a result of harmful content, harmful behaviours and harmful contact online. Those things have a cost in the same way that implementing data privacy has a cost. To come back to Lulu's point, it has to be proportionate to the business.

Q126 Mrs Miller: Adam, you said a few moments ago that companies are starting to put safety at the core of what they do, which will be welcome to us all—maybe it should have happened a lot earlier. I know you have worked a lot in that area. Regulators and company owners will have to depend on an ethical culture in their organisations if they are going to abide by the new regulations, because they cannot micromanage and regulators cannot micromanage. Will the Bill do enough to drive that ethical culture? If not, what more could it do or could the industry do? I would be really interested in everybody's answer to this one, but I will start with Adam.

Adam Hildreth: What we are seeing from the people that are getting really good at this and that really understand it is that they are treating this as a proper risk assessment, at a very serious level, across the globe. When we are talking about tier 1s, they are global businesses. When they do it really well, they understand risk and how they are going to roll out systems, technology, processes and people in order to address that. That can take time. Yes, they understand the risk, who it is impacting and what they are going to do about it, but they still need to train people and develop processes and maybe buy or build technology to do it.

We are starting to see that work being done really well. It is done almost in the same way that you would risk assess anything else: corporate travel, health and safety in the workplace—anything. It should really become one of those pillars. All those areas I have just gone through are regulated. Once you have regulation there, it justifies why someone is doing a risk assessment, and you will get businesses and corporates going through that risk assessment process. We are seeing others that do not do the same level of risk assessment and they do not have that same buy-in.

Q127 Mrs Miller: Lulu, how do you drive a culture change?

Lulu Freemont: TechUK's membership is really broad. We have cyber and defence companies in our membership, and large platforms and telcos. We speak on behalf of the sector. We would say that there is a real commitment to safety and security.

To bring it back to regulation, the risk-based approach is very much the right one—one that we think has the potential to really deliver—but we have to think about the tech ecosystem and its diversity. Lots of TechUK members are on the business-to-business side and are thinking about the role that they play in supporting the infrastructure for many of the platforms to operate. They are not entirely clear that they are exempt in the Bill. We understand that it is a very clear policy intention to exempt those businesses, but they do not have the level of legal clarity that they need to understand their role as access facilities within the tech.

That is just one example of a part of the sector that you would not expect to be part of this culture change or regulation but which is being caught in it slightly as an unintended consequence of legal differences or misinterpretations. Coming from that wide-sector perspective, we think that we need clarity on those issues to understand the different functionalities, and each platform and service will be different in their approach to this stuff.

Q128 Mrs Miller: Ian, how do you drive a culture change in the sector?

Ian Stevenson: I think you have to look at the change you are trying to effect. For many people in the sector, there is a lack of awareness about what happens when the need to consider safety in building features is not put first. Even when you realise how many bad things can happen online, if you do not know what to do about it, you tend not to be able to do anything about it.

If we want to change culture—it is the same for individual organisations as for the sector as a whole—we have to educate people on what the problem is and give them the tools to feel empowered to do something about it. If you educate and empower people, you remove the barrier to change. In some places, an extremely ethical people-centric and safety-focused culture very naturally emerges, but in others, less so. That is precisely where making it a first-class citizen in terms of risk assessment for boards and management becomes so important. When people see management caring about things, that gets pushed out through the organisations.

Q129 Kim Leadbeater: In your view, what needs to be added or taken away from the Bill to help it achieve the Government's aim of making the UK
"the safest place in the world to be online"?

Lulu Freemont: First, I want to outline that there are some strong parts in the Bill that the sector really supports. I think the majority of stakeholders would agree that the objectives are the right ones. The Bill tries to strike a balance between safety, free speech and encouraging innovation and investment in the UK's digital economy. The approach—risk-based, systems-led and proportionate—is the right one for the 25,000 companies that are in scope. As it does not focus on individual pieces of content, it has the potential to be future-proof and to achieve longer-term outcomes.

The second area in the Bill that we think is strong is the prioritisation of illegal content. We very much welcome the clear definitions of illegal content on the face of the Bill, which are incredibly useful for businesses as they start to think about preparing for their risk assessment on illegal content. We really support Ofcom as the appropriate regulator.

There are some parts of the Bill that need specific focus and, potentially, amendments, to enable it to deliver on those objectives without unintended consequences. I have already mentioned a few of those areas. The first is defining harmful content in primary legislation. We can leave it to codes to identify the interpretations around that, but we need definitions of harmful content so that businesses can start to understand what they need to do.

Secondly, we need clarity that businesses will not be required to monitor every piece of content as a result of the Bill. General monitoring is prohibited in other regions, and we have concerns that the Online Safety Bill is drifting away from those norms. The challenges of general monitoring are well known: it encroaches on individual rights and could result in the over-removal of content. Again, we do not think that the intention is to require companies of all sizes to look at every piece of content on their site, but it might be one of the unintended consequences, so we would like an explicit prohibition of general monitoring on the face of the Bill.

We would like to remove the far-reaching amendment powers of the Secretary of State. We understand the need for technical powers, which are best practised within regulation, but taking those further so that the Secretary of State can amend the regime in such an extreme way to align with public policy is of real concern, particularly to smaller businesses looking to confidently put in place systems and processes. We would like some consideration of keeping senior management liability as it is. Extending that further is only going to increase the chilling impact that it is having and the environment it is creating within UK investment. The final area, which I have just spoken about, is clarifying the scope. The business-to-business companies in our membership need clarity that they are not in scope and for that intention to be made clear on the face of the Bill.

We really support the Bill. We think it has the potential to deliver. There are just a few key areas that need to be changed or amended slightly to provide businesses with clarity and reassurances that the policy intentions are being delivered on.

Adam Hildreth: To add to that—Lulu has covered absolutely everything, and I agree—the critical bit is not monitoring individual pieces of content. Once you have done your risk assessment and put in place your systems, processes, people and technology, that is what people are signing up for. They are not signing up for this end assessment where, because you find that one piece of harmful content exists, or maybe many, you have failed to abide by what you are really signing up to.

That is the worry from my perspective: that people do a full risk assessment, implement all the systems, put in place all the people, technology and processes that they need, do the best job they can and have understood what investment they are putting in, and someone comes along and makes a report to a regulator—Ofcom, in this sense—and says, "I found this piece of content there." That may expose weaknesses, but the very best risk assessments are ongoing ones anyway, where you do not just put it away in a filing cabinet somewhere and say, "That's done." The definitions of online harms and harmful content change on a daily basis, even for the biggest social media platforms; they change all the time. There was talk earlier about child sexual abuse material that appears as cartoons, which would not necessarily be defined by certain legislation as illegal. Hopefully the legislation will catch up, but that is where that risk assessment needs to be made again, and policies may need to be changed and everything else. I just hope we do not get to the point where the individual monitoring of content, or content misses, is the goal of the Bill—that the approach taken to online safety is this overall one.

The Chair: Thank you. I call the Minister.

Q130 Chris Philp: Thank you, Sir Roger, and thank you very much indeed for joining us for this afternoon's session. Adam, we almost met you in Leeds last October or November, but I think you were off with covid at the time.

Adam Hildreth: I had covid at the time, yes.

Chris Philp: Covid struck. I would like to ask Adam and Ian in particular about the opportunities provided by emerging and new technology to deliver the Bill's objectives. I would like you both to give examples of where you think new tech can help deliver these safety duties. I ask you to comment particularly on what it

might do on, first, age assurance—which we debated in our last session—and secondly, scanning for child sexual abuse images in an end-to-end encrypted environment. Adam, do you want to go first?

Adam Hildreth: Well, if Ian goes first, the second question would be great for him to answer, because we worked on it together.

Chris Philp: Fair enough. Ian?

Ian Stevenson: Yes, absolutely. The key thing to recognise is that there is a huge and growing cohort of companies, around the world but especially in the UK, that are working on technologies precisely to try to support those kinds of safety measures. Some of those have been supported directly by the UK Government, through the safety tech challenge fund, to explore what can be done around end-to-end encrypted messaging. I cannot speak for all the participants, but I know that many of them are members of the safety tech industry association.

Between us, we have demonstrated a number of different approaches. My own company, Cyacomb, demonstrated technology that could block known child abuse within encrypted messaging environments without compromising the privacy of users' messages and communications. Other companies in the UK, including DragonflAI and Yoti, demonstrated solutions based on detecting nudity and looking at the ages of the people in those images, which are again hugely valuable in this space. Until we know exactly what the regulation is going to demand, we cannot say exactly what the right technology to solve it is.

However, I think that the fact that that challenge alone produced five different solutions looking at the problem from different angles shows just how vibrant the innovation ecosystem can be. My background in technology is long and mixed, but I have seen a number of sectors emerge—including cyber-security and fintech—where, once the foundations for change have been created, the ability of innovators to come up with answers to difficult questions is enormous. The capacity to do that is enormous.

There are a couple of potential barriers to that. The strength of the regulation is that it is future proof. However, until we start answering the question, "What do we need to do and when? What will platforms need to do and when will they need to do it?" we do not really create in the commercial market the innovation drivers for the technical solutions that will deliver this. We do not create the drivers for investment. It is really important to be as specific as we can about what needs to be done and when.

The other potential barrier is regulation. We have already had a comment about how there should be a prohibition of general monitoring. We have seen what has happened in the EU recently over concerns about safety technologies that are somehow looking at traffic on services. We need to be really clear that, while safety technologies must protect privacy, there needs to be a mechanism so that companies can understand when they can deploy safety technologies. At the moment there are situations where we talk to potential customers for safety technologies and they are unclear as to whether it would be proportionate to deploy those under, for example, data protection law. There are areas, even within the safety tech challenge fund work on end-to-end encrypted messaging, where it was unclear whether some of the technologies—however brilliant they were at preventing child abuse in those encrypted environments—would be deployable under current data protection and privacy of electronic communications regulations.

There are questions there. We need to make sure that when the Online Safety Bill comes through, it makes clear what is required and how it fits together with other regulations to enable that. Innovators can do almost anything if you give them time and space. They need the certainty of knowing what is required, and an environment where solutions can be deployed and delivered.

Q131 Chris Philp: Ian, thank you very much. I am encouraged by your optimism about what innovation can ultimately deliver. Adam, let me turn to you.

Adam Hildreth: I agree with Ian that the level of innovation is amazing. If we start talking about age verification and end-to-end encryptions, for me—I am going to say that same risk assessment phrase again—it absolutely depends on the type of service, who is using the service and who is exploiting the service, as to which safety technologies should be employed. I think it is dangerous to say, "We are demanding this type of technology or this specific technology to be deployed in this type of instance," because that removes the responsibility from the people who are creating it.

Q132 Chris Philp: Sorry to interject, but to be clear, the Bill does not do that. The Bill specifies the objectives, but it is tech agnostic. The manner of delivering those is, of course, not specified, either in the Bill or by Ofcom.

Adam Hildreth: Absolutely. Sorry, I was saying that I agree with how it has been worded. We know what is available, but technology changes all the time and solutions change all the time—we can do things in really innovative ways. However, the risk assessment has to bring together freedom of speech versus the types at risk of abuse. Is it children who are at risk, and if so, what are they at risk from? That changes the space massively when compared with some adult gaming communities, where what is harmful to them is very different from what harms other audiences. That should dictate for them what system and technology is deployed. Once we understand what best of breed looks like for those types of companies, we should know what good is.

Q133 Chris Philp: Thank you, Adam. We only have one minute left, so what is your prediction for the potential possibilities that emerging tech presents to deal with the issues of age assurance, which are difficult, and CSEA scanning, given end-to-end encrypted environments?

Adam Hildreth: The technology is there. It exists and it is absolutely deployable in the environments that need it. I am sure Ian would agree; we have seen it and done a lot of testing on it. The technology exists in the environments that need it.

Q134 Chris Philp: Including inside the end-to-end encrypted environment, rather than just at the device level? Quite a few of the safety challenge solutions that Ian mentioned are at the device level; they are not inside the encryption.

Adam Hildreth: There are ways that can work. Again, it brings in freedom of expression, global businesses and some other areas, so it is more about regulation and consumer concerns about the security of data, rather than whether technological solutions are available.

The Chair: Ms Freemont, Mr Hildreth and Mr Stevenson, thank you all very much indeed. We have run out of time. As ever, if you have any further observations that you wish to make, please put them in writing and let the Committee have them; we shall welcome them. Thank you for your time this afternoon. We are very grateful to you.

Examination of Witnesses

Jared Sine, Nima Elmi and Dr Rachel O'Connell gave evidence.

4.16 pm

The Chair: We are now going to hear from Jared Sine, who is the chief business affairs and legal officer at Match Group, and Nima Elmi, the head of public policy in Europe at Bumble, who is appearing by Zoom. Thank you for joining us. I hope you can hear us all right. Wave if you can.

Nima Elmi *indicated assent.*

The Chair: We also have Dr Rachel O'Connell, who is the CEO of TrustElevate. Good afternoon.

Q135 Barbara Keeley: Does the Bill differentiate enough between services that have different business models? If not, what do you think are the consequences of the lack of differentiation, and where could more differentiation be introduced? Shall we start with you, Jared Sine?

Jared Sine: Sure—thank you for the question. Business models play a pretty distinct role in the incentives of the companies. When we talk to people about Match Group and online dating, we try to point out a couple of really important things that differentiate what we do in the dating space from what many technology companies are doing in the social media space. One of those things is how we generate our revenue. The overwhelming majority of it is subscription-based, so we are focused not on time on platform or time on device, but on whether you are having a great experience, because if you are, you are going to come back and pay again, or you are going to continue your subscription with us. That is a really big differentiator, in terms of the business model and where incentives lie, because we want to make sure they have a great experience.

Secondly, we know we are helping people meet in real life. Again, if people are to have a great experience on our platforms, they are going to have to feel safe on them, so that becomes a really big focus for us.

Finally, we are more of a one-to-one platform, so people are not generally communicating to large groups, so that protects us from a lot of the other issues you see on some of these larger platforms. Ultimately, what that means is that, for our business to be successful, we really have to focus on safety. We have to make sure users come, have a good, safe experience, and we have to have tools for them to use and put in place to empower themselves so that they can be safe and have a great experience. Otherwise, they will not come back and tell their friends.

The last thing about our platforms is that ultimately, if they are successful, our users leave them because they are engaged in a relationship, get married or just decide they are done with dating all together—that happens on occasion, too. Ultimately, our goal is to make sure that people have that experience, so safety becomes a core part of what we do. Other platforms are more focused on eyeballs, advertising sales and attention—if it bleeds, it leads—but those things are just not part of the equation for us.

Q136 Barbara Keeley: And do you think the Bill differentiates enough? If not, what more could be done in it?

Jared Sine: We are very encouraged by the Bill. We think it allows for different codes of conduct or policy, as it relates to the various different types of businesses, based on the business models. That is exciting for us because we think that ultimately those things need to be taken into account. What are the drivers and the incentives in place for those businesses? Let us make sure that we have regulations in place that address those needs, based on the approaches of the businesses.

The Chair: Nima, would you like go next?

Nima Elmi: Thank you very much for inviting me along to this discussion. Building on what Jared said, currently the Bill is not very clear in terms of references to categorisations of services. It clusters together a number of very disparate platforms that have different platform designs, business models and corporate aims. Similarly to Match Group, our platform is focused much more on one-to-one communications and subscription-based business models. There is an important need for the Bill to acknowledge these different types of platforms and how they engage with users, and to ensure appropriate guidance from Ofcom on how they should be categorised, rather than clustering together a rather significant amount of companies that have very different business aims in in this space.

The Chair: Dr O'Connell, would you like to answer?

Dr Rachel O'Connell: Absolutely. I think those are really good points that you guys have raised. I would urge a little bit of caution around that though, because I think about Yellow Tinder, which was the Tinder for teens, which has been rebranded as Yubo. It transgresses: it is a social media platform; it enables livestreaming of teens to connect with each other; it is ultimately for dating. So there is a huge amount of risk. It is not a subscription-based service.

I get the industry drive to say, "Let's differentiate and let's have clarity", but in a Bill, essentially the principles are supposed to be there. Then it is for the regulator, in my view, to say, at a granular level, that when you conduct a risk impact assessment, you understand whether the company has a subscription-based business model, so the risk is lower, and also if there is age checking to make sure those users are 18-plus. However, you must also consider that there are teen dating sites, which would definitely fall under the scope of this Bill and the provisions that it is trying to make to protect kids and to reduce the risk of harm.

While I think there is a need for clarity, I would urge caution. For the Bill to have some longevity, being that specific about the categorisations will have some potential unintended consequences, particularly as it relates to children and young people.

Q137 Barbara Keeley: The next question is really about age verification, which you have touched on, so let us start with you, Dr O'Connell. What do you think the Bill should contain to enable age verification or the age assurance needed to protect children online?

Dr Rachel O'Connell: There is a mention of age assurance in the Bill. There is an opportunity to clarify that a little further, and also to bring age verification services under the remit of the Bill, as they are serving and making sure that they are mitigating risk. There was a very clear outline by Elizabeth Denham when we were negotiating the Digital Economy Act in relation to age verification and adult content sites; she was very specific when she came to Committee and said it should be a third party conducting the checks. If you want to preserve privacy and security, it should be a third-party provider that runs the checks, rather than companies saying, "You know what? We'll track everybody for the purposes of age verification."

There needs to be a clear delineation, which currently in clause 50 is not very clear. I would recommend that that be looked at again and that some digital identity experts be brought into that discussion, so that there is a full appreciation. Currently, there is a lot of latitude for companies to develop their own services in-house for age verification, without, I think, a proper risk assessment of what that might mean for end users in terms of eroding their privacy.

Q138 Barbara Keeley: TikTok were talking to us earlier about their age verification. If companies do it themselves rather than it being a third party, where does that fall down?

Dr Rachel O'Connell: That means you have to track and analyse people's activities and you are garnering a huge amount of data. If you are then handling people under the age of 13, under the Data Protection Act, you must obtain parental consent prior to processing data. By definition, you have to gather the data from parents. I have been working in this space for 25 years. I remember, in 2008, when the Attorneys General brought all the companies together to consider age verification as part of the internet safety technical task force, the arguments of industry—I was in industry at the time—were that it would be overly burdensome and a privacy risk. Looking back through history, industry has said that it does not want to do that. Now, there is an incentive to potentially do that, because you do not have to pay for a third party to do it, but what are the consequences for the erosion of privacy and so on?

I urge people to think carefully about that, in particular when it comes to children. It would require tracking children's activities over time. We do not want our kids growing up in a surveillance society where they are being monitored like that from the get-go. The advantage of a third-party provider is that they can have a zero data model. They can run the checks without holding the data, so you are not creating a data lake. The parent or child provides information that can be hashed on the device and checked against data sources that are hashed, which means there is no knowledge. It is a zero data model.

The information resides on the user's device, which is pretty cool. The checks are done, but there is no exposure and no potential for man-in-the-middle checks. The company then gets a token that says "This person is over 18", or "This person is below 12. We have verified parental responsibility and that verified parent has given consent." You are dealing with tokens that do not contain any personal information, which is a far better approach than companies developing things in-house.

Q139 Barbara Keeley: I think the TikTok example was looking at materials and videos and seeing whether they mention school or birthdays as a way of verifying age. As you say, that does involve scanning the child's data.

Q140 The Chair: Can I see if Ms Elmi wants to come in? She tends to get left out on a limb, on the screen. Are you okay down there? Do you need to come in on this, or are you happy?

Nima Elmi: Yes, I am. I have nothing to add.

Q141 Barbara Keeley: Jared Sine, did you have anything to add?

Jared Sine: Sure. I would add a couple of thoughts. We run our own age verification scans, which we do through the traditional age gate but also through a number of other scans that we run.

Again, online dating platforms are a little different. We warn our users upfront that, as they are going to be meeting people in real life, there is a fine balance between safety and privacy, and we tend to lean a little more towards safety. We announce to our users that we are going to run message scans to make sure there is no inappropriate behaviour. In fact, one of the tools we have rolled out is called "Are you sure? Does this bother you?", through which our AI looks at the message a user is planning to send and, if it is an inappropriate message, a flag will pop up that says, "Are you sure you want to send this?" Then, if they go ahead and send it, the person receiving it at the other end will get a pop-up that says, "This may not be something you want to see. Go ahead and click here if you want to." If they open it, they then get another pop-up that asks "Does this bother you?" and, if it does, you can report the user immediately.

We think that is an important step to keep our platform safe. We make sure our users know that it is happening, so it is not under the table. However, we think there has to be a balance between safety and privacy, especially when we have users who are meeting in person. We have actually demonstrated on our platforms that this reduces harassment and behaviour that would otherwise be untoward or that you would not want on the platform.

We think that we have to be careful not to tie the hands of industry to be able to come up with technological solutions and advances that can work side by side with third-party tools and solutions. We have third-party ID verification tools that we use. If we identify or believe a user is under the age of 18, we push them through an ID verification process.

The other thing to remember, particularly as it relates to online dating, is that companies such as ours and Bumble have done the right thing by saying "18-plus only on our platforms". There is no law that says that an online dating platform has to be 18-plus, but we think it is right thing to do. I am a father of five kids; I would not want kids on my platform. We are very vigilant in taking steps to make sure we are using the latest and greatest tools available to try to make sure that our platforms are safe.

Q142 Mrs Miller: Rachel, we have, in you, what we are told is a leading, pre-eminent authority on the issue of age verification, so we are listening very carefully to what you say. I am thinking about the evidence we had earlier today, which said that it is reasonably straightforward for a large majority of young people to subvert age verification through the use of VPNs. You have been advocating third-party verification. How could we also deal with this issue of subverting the process through the use of the VPNs?

Dr Rachel O'Connell: I am the author of the technical standard PAS 1296, an age checking code of practice, which is becoming a global standard at the moment. We worked a lot with privacy and security and identity experts. It should have taken nine months, but it took a bit longer. There was a lot of thought that went into it. Those systems were developed to, as I just described, ensure a zero data, zero knowledge kind of model. What they do is enable those verifications to take place and reduce the requirement. There is a distinction between monitoring your systems, as was said earlier, for age verification purposes and abuse management. They are very different. You have to have abuse management systems. It is like saying that if you have a nightclub, you have to have bouncers. Of course you have to check things out. You need bouncers at the door. You cannot let people go into the venue, then afterwards say that you are spotting bad behaviour. You have to check at the door that they are the appropriate age to get into the venue.

Q143 Mrs Miller: Can they not just hop on a VPN and bypass the whole system anyway?

Dr Rachel O'Connell: I think you guys will be aware of the DCMS programme of work about the verification of children last year. As part of that, there was a piece of research that asked children what they would think about age verification. The predominant thing that came across from young children is that they are really tired of having to deal with weirdos and pervs. It is an everyday occurrence for them.

To just deviate slightly to the business model, my PhD is in forensics and tracking paedophile activity on the internet way back in the '90s. At that time, guys would have to look for kids. Nowadays, on TikTok and various livestream platforms, the algorithms recognise that an individual—a man, for example—is very interested in looking at content produced by kids. The algorithms see that a couple of times and go, "You don't have to look anymore. We are going to seamlessly connect you with kids who livestream. We are also going to connect you with other men that like looking at this stuff."

If you are on these livestream sites at 3 o'clock in the morning, you can see these kids who are having sleepovers or something. They put their phone down to record whatever the latest TikTok dance is, and they think that they are broadcasting to other kids. You would assume that, but what they then hear is the little pops of love hearts coming on to the screen and guys' voices saying, "Hey sweetie, you look really cute. Lick your lips. Spread your legs." You know where I am going with this.

The Online Safety Bill should look at the systems and processes that underpin these platforms, because there is gamification of kids. Kids want to become influencers—maybe become really famous. They see the views counter and think, "Wow, there are 200 people looking at us." Those people are often men, who will co-ordinate their activities at the back. They will push the boys a little bit further, and if a girl is on her own, they will see. If the child does not respond to the request, they will drop off. The kid will think, "Oh my God. Well, maybe I should do it this one time."

What we have seen is a quadrupling of child sexual abuse material online that has been termed "self-generated", because the individual offender hasn't actually produced it. From a psychological perspective, it is a really bad name, but that is a separate topic. Imagine if that was your kid who had been coerced into something that had then been labelled as "self-generated". The businesses models that underpin those processes that happen online are certainly something that should be really within scope.

We do not spend enough time thinking about the implications of the use of recommendation engines and so on. I think the idea of the VPN is a bit of a red herring. Children want safety. They do not want to have to deal with this sort of stuff online. There are other elements. If you were a child and felt that you might be a little bit fat, you could go on YouTube and see whether you could diet or something. The algorithms will pick that up also. There is a tsunami of dieting and thinspiration stuff. There is psychological harm to children as a result of the systems and processes that these companies operate.

There was research into age verification solutions and trials run with BT. Basically, the feedback from both parents and children was, "Why doesn't this exist already?". If you go into your native EE app where it says, "Manage my family" and put in your first name, last name and mobile number and your child's first name, last name and date of birth, it is then verified that you are their parent. When the child goes on Instagram or TikTok, they put in their first and last name. The only additional data point is the parent's mobile number. The parent gets a notification and they say yes or no to access.

There are solutions out there. As others have mentioned, the young people want them and the parents want them. Will people try to work around them? That can happen, but if it is a parent-initiated process or a child-initiated process, you have the means to know the age bands of the users. From a business perspective, it makes a lot of sense because you can have a granular approach to the offerings you give to each of your customers in different age bands.

Nima Elmi: Just to add to what Rachel has said, I think she has articulated extremely well the complexities of the issues around not only age verification, but business models. Ultimately, this is such a complex matter that it requires continued consultation across industry, experts and civil society to identity pragmatic

recommendations for industry when it comes to not only verifying the age of their users, but thinking about the nuanced differences between platforms, purposes, functionality and business models, and what that means.

In the context of the work we do here at Bumble, we are clear about our guidelines requiring people to be 18-plus to download our products from app stores, as well as ensuring that we have robust moderation processes to identify and remove under-18s from our platforms. There is an opportunity here for the Bill to go further in providing clarity and guidance on the issue of accessibility of children to services.

Many others have said over the course of today's evidence that there needs to be a bit more colour put into definitions, particularly when certain sections of the Bill refer to what constitutes a "significant number of users" for determining child accessibility to platforms. Coupled with the fact that age verification or assurance is a complex area in and of itself and the nuance between how social media may engage with it versus a dating or social networking platform, I think that more guidance is very much needed and a much more nuanced approach would be welcome.

The Chair: I have three Members and the Minister to get in before 5 o'clock, so I urge brief questions and answers please.

Q144 Kirsty Blackman: Is it technically possible—I do not need to know how—to verify the age of children who are under 16, for example?

Dr Rachel O'Connell: Yes.

Q145 Kirsty Blackman: So technology exists out there for that to happen.

Dr Rachel O'Connell: Yes.

Q146 Kirsty Blackman: Once we have the verification of those ages, do you think it would be possible or desirable to limit children's interactions to only with other children? Is that the direction you were going in?

Dr Rachel O'Connell: I will give an example. If you go to an amusement park, kids who are below four feet, for example, cannot get on the adult rides, so the equivalent would be that they should not be on an 18-plus dating site. The service can create it at a granular level so the kids can interact with kids in the same age group or a little bit older, but they can also interact with family. You can create circles of trust among verified people.

Kirsty Blackman: For a game like Roblox, which is aimed at kids—it is a kids platform—if you had the age verification and if that worked, you could have a situation where a 13-year-old on Roblox could only interact with children who are between 12 and 14. Does the technology exist to make that work?

Dr Rachel O'Connell: You could do. Then if you were using it in esports or there was a competition, you could broaden it out. The service can set the parameters, and you can involve the parents in making decisions around what age bands their child can play with. Also, kids are really into esports and that is their future, so there are different circumstances and contexts that the technology could enable.

Q147 Kirsty Blackman: Finally, do you think it would be desirable for Ofcom to consider a system with more consistency in parental controls, so that parents can always ensure that their children cannot talk to anybody outside their circle? Would that be helpful?

Dr Rachel O'Connell: There is a history of parental controls, and only 36% of parents use them. Ofcom research consistently says that it is 70%, but in reality, it is lower. When using age verification, the parents are removing the ability to watch everything. It is a platform; they are providing the digital playground. In the same way, when you go on swings and slides, there is bouncy tarmac because you know the kids are going to use them. It is like creating that health and safety environment in a digital playground.

When parents receive a notification that their child wants to access something, there could be a colour-coded nutrition-style thing for social media, livestreaming and so on, and the parents could make an informed choice. It is then up to the platform to maintain that digital playground and run those kinds of detection systems to see if there are any bad actors in there. That is better than parental controls because the parent is consenting and it is the responsibility of the platform to create the safer environment. It is not the responsibility of the parent to look over the child's shoulder 24/7 when they are online.

Q148 Kim Leadbeater: The age verification stuff is really interesting, so thank you to our witnesses. On violence against women and girls, clauses 150 to 155 set out three new communications offences. Do you think those offences will protect women from receiving offensive comments, trolling and threats online? What will the Bill mean for changing the way you manage those risks on your platforms?

Jared Sine: I do not know the specific provisions but I am familiar with the general concept of them. Any time you put something in law, it can either be criminalised or have enforcement behind it, and I think that helps. Ultimately, it will be up to the platforms to come up with innovative technologies or systems such as "Are You Sure?" and "Does This Bother You?" which say that although the law says x, we are going to go beyond that to find tools and systems that make it happen on our platform. Although I think it is clearly a benefit to have those types of provisions in law, it will really come down to the platforms taking those extra steps in the future. We work with our own advisory council, which includes the founder of the #MeToo movement, REIGN and others, who advise us on how to make platforms safer for those things. That is where the real bread gets buttered, so to speak.

Q149 Kim Leadbeater: Do you think that is consistent across your industry? It sounds like you are taking a very proactive approach to it.

Jared Sine: We are proactive about it, and I know our colleagues and friends over at Bumble are proactive about it as well. Our heads of trust and safety both came from the same company—Uber—before coming to us, so I know that they compare notes quite regularly. Because of the way the legislation is set up, there can be codes of conduct applying specifically to online dating, and to the extent that that technology exists, you need to deploy it.

The Chair: Shall we ask our friends at Bumble if they would like to come in?

Nima Elmi: It is a great question. There are three points that I want to address, and I will try to be brief. First, Bumble is very much a uniquely female-founded and female-led tech company that adopts a safety-by-design ethos. It is baked within our DNA. The majority of our board are women, and they are public figures who, unfortunately, have to some extent experienced online harms targeting women.

We believe it is incredibly important that the Bill acknowledges that women are disproportionately impacted by online harms. Some studies have found that women are 27 times more likely than men to suffer online harassment and online harms. Currently, the Bill does not acknowledge or reference gender or women at all, so a lot more can be done, and we have submitted some recommendations.

Not every company in our industry or across the tech sector is female-founded and female-led, and they prioritise the harms that they want to tackle on their platforms very differently—that is important. Our systems-based approach, which bakes in safety-by-design principles, puts women at the centre of how our products are designed and used. We deploy corrective action and safety tools to make sure that our female members feel not only safe but empowered on our platforms. When it comes to managing risk, it is central to us to ensure that women feel safe on our products and services. We are here advocating for the fact that it should not just be our products that are safe for women—it should be the internet as a whole. In our view, the Bill does not currently go far enough to make sure that that happens.

We welcome the inclusion of the miscommunication offences in clauses 150 to 155 and also welcome the offence of cyber-flashing, the inclusion of which we have been advocating for publicly for several months. However, in both instances, and particularly with cyber-flashing, the Bill does not go far enough in acknowledging that it is an offence, as Professor McGlynn has highlighted, that should be grounded on consent rather than the motivation of the perpetrator.

Essentially, there are a number of inclusions that are a step in the right direction, but we would welcome significant changes to the Bill, predominantly through including a safety duty for women, to ensure that all platforms are consistent in their approach and prioritise how their female users engage with their services, so that they feel protected, and to ensure that determining those features is not predicated on the composition of the board or who the founder is.

The Chair: Right. For once, we seem to have run out of questions. Minister, do you wish to contribute?

Chris Philp: Everything I was going to ask has already been asked by my colleagues, so I will not duplicate that.

Q150 The Chair: In that case, given that we have the time, rather than doing what I normally do and inviting you to make any further submissions in writing, if there are any further comments that you would like to make about the Bill, the floor is yours. Let us start with Mr Sine.

Jared Sine: I would just make one brief comment. I think it has been mentioned by everyone here. Everyone has a role to play. Clearly, the Government have a role in proposing and pushing forward the legislation. The platforms that have the content have an obligation and a responsibility to try to make sure that their users are safe. One of the things that Dr O'Connell mentioned is age verification and trying to make sure that we keep young kids off platforms where they should not be.

I think there is a big role to play for the big tech platforms—the Apples and Googles—who distribute our apps. Over the years, we have said again and again to both of those companies, "We have age-gated our apps at 18, yet you will allow a user you know is 15, 14, 16—whatever it is—to download that app. That person has entered that information and yet you still allow that app to be downloaded." We have begged and pleaded with them to stop and they will not stop. I am not sure that that can be included in the Bill, but if it could be, it would be powerful.

If Apple and Google could not distribute any of our apps—Hinge, Match, Tinder—to anyone under the age of 18, that solves it right there. It is the same methodology that has been used at clubs with bouncers—you have a bouncer at the door who makes sure you are 21 before you go in and have a drink. It should be the same thing with these technology platforms. If they are going to distribute and have these app stores, the store should then have rules that show age-gated apps—"This is for 17-plus or 18-plus"—and should also enforce that. It is very unfortunate that our calls on this front have gone unanswered. If the Bill could be modified to include that, it would really help to address the issue.

Dr Rachel O'Connell: Absolutely. I 100% support that. There is a tendency for people to say, "It is very complex. We need a huge amount of further consultation." I started my PhD in 1996. This stuff has been going on for all that time. In 2008, there was a huge push by the Attorneys General, which I mentioned already, which brought all of the industry together. That was 2008. We are in 2022 now. 2017 was the Internet Safety Strategy Green Paper. We know what the risks are. They are known; we understand what they are. We understand the systems and processes that facilitate them. We understand what needs to be done to mitigate those risks and harms. Let's keep on the track that we are going on.

Regarding industry's concerns, a lot of them will be ironed out when companies are required to conduct risk assessments and impact assessments. They might ask, what are the age bands of your users? What are the risks associated with the product features that you are making available? What are the behaviour modification techniques that you are using, like endless scroll and loot boxes that get kids completely addicted? Are those appropriate for those ages? Then you surface the decision making within the business that results in harms and also the mitigations.

I urge you to keep going on this; do not be deterred from it. Keep the timeframe within which it comes into law fairly tight, because there are children out there who are suffering. As for the harassment—I have experienced it myself, it is horrible.

Those would be my final words.

The Chair: Thank you. Finally, Nima Elmi, please.

Nima Elmi: Thank you again for your time. I want to re-emphasise a couple of points, since we have a few minutes.

First, on the point around gendered harms, I think it is important for the Committee to really think about whether this is an opportunity to make reference in the Bill to acknowledge that women are experiencing online harms at a significantly higher rate than men. That is meant to futureproof the Bill, as new forms of online harms are, unfortunately, usually felt by women first. I know that Maria Miller, for example, has been doing extensive work around the issue of AI nudification tools, which, in the current framing of the Bill, would not be captured.

We would certainly urge that there is a greater focus in the Bill on gendered harms, whether that is through a specific safety duty, acknowledgement as a category within risk assessment, a designated code of practice—which I know Clare McGlynn, Refuge and EVAW have also advocated for—or acknowledgement of gender-based violence in transparency reporting.

Right now, the nature of moderation of technology platforms is very much grounded in the prioritisation of issues based on the leadership and usage of certain platforms, and this is an opportunity for the Government and Parliament to provide a standard setting that ensures consistency across the board while acknowledging the nuanced differences between the platforms and their business models, and their end goals. I would really like to emphasise that point.

The second point I want to emphasise, on cyber-flashing in particular, is the fact that we have an opportunity to bake in what should be societal standards that we want to hold people accountable to, both offline and online. Offences captured by the Bill that do not create a threshold where you will see prosecutions and a change in behaviour—for example, in the current formulation of the cyber-flashing offence, which is grounded in the perpetrator's motivation rather than in consent—will have little impact in changing the hearts and minds of individuals and stopping that behaviour, because the threshold will be so high.

We would definitely encourage the Committee to reflect on the pragmatic ways in which the Bill can be refined. In particular, I want to emphasise that it will be important to acknowledge that online harms are sadly very much experienced by women—both emerging forms and existing forms of harms. I welcome this opportunity to share this feedback with the Committee.

The Chair: Ms Elmi, Dr O'Connell and Mr Sine, thank you all very much indeed; the Committee is indebted to you. Thank you so much.

Examination of Witnesses

Rhiannon-Faye McDonald and Susie Hargreaves OBE gave evidence.

4.55 pm

The Chair: We will now hear from Rhiannon-Faye McDonald, victim and survivor advocate at the Marie Collins Foundation, and Susie Hargreaves, chief executive at the Internet Watch Foundation. Thank you for joining us this afternoon; first question, please.

Q151 Alex Davies-Jones: Thank you both for joining us this afternoon. One of the key objectives of the legislation is to ensure that a high level of protection for children and adults is in place. In your view, does the Bill in its current form achieve that?

Susie Hargreaves: Thank you very much for inviting me today. I think the Bill is working in the right direction. Obviously, the area that we at the IWF are concerned with is child sexual abuse online, and from our point of view, the Bill does need to make a few changes in order to put those full protections in place for children.

In particular, we have drafted an amendment to put co-designation on the face of the Bill. When it comes to child sexual abuse, we do not think that contracting out is an acceptable approach, because we are talking about the most egregious form of illegal material—we are talking about children—and we need to ensure that Ofcom is not just working in a collaborative way, but is working with experts in the field. What is really important for us at the moment is that there is nothing in the Bill to ensure that the good work that has been happening over 25 years in this country, where the IWF is held up as a world leader, is recognised, and that that expertise is assured on the face of the Bill. We would like to see that amendment in particular adopted, because the Bill needs to ensure that there are systems and processes in place for dealing with illegal material. The IWF already works with internet companies to ensure they take technical services.

There needs to be a strong integration with law enforcement—again, that is already in place with the memorandum of understanding between CPS, the National Police Chiefs' Council and the IWF. We also need clarity about the relationship with Ofcom so that child sexual abuse, which is such a terrible situation and such a terrible crime, is not just pushed into the big pot with other harms. We would like to see those specific changes.

Rhiannon-Faye McDonald: Generally, we think the Bill is providing a higher standard of care for children, but there is one thing in particular that I would like to raise. Like the IWF, the Marie Collins Foundation specialises in child sexual abuse online, specifically the recovery of people who have been affected by child sexual abuse.

The concern I would like to raise is around the contextual CSA issue. I know this has been raised before, and I am aware that the Obscene Publications Act 1959 has been brought into the list of priority offences. I am concerned that that might not cover all contextual elements of child sexual abuse: for example, where images are carefully edited and uploaded to evade content moderation, or where there are networks of offenders who are able to gain new members, share information with each other, and lead other people to third-party sites where illegal content is held. Those things might not necessarily be caught by the illegal content provisions; I understand that they will be dealt with through the "legal but harmful" measures.

My concern is that the "legal but harmful" measures do not need to be implemented by every company, only those that are likely to be accessed by children. There are companies that can legitimately say that the majority of their user base is not children, and therefore would not have to deal with that, but that provides a space for this contextual CSA to happen. While those platforms may not be accessed by children as much as other platforms, it still provides a place for this to happen—the harm can still occur, even if children do not come across it as much as they would elsewhere.

Q152 Alex Davies-Jones: On that point, one of the concerns that has been raised by other stakeholders is about the categorisation of platforms—for example,

category 1 and category 2B have different duties on them, as Ofcom is the regulator. Would you rather see a risk-based approach to platforms, rather than categorisation? What are your thoughts on that?

Susie Hargreaves: We certainly support the concept of a risk-based approach. We host very little child sexual abuse content in the UK, with the majority of the content we see being hosted on smaller platforms in the Netherlands and other countries. It is really important that we take a risk-based approach, which might be in relation to where the content is—obviously, we are dealing with illegal content—or in relation to where children are. Having a balance there is really important.

Q153 Alex Davies-Jones: A final question from me. We heard concerns from children's charities and the Children's Commissioner that the Bill does not account for breadcrumbing—the cross-platform grooming that happens on platforms. What more could the Bill do to address that, and do you see it as an omission and a risk?

Susie Hargreaves: I think we probably have a slightly different line from that of some of the other charities you heard from this morning, because we think it is very tricky and nuanced. What we are trying to do at the moment is define what it actually means and how we would have to deal with it, and we are working very closely with the Home Office to go through some of those quite intense discussions. At the moment, "harmful" versus "illegal" is not clearly defined in law, and it could potentially overwhelm certain organisations if we focus on the higher-level harms and the illegal material. We think anything that protects children is essential and needs to be in the Bill, but we need to have those conversations and to do some more work on what that means in reality. We are more interested in the discussions at the moment about the nuance of the issue, which needs to be mapped out properly.

One of the things that we are very keen on in the Bill as a whole is that there should be a principles-based approach, because we are dealing with new harms all the time. For example, until 2012 we had not seen self-generated content, which now accounts for 75% of the content we remove. So we need constantly to change and adapt to new threats as they come online, and we should not make the Bill too prescriptive.

The Chair: Ms McDonald?

Rhiannon-Faye McDonald: I was just thinking of what I could add to what Susie has said. My understanding is that it is difficult to deal with cross-platform abuse because of the ability to share information between different platforms—for example, where a platform has identified an issue or offender and not shared that information with other platforms on which someone may continue the abuse. I am not an expert in tech and cannot present you with a solution to that, but I feel that sharing intelligence would be an important part of the solution.

Q154 Mrs Miller: What risks do end-to-end encrypted platforms pose to children, and how should the Bill seek to mitigate those risks specifically?

Susie Hargreaves: We are very clear that end-to-end encryption should be within scope, as you have heard from other speakers today. Obviously, the huge threat on the horizon is the end-to-end encryption on Messenger, which would result in the loss of millions of images of child sexual abuse. In common with previous speakers, we believe that the technology is there. We need not to demonise end-to-end encryption, which in itself is not bad; what we need to do is ensure that children do not suffer as a consequence. We must have mitigations and safety mechanisms in place so that we do not lose these child sexual abuse images, because that means that we will not be able to find and support those children.

Alongside all the child protection charities, we are looking to ensure that protections equivalent to the current ones are in place in the future. We do not accept that the internet industry cannot put them in place. We know from experts such as Dr Hany Farid, who created PhotoDNA, that those mechanisms and protections exist, and we need to ensure that they are put in place so that children do not suffer as a consequence of the introduction of end-to-end encryption. Rhiannon has her own experiences as a survivor, so I am sure she would agree with that.

Rhiannon-Faye McDonald: I absolutely would. I feel very strongly about this issue, which has been concerning me for quite some time. I do not want to share too much, but I am a victim of online grooming and child sex abuse. There were images and videos involved, and I do not know where they are and who has seen them. I will never know that. I will never have any control over it. It is horrifying. Even though my abuse happened 19 years ago, I still walk down the street wondering whether somebody has seen those images and recognises me from them. It has a lifelong impact on the child, and it impacts on recovery. I feel very strongly that if end-to-end encryption is implemented on platforms, there must be safeguards in place to ensure we can continue to find and remove these images, because I know how important that is to the subject of those images.

Q155 Mrs Miller: So what needs to change in the Bill to make sure that happens? I am not clear.

Susie Hargreaves: We just want to make sure that the ability to scan in an end-to-end encrypted environment is included in the Bill in some way.

Q156 Kirsty Blackman: The ability to scan is there right now—we have got that—so you are just trying to make sure we are standing still, basically. Am I correct in my understanding?

Susie Hargreaves: I think with technology you can never stand still. We do not know what is coming down the line. We have to deal with the here and now, but we also need to be prepared to deal with whatever comes down the line. The answer, "Okay, we will just get people to report," is not a good enough replacement for the ability to scan for images.

When the privacy directive was introduced in Europe and Facebook stopped scanning for a short period, we lost millions of images. What we know is that we must continue to have those safety mechanisms in place. We need to work collectively to do that, because it is not acceptable to lose millions of images of child sexual abuse and create a forum where people can safely share them without any repercussions, as Rhiannon says. One survivor we talked to in this space said that one of her images had been recirculated 70,000 times. The ability

to have a hash of a unique image, go out and find those duplicates and make sure they are removed means that people are not re-victimised on a daily basis. That is essential.

Q157 Kirsty Blackman: Focusing on thinking about how to prevent grooming behaviour, does the Bill have enough in place to protect children from conversations that they may have adults, or from facing grooming behaviour online?

Rhiannon-Faye McDonald: There is one specific point that I would like to raise about this. I am concerned about private communications. We know that many offenders identify and target children on more open platforms, and then very quickly move them to more private platforms to continue the grooming and abuse. We were very pleased to see that private communications were brought in scope. However, there is a difficulty in the code of practice. When that is drafted, Ofcom is not going to be able to require proactive tools to be used to identify. That includes things like PhotoDNA and image and text-based classifiers.

So although we have tools that we can use currently, which can identify conversations where grooming is happening, we are not going to be using those immediately on private platforms, on private communications where the majority of grooming is going to happen. That means there will be a delay while Ofcom establishes that there is a significant problem with grooming on the platform, and then issues are noticed to require those tools to be used.

Q158 Kirsty Blackman: You mentioned the reporting mechanisms that are in place, Susie. Yes, they are not the only tool, and should not be the only tool—many more things should be happening—but are the reporting mechanisms that will be in place, once the Bill has come in and is being embedded, sufficient, or do they need to be improved as well; as requirements for platforms to have reporting mechanisms?

Susie Hargreaves: An awful lot of work has already gone into this over the past few years. We have been working closely with Departments on the draft code of practice. We think that, as it stands, it is in pretty good shape. We need to work more closely with Ofcom as those codes are developed—us and other experts in the field. Again, it needs to be very much not too directing, in the sense that we do not want to limit people, and to be available for when technology changes in the future. It is looking in the right shape, but of course we will all be part of the consultation and of the development of those practices as they go. It requires people to scan their networks, to check for child sexual abuse and—I guess for the first time, the main thing—to report on it. It is going to be a regulated thing. In itself, that is a huge development, which we very much welcome.

Q159 Kirsty Blackman: I have one last question. Rhiannon, a suggestion was made earlier by Dr Rachel O'Connell about age verification and only allowing children to interact with other children whose age is verified within a certain area. Do you think that would help to prevent online grooming?

Rhiannon-Faye McDonald: It is very difficult. While I am strongly about protecting children from encountering perpetrators, I also recognise that children need to have freedoms and the ability to use the internet in the ways that they like. I think if that was implemented and it was 100% certain that no adult could pose as a 13-year-old and therefore interact with actual 13-year-olds, that would help, but I think it is tricky.

Susie Hargreaves: One of the things we need to be clear about, particularly where we see children groomed —we are seeing younger and younger children—is that we will not ever sort this just with technology; the education piece is huge. We are now seeing children as young as three in self-generated content, and we are seeing children in bedrooms and domestic settings being tricked, coerced and encouraged into engaging in very serious sexual activities, often using pornographic language. Actually, a whole education piece needs to happen. We can put filters and different technology in place, but remember that the IWF acts after the event—by the time we see this, the crime has been committed, the image has been shared and the child has already been abused. We need to bump up the education side, because parents, carers, teachers and children themselves have to be able to understand the dangers of being online and be supported to build their resilience online. They are definitely not to be blamed for things that happen online. From Rhiannon's own story, how quickly it can happen, and how vulnerable children are at the moment—I don't know.

Rhiannon-Faye McDonald: For those of you who don't know, it happened very quickly to me, within the space of 24 hours, from the start of the conversation to the perpetrator coming to my bedroom and sexually assaulting me. I have heard other instances where it has happened much more quickly than that. It can escalate extremely quickly.

Just to add to Susie's point about education, I strongly believe that education plays a huge part in this. However, we must be very careful in how we educate children, so that the focus is not on how to keep themselves safe, because puts the responsibility on them, which in turn increases the feelings of responsibility when things do go wrong. That increased feeling of responsibility makes it less likely that they will disclose that something has happened to them, because they feel that they will be blamed. It will decrease the chance that children will tell us that something has happened.

Q160 Barbara Keeley: Just to follow up on a couple of things, mainly with Susie Hargreaves. You mentioned reporting mechanisms and said that reporting will be a step forward. However, the Joint Committee on the draft Bill recommended that the highest-risk services should have to report quarterly data to Ofcom on the results of their child sexual exploitation and abuse removal systems. What difference would access to that kind of data make to your work?

Susie Hargreaves: We already work with the internet industry. They currently take our services and we work closely with them on things such as engineering support. They also pay for our hotline, which is how we find child sexual abuse. However, the difference it would make is that we hope then to be able to undertake work where we are directly working with them to understand the level of their reports and data within their organisations.

At the moment, we do not receive that information from them. It is very much that we work on behalf of the public and they take our services. However, if we

were suddenly able to work directly with them—have information about the scale of the issue within their own organisations and work more directly on that—then that would help to feed into our work. It is a very iterative process; we are constantly developing the technology to deal with the current threats.

It would also help us by giving us more intelligence and by allowing us to share that information, on an aggregated basis, more widely. It would certainly also help us to understand that they are definitely tackling the problem. We do believe that they are tackling the problem, because it is not in their business interests not to, but it just gives a level of accountability and transparency that does not exist at the moment.

Q161 Barbara Keeley: You also said earlier that there was nothing in the Bill on co-designation—nothing to recognise the Internet Watch Foundation's 25 years of experience. Do you still expect to be co-designated as a regulator by Ofcom, and if so, what do you expect your role to be?

Susie Hargreaves: At the moment, there is nothing on the face of the Bill on co-designation. We do think that child sexual abuse is different from other types of harm, and when you think about the huge number of harms, and the scale and complexity of the Bill, Ofcom has so much to work with.

We have been working with Ofcom for the past year to look at exactly what exactly our role would be. However, because we are the country's experts on dealing with child sexual abuse material, because we have the relationships with the companies, and because we are an internationally renowned organisation, we are able to have that trusted relationship and then undertake a number of functions for Ofcom. We could help to undertake specific investigations, help update the code, or provide that interface between Ofcom and the companies where we undertake that work on their behalf.

We very much feel that we should be doing that. It is not about being self-serving, but about recognising the track record of the organisation and the fact that the relationships and technology are in place. We are already experts in this area, so we are able to work directly with those companies because we already work with them and they trust us. Basically, we have a memorandum of understanding with the CPS and the National Police Chiefs' Council that protects our staff from prosecution but the companies all work with us on a voluntary basis. They already work with us, they trust our data, and we have that unique relationship with them.

We are able to provide that service to take the pressure off Ofcom because we are the experts in the field. We would like that clarified because we want this to be right for children from day one—you cannot get it wrong when dealing with child sexual abuse. We must not undo or undermine the work that has happened over the last 25 years.

Q162 Barbara Keeley: Just to be clear, is there uncertainty somewhere in there? I am just trying to comprehend.

Susie Hargreaves: There is uncertainty, because we do not know exactly what our relationship with Ofcom is going to be. We are having discussions and getting on very well, but we do not know anything about what the relationship will be or what the criteria and timetable for the relationship are. We have been working on this for nearly five years. We have analysts who work every single day looking at child sexual abuse; we have 70 members of staff, and about half of them look at child sexual abuse every day. They are dealing with some of the worse material imaginable, they are already in a highly stressful situation and they have clear welfare needs; uncertainty does not help. What we are looking for is certainty and clarity that child sexual abuse is so important that it is included on the face of the Bill, and that should include co-designation.

The Chair: Thank you. One question from Kim Leadbeater.

Q163 Kim Leadbeater: Thank you for your very powerful testimony, Rhiannon. I appreciate that could not have been easy. Going back to the digital literacy piece, it feels like we were talking about digital literacy in the Bill when it started coming through, and that has been removed now. How important do you think it is that we have a digital literacy strategy, and that we hold social media providers in particular to having a strategy on digital education for young people?

Rhiannon-Faye McDonald: It is incredibly important that we have this education piece. Like Susie said, we cannot rely on technology or any single part of this to solve child sexual abuse, and we cannot rely on the police to arrest their way out of the problem. Education really is the key. That is education in all areas—educating the child in an appropriate way and educating parents. We hold parenting workshops. Parents are terrified; they do not know what to do, what platforms are doing what, or what to do when things go wrong. They do not even know how to talk to children about the issue; it is embarrassing for them and they cannot bring it up. Educating parents is a huge thing. Companies have a big responsibility there. They should have key strategies in place on how they are going to improve education.

Q164 Chris Philp: Can I start by thanking both Rhiannon-Faye and Susie for coming and giving evidence, and for all the work they are doing in this area? I know it has been done over many years in both cases.

I would like to pick up on a point that has arisen in the discussion so far—the point that Susie raised about the risks posed by Meta introducing end-to-end encryption, particularly on the Facebook Messenger service. You have referenced the fact that huge numbers of child sexual exploitation images are identified by scanning those communications, leading to the arrests of thousands of paedophiles each year. You also referenced the fact that when this was temporarily turned off in Europe owing to the privacy laws there—briefly, thankfully—there was a huge loss of information. We will come on to the Bill in a minute, but as technology stands now, if Meta did proceed with end-to-end encryption, would that scanning ability be lost?

Susie Hargreaves: Yes. It would not affect the Internet Watch Foundation, but it would affect the National Centre for Missing and Exploited Children. Facebook, as a US company, has a responsibility to do mandatory reporting to NCMEC, which will be brought in with the Bill in this country. Those millions of images would be lost, as of today, if they brought end-to-end encryption in now.

Q165 Chris Philp: Why would it not affect the Internet Watch Foundation?

Susie Hargreaves: Because they are scanning Facebook—sorry, I am just trying to unpack the way it works. It will affect us, actually. Basically, when we provide our hash list to Facebook, it uses that to scan Messenger, but the actual images that are found—the matches—are not reported to us; they are reported into NCMEC. Facebook does take our hash list. For those of you who do not know about hashing, it is a list of digital fingerprints—unique images of child sexual abuse. We currently have about 1.3 million unique images of child sexual abuse. Facebook does use our hash list, so yes it does affect us, because it would still take our hash list to use on other platforms, but it would not use it on Messenger. The actual matches would go into NCMEC. We do not know how many matches it gets against our hash list, because it goes into NCMEC.

Q166 Chris Philp: But its ability to check images going across Messenger against your list would effectively terminate.

Susie Hargreaves: Yes, sorry—I was unclear about that. Yes, it would on Messenger.

Q167 Chris Philp: Clearly the Bill cannot compel the creation of technology that does not exist yet. It is hoped that there will be technology—we heard evidence earlier suggesting that it is very close to existing—that allows scanning in an end-to-end encrypted environment. Do you have any update on that that you can give the Committee? If there is no such technology, how do you think the Bill should address that? Effectively there would be a forced choice between end-to-end encryption and scanning for CSEA content.

Susie Hargreaves: As I said before, it is essential that we do not demonise end-to-end encryption. It is really important. There are lots of reasons why, from a security and privacy point of view, people want to be able to use end-to-end encryption.

In terms of whether the technology is there, we all know that there are things on the horizon. As Ian said in the previous session, the technology is there and is about to be tried out. I cannot give any update at this meeting, but in terms of what we would do if end-to-end encryption is introduced and there is no ability to scan, we could look at on-device scanning, which I believe you mentioned before, Minister.

Chris Philp: Yes.

Susie Hargreaves: That is an option. That could be a backstop position. I think that, at the moment, we should stand our ground on this and say, "No, we need to ensure that we have some form of scanning in place if end-to-end encryption is introduced."

Q168 Chris Philp: For complete clarity, do you agree that the use of end-to-end encryption cannot be allowed at the expense of child safety?

Susie Hargreaves: I agree 100%.

Chris Philp: Good. Thank you.

The Chair: Thank you very much indeed, Ms McDonald and Ms Hargreaves. We are most grateful to you; thank you for your help.

Examination of Witnesses

Ellen Judson and Kyle Taylor gave evidence.

5.29 pm

The Chair: Finally this afternoon, we will hear from Ellen Judson, who is the lead researcher at the Centre for the Analysis of Social Media at Demos, and Kyle Taylor, who is the founder and director of Fair Vote. Thank you for joining us this afternoon.

Q169 Alex Davies-Jones: Thank you both for joining us, and for waiting until the end of a very long day. We appreciate it.

There is a wide exemption in the Bill for the media and for journalistic content. Are you concerned that that is open to abuse?

Kyle Taylor: Oh, absolutely. There are aspects of the Bill that are extremely worrying from an online safety perspective: the media exemption, the speech of democratic importance exemption, and the fact that a majority of paid ads are out of scope. We know that a majority of harmful content originates from or is amplified by entities that meet one of those exceptions. What that means is that the objective of the Bill, which is to make the online world safer, might not actually be possible, because platforms, at least at present, are able to take some actions around these through their current terms and conditions, but this will say explicitly that they cannot act.

One real-world example is the white supremacist terror attack just last week in Buffalo, in the United States. The "great replacement" theory that inspired the terrorist was pushed by Tucker Carlson of Fox News, who would meet the media exemption; by right-wing blogs, which were set up by people who claim to be journalists and so would meet the journalistic standards exemption; by the third-ranking House Republican, who would meet the democratic importance exemption; and it was even run as paid ads by those candidates. In that one example, you would not be able to capture a majority of the way that harm spreads online.

Q170 Alex Davies-Jones: Is there a way in which the exemptions could be limited to ensure that the extremists you have mentioned cannot take advantage of them?

Ellen Judson: I think there are several options. The primary option, as we would see it, is that the exemptions are removed altogether, on the basis that if the Bill is really promoting a systems-based approach rather than focusing on individual small categories of content, then platforms should be required to address their systems and processes whenever those lead to an increased risk of harm. If that leads to demotion of media content that meets those harmful thresholds, that would seem appropriate within that response.

If the exemptions are not to be removed, they could be improved. Certainly, with regard to the media exemption specifically, I think the thresholds for who qualifies as a recognised news publisher could be raised to make it more difficult for bad actors and extremists, as Kyle mentioned, simply to set up a website, add a complaints policy, have an editorial code of conduct and then say that they are a news publisher. That could involve linking to existing publishers that are already registered with existing regulators, but I think there are various ways that could be strengthened.

On the democratic importance and journalism exemptions, I think the issue is that the definitions are very broad and vague; they could easily be interpreted in any way. Either they could be interpreted very narrowly, in which case they might not have much of an impact on how platforms treat freedom of expression, as I think they were intended to do; or they could be interpreted very broadly, and then anyone who thinks or who can claim to think that their content is democratically important or journalistic, even if it is clearly abusive and breaches the platform's terms and conditions, would be able to claim that.

One option put forward by the Joint Committee is to introduce a public interest exemption, so that platforms would have to think about how they are treating content that is in the public interest. That would at least remove some of the concerns. The easiest way for platforms to interpret what is democratically important speech and what is journalistic speech is based on who the user is: are they a politician or political candidate, or are they a journalist? That risks them privileging certain people's forms of speech over that of everyday users, even if that speech is in fact politically relevant. I think that having something that moves the threshold further away from focusing on who a user is as a proxy for whether their speech is likely to deserve extra protection would be a good start.

Kyle Taylor: It is basically just saying that content can somehow become less harmful depending on who says it. A systems-based approach is user-neutral, so its only metric is: does this potentially cause harm at scale? It does not matter who is saying it; it is simply a harm-based approach and a system solution. If you have exemptions, exceptions and exclusions, a system will not function. It suggests that a normal punter with six followers saying that the election was stolen is somehow more harmful than the President of the United States saying that an election is stolen. That is just the reality of how online systems work and how privileged and powerful users are more likely to cause harm.

Q171 Alex Davies-Jones: You are creating a two-tier internet, effectively, between the normal user and those who are exempt, which large swathes of people will be because it is so ambiguous. One of the other concerns that have been raised is the fact that the comments sections on newspaper websites are exempt from the Bill. Do you see an issue with that?

Ellen Judson: There is certainly an issue as that is often where we see a lot of abuse and harm, such that if that same content were replicated on a social media platform, it would almost certainly be within the scope of the Bill. There is a question, which is for Ofcom to consider in its risk profiles and risk registers, about where content at scale has the potential to cause the most harm. The reach of a small news outlet's comments section would be much less than the reach of Donald Trump's Twitter account, for instance. Certainly, if the risk assessments are done and comments sections of news websites have similar reach and scale and could cause significant harm, I think it would be reasonable for the regulator to consider that.

Kyle Taylor: It is also that they are publicly available. I can speak from personal experience. Just last week, there was a piece about me. The comments section simultaneously said that I should be at Nuremberg 2.0 because I was a Nazi, but also that I should be in a gas chamber. Hate perpetuates in a comments section just as it does on a social media platform. The idea that it is somehow less harmful because it is here and not there is inconsistent and incoherent with the regime where the clue is in the name: the Online Safety Bill. We are trying to make the online world safer.

On media I would add that we have to think about how easy it is, based on the criteria in the Bill, to become exempt as a media entity. We can think about that domestically, but what happens when a company is only meant to enforce their terms and conditions in that country, but can broadcast to the world? The UK could become the world's disinformation laundromat because you can come here, meet the media exemption and then blast content to other places in the world. I do not think that is something that we are hoping to achieve through this Bill. We want to be the safest place in the world to go online and to set a global benchmark for what good regulation looks like.

Q172 Alex Davies-Jones: I suppose, yes. Under the current media carve-out, how do you see platforms being able to detect state actors that are quoting misinformation or perpetuating disinformation on their platforms?

Ellen Judson: I think it is a real challenge with the media exemptions, because it is a recognised tactic of state-based actors, state-aligned actors and non-state actors to use media platforms as ways to disseminate this information. If you can make a big enough story out of something, it gets into the media and that perpetuates the campaign of abuse, harassment and disinformation. If there are protections in place, it will not take disinformation actors very long to work out that if there are ways that they can get stories into the press, they are effectively covered.

In terms of platform enforceability, if platforms are asked to, for instance, look at their systems of amplification and what metrics they use to recommend or promote content to users, and to do that from a risk-based perspective and based on harm except when they are talking about media, it all becomes a bit fuzzy what a platform would actually be expected to do in terms of curating those sorts of content.

Kyle Taylor: As an example, Russia Today, until its broadcast licence was revoked about three months ago, would have qualified for the media exemption. Disinformation from Russia Today is not new; it has been spreading disinformation for years and years, and would have qualified for the media exemption until very recently.

Q173 Alex Davies-Jones: So as a result of these exemptions, the Bill as it stands could make the internet less safe than it currently is.

Kyle Taylor: The Bill as it stands could absolutely make the internet less safe than it currently is.

Q174 Kirsty Blackman: You have done a really good job of explaining the concerns about journalistic content. Thinking about the rest of the Bill for a moment, do you think the balance between requiring the removal of content and the prioritisation of content is right? Do you think it will be different from how things are now? Do you think there is a better way it could be done in the Bill?

Ellen Judson: The focus at the moment is too heavily on content. There is a sort of tacit equation of content removal—sometimes content deprioritisation, but primarily content removal—as the way to protect users from harm, and as the threat to freedom of expression. That is where the tension comes in with how to manage both those things at once. What we would want from a Bill that was taking more of a systems approach is thinking: where are platforms making decisions about how they are designing their services, and how they are operating their services at all levels? Content moderation policy is certainly included, but it goes back to questions of how a recommendation algorithm is designed and trained, who is involved in that process, and how human moderators are trained and supported. It is also about what functionality users are given and what behaviour is incentivised and encouraged. There is a lot of mitigation that platforms can put in place that does not talk about directly affecting user content.

I think we should have risk assessments that focus on the risks of harms to users, as opposed to the risk of users encountering harmful content. Obviously there is a relationship, but one piece of content may have very different effects when it is encountered by different users. It may cause a lot of harm to one user, whereas it may not cause a lot of harm to another. We know that when certain kinds of content are scaled and amplified, and certain kinds of behaviour are encouraged or incentivised, we see harms at a scale that the Bill is trying to tackle. That is a concern for us. We want more of a focus on some things that are mentioned in the Bill—business models, platform algorithms, platform designs and systems and processes. They often take a backseat to the issues of content identification and removal.

Kyle Taylor: I will use the algorithm as an example, because this word flies around a lot when we talk about social media. An algorithm is a calculation that is learning from people's behaviour. If society is racist, an algorithm will be racist. If society is white, an algorithm will be white. You can train an algorithm to do different things, but you have to remember that these companies are for-profit businesses that sell ad space. The only thing they are optimising for in an algorithm is engagement.

What we can do, as Ellen said, through a system is force optimisation around certain things, or drive algorithms away from certain types of content, but again, an algorithm is user-neutral. An algorithm does not care what user is saying what; it is just "What are people clicking on?", regardless of what it is or who said it. An approach to safety has to follow the same methodology and say, "We are user-neutral. We are focused entirely on propensity to cause harm."

The second piece is all the mitigation measures you can take once a post is up. There has been a real binary of "Leave it up" and "Take it down", but there is a whole range of stuff—the most common word used is "friction"—to talk about what you can do with content once it is in the system. You have to say to yourself, "Okay, we absolutely must have free speech protections that exceed the platform's current policies, because they are not implemented equally." At the same time, you can preserve someone's free expression by demonetising content to reduce the incentive of the company to push that content or user through its system. That is a way of achieving both a reduction in harm and the preservation of free expression.

Kirsty Blackman: May I just ask one more question, Chair?

The Chair: Briefly, because there are two other Members and the Minister wishing to ask questions.

Q175 Kirsty Blackman: Thanks. On the propensity to cause harm, we heard earlier that a company might create a great new feature and put it out, but then there is a period—a lag, if you like—before they realise the harm that is being caused. Do you trust that companies would have the ability to understand in advance of doing something what harm it may cause, and adequately to assess that?

Ellen Judson: I think there are a lot of things that companies could be doing. Some of these things are in research that they probably are conducting. As we have seen from the Facebook files, companies are conducting that sort of research, but we aren't privy to the results. I think there a couple of things we want to see. First, we want companies to have to be more transparent about what kind of testing they have done, or, if not testing, about who they have consulted when designing these products. Are they consulting human rights experts? Are they consulting people who are affected by identity-based harm, or are they just consulting their shareholders? Even that would be a step in the right direction, and that is why it is really important.

We feel that there need to be stronger provisions in the Bill for independent researcher and civil society access to data. Companies will be able to do certain amounts of things, and regulators will have certain powers to investigate and do their own research, but it requires the added efforts of civil society properly to hold companies to account for the effects of certain changes they have made—and also to help them in identifying what the effects of those changes to design have been. I think that is really crucial.

The Chair: We are playing "Beat the clock". I am going to ask for brief answers and brief questions, please. I will take one question from Kim Leadbeater and one from Barbara Keeley.

Q176 Kim Leadbeater: Gosh, right. I think we are clear that your view is that these two exceptions could potentially do more harm than good. The ideal scenario from your perspective would be to remove them, but again, the challenge is how we balance the freedom of speech issue with protecting the rights of people online who are vulnerable to abuse and harassment. How would you respond to those who say that the Bill risks setting an unwitting precedent for non-democratic countries that would seek to restrict the freedom of expression of their citizens?

Ellen Judson: There is absolutely a risk of over-moderation, and of the Bill incentivising over-moderation, particularly because of the very heavy content focus. Even with illegal content, there is a very broad range of content that companies are expected proactively to monitor for, even when the technical systems to identify

that content reliably at scale are perhaps not in place. I absolutely understand and share the concern about over-moderation.

Our response would be that we should look to strengthen the freedom of expression duties currently in the Bill. At the moment, there is a quite vague duty to have regard to the importance of freedom of expression, but it is not at all clear what that would actually mean, and what would be expected from the platforms. One change we would want would be for rights—including freedom of expression and privacy—to be included in the online safety objectives, and to establish that part of the purpose of this regime is to ensure that services are being designed to protect and promote human rights, including freedom of expression. We think that would be a way to bring freedom of expression much more into the centre of the regime and the focus of the Bill, without having to have those add-on exemptions after the fact.

Kyle Taylor: And it creates a level playing field—it says, "These rules apply to everyone equally."

On the second point, authoritarian—absolutely—but the other area that is really important is fragile democracies. For example, if you look at Hungary, just last week Viktor Orbán said, "You know what you need? Your own media." If we are setting a standard that says it is totally fine to exempt people in politics and media, then for those fragile democracies that control most aspects of information sharing, we are explicitly saying that it is okay to privilege them over others. That is a very dangerous precedent to set when we have the opportunity to set best global standards here with the Bill.

The Chair: Barbara Keeley?

Q177 Barbara Keeley: I have a really simple question. You have touched on the balance between free speech rights and the rights of people who are experiencing harassment, but does the Bill do enough to protect human rights?

Ellen Judson: At the moment, no. The rights that are discussed in the Bill at the minute are quite limited: primarily, it is about freedom of expression and privacy, and the way that protections around privacy have been drafted is less strong than for those around freedom of expression. Picking up on the question about setting precedents, if we have a Bill that is likely to lead to more content moderation and things like age verification and user identity verification, and if we do not have strong protections for privacy and anonymity online, we are absolutely setting a bad precedent. We would want to see much more integration with existing human rights legislation in the Bill.

Kyle Taylor: All I would add is that if you look at the exception for content of democratic importance, and the idea of "active political issue", right now, conversion therapy for trans people—that has been described by UN experts as torture—is an active political issue. Currently, the human rights of trans people are effectively set aside because we are actively debating their lives. That is another example of how minority and marginalised people can be negatively impacted by this Bill if it is not more human rights-centred.

Q178 Chris Philp: Let me start with this concept—this suggestion, this claim—that there is special protection for politicians and journalists. I will come to clause 50, which is the recognised news publisher exemption, in a moment, but I think you are referring to clauses 15 and 16. If we turn to those clauses and read them carefully, they do not specifically protect politicians and journalists, but "content of democratic importance" and "journalistic content". It is about protecting the nature of the content, not the person who is speaking it. Would you accept that?

Ellen Judson: I accept that that is what the Bill currently says. Our point was thinking about how it will be implemented in practice. If platforms are expected to prove to a regulator that they are taking certain steps to protect content of democratic importance—in the explanatory notes, that is content related to Government policy and political parties—and they are expected to prove that they are taking a special consideration of journalistic content, the most straightforward way for them to do that will be in relation to journalists and politicians. Given that it is such a broad category and definition, that seems to be the most likely effect of the regime.

Kyle Taylor: It is potentially—

Q179 Chris Philp: Sorry, Kyle, do come in in a second, but I just want to come back on that point.

Is it not true that a member of the public or anyone debating a legitimate political topic would also benefit from these measures? It is likely that MPs would automatically benefit—near automatically—but a member of the public might equally benefit if the topic they are talking about is of democratic or journalistic importance.

Ellen Judson: Our concern is that defining what is a legitimate political debate is itself already privileging. As you said, an MP is very likely automatically to benefit.

Chris Philp: Well, it is likely; I would not say it is guaranteed.

Ellen Judson: A member of the public may be discussing something—for example, an active political debate that is not about the United Kingdom, which I believe would be out of scope of that protection. They would be engaged in political discussion and exercising freedom of expression, and if they were not doing so in a way that met the threshold for action based on harm, their speech should also come under those protections.

Kyle Taylor: I would add that the way in which you have described it would be so broad as to effectively be meaningless in the context of the Bill, and that instead we should be looking for universal free expression protections in that part of the Bill, and removing this provision. Because what is not, in a liberal democracy, speech of democratic importance? Really, that is everything. When does it reach the threshold where it is an active political debate? Is it when enough people speak about it or enough politicians bring it up? It is so subjective and so broad effectively to mean that everything could qualify. Again, this is not taking a harms-based approach to online safety, because the question is not "Who is saying it?" or "In what context?"; the question is, "Does this have the propensity to cause harm at scale?"

Q180 Chris Philp: The harms are covered elsewhere in the Bill. This is saying what you have to take into account. In fact, at the very beginning of your remarks,

Kyle, you said that some of the stuff in the US a week or two ago might have been allowed to stand under these provisions, but the provision does not provide an absolute protection; it simply says that the provider has to take it into account. It is a balancing exercise. Other parts of the Bill say, "You've got to look at the harm on a systemic basis." This is saying, "You've got to take into account whether the content is of democratic or journalistic importance." You made a point a second ago about general protection on free speech, which is in clause 19(2).

Kyle Taylor: Can I respond to that?

Chris Philp: Yes, sure.

Kyle Taylor: My point is that if there is a provision in the Bill about freedom of expression, it should be robust enough that this protection does not have to be in the Bill. To me, this is saying, "Actually, our free expression bit isn't strong enough, so we're going to reiterate it here in a very specific context, using very select language". That may mean that platforms decide not to act for fear of reprisal, as opposed to pursuing online safety. I suggest strengthening the freedom of expression section so that it hits all the points that the Government intend to hit, and removing those qualifiers that create loopholes and uncertainty for a regime that, if it is systems-based, does not have loopholes.

Q181 Chris Philp: I understand the point you are making, logically. Someone mentioned the human rights element earlier. Of course, article 10 of the European convention on human rights expresses the right to freedom of speech. The case law deriving from that ECHR article provides an enhanced level of protection, particularly for freedom of the press relative to otherwise, so there is some established case law which makes that point. You were talking about human rights earlier, weren't you?

Ellen Judson: We absolutely recognise that. There is discussion in terms of meeting certain standards of responsible journalism in relation to those protections. Our concern is very much that the people and actors who would most benefit from the journalistic protections specifically would be people who do not meet those standards and cannot prove that they meet those standards, because the standards are very broad. If you intend your content to be journalistic, you are in scope, and that could apply to extremists as much as to people meeting standards of responsible journalism.

Q182 Chris Philp: If you are talking about clause 16, it is not that you intend it to be journalistic content; it is that it is journalistic content. You might be talking about clause 50, which is the general exemption to recognise news publishers from the provisions of the Bill. That of course does not prevent social media platforms from choosing to apply their terms and conditions to people who are recognised news publishers; it is just that the Bill is not compelling them. It is important to make that clear—that goes back to the point you made right at the beginning, Kyle. A couple of times in your testimony so far, you have said that you think the way the definition of "recognised news publisher" is drafted in clause 50 is too wide, and potentially susceptible to, basically, abuse by people who are in essence pretending to be news publishers, but who are not really. They are using this as a way to get a free pass from the provisions of the Bill. I completely understand that concern. Do you have any specific suggestions for the Committee about how that concern might be addressed? How could we change the drafting of the Bill to deal with that issue?

Kyle Taylor: Remove the exemption.

Q183 Chris Philp: You mean completely? Just delete it?

Kyle Taylor: Well, I am struggling to understand how we can look at the Bill and say, "If this entity says it, it is somehow less harmful than if this entity says it." That is a two-tiered system and that will not lead to online safety, especially when those entities that are being given privilege are the most likely and largest sources and amplifiers of harmful content online. We sit on the frontlines of this every day, looking at social media, and we can point to countless examples from around the world that will show that, with these exemptions, exceptions and exclusions, you will actually empower those actors, because you explicitly say that they are special. You explicitly say that if they cause harm, it is somehow not as bad as if a normal user with six followers on Twitter causes harm. That is the inconsistency and incoherency in the Bill.

Chris Philp: We are talking here about the press, not about politicians—

Kyle Taylor: Yes, but the press and media entities spread a lot of disinformation—

Q184 Chris Philp: I get that. You have mentioned Victor Orbán and the press already in your comments. There is a long-standing western tradition of treating freedom of the press as something that is sacrosanct and so foundational to the functioning of democracy that you should not infringe or impair it in any way. That is the philosophy that underpins this exclusion.

Kyle Taylor: Except that that is inconsistent in the Bill, because you are saying that for broadcast, they must have a licence, but for print press, they do not have to subscribe to an independent standards authority or code. Even within the media, there is this inconsistency within the Bill.

Chris Philp: That is a point that applies regardless of the Bill. The fact is that UK broadcast is regulated whereas UK newspapers are not regulated, and that has been the case for half a century. You can debate whether that is right or wrong, but—

Kyle Taylor: We are accepting that newspapers are not regulated then.

Q185 Chris Philp: That matter stands outside the scope of the Bill. If one was minded to tighten this up—I know that you have expressed a contrary view to the thing just being deleted—and if you were to accept that the freedom of the press is something pretty sacrosanct, but equally you don't want it to be abused by people using it as a fig leaf to cover malfeasant activity, do you have any particular suggestions as to how we can improve the drafting of that clause?

Kyle Taylor: I am not suggesting that the freedom of the press is not sacrosanct. Actually, I am expressing the opposite, which is that I believe that it is so sacrosanct

that it should be essential to the freedom-of-expression portion of the Bill, and that the press should be set to a standard that meets international human rights and journalistic standards. I want to be really clear that I absolutely believe in freedom of the press, and it is really important that we don't leave here suggesting that we don't think that the press should be free—

Q186 Chris Philp: I got that, but as I say, article 10 case law does treat the press a little differently. We are about to run out of time. I wanted to ask about algorithms, which I will probably not have a chance to do, but are there any specific changes to the clause that you would urge us to make?

Ellen Judson: To the media exemption—

Chris Philp: To clause 50, "Recognised news publisher".

Ellen Judson: One of the changes that the Government have indicated that they are minded to make—please correct me if I misunderstood—is to introduce a right to appeal.

Chris Philp: Correct.

Ellen Judson: Content having to stay online while the appeal was taking place I would very much urge not to be introduced, on the grounds that the content staying online might then be found to be incredibly harmful, and by the time you have got through an appeals process, it will already have done the damage it was going to do. So, if there is a right to appeal—I would urge there not to be a particular right to appeal beyond what is already in the Bill, but if that is to be included, not having the restriction that the platforms must carry the content while the appeal process is ongoing would be important.

Kyle Taylor: You could require an independent standards code as a benchmark at least.

The Chair: Order. I am afraid that brings us to the end of the time allotted for the Committee to ask questions. It also brings us to the end of the day's sitting. On behalf of the Committee, I thank the witnesses for your evidence. As you ran out of time and the opportunity to frame answers, if you want to put them in writing and offer them to the Minister, I am sure they will be most welcome. The Committee will meet again on Thursday at 11.30 am in this room to hear further evidence on the Bill.

Ordered, That further consideration be now adjourned. *—(Steve Double.)*

6 pm

Adjourned till Thursday 26 May at half-past Eleven o'clock.

Written evidence to be reported to the House

OSB01 Professional Publishers Association (PPA)

OSB02 Neil Kendall and others

OSB03 Girlguiding

OSB04 Alliance to Counter Crime Online and the World Parrot Trust (joint submission)

OSB05 Which?

OSB06 Index on censorship

OSB07 Alliance for intellectual property

OSB08 Internet Services Providers' Association (ISPA UK)

OSB09 International Justice Mission (IJM UK)

OSB10 Local Government Association (LGA)

OSB11 Russ Elliott

OSB12 SWGfL - Safety & Security Online

OSB13 Action for Primates and Lady Freethinker

OSB14 Association of British Insurers (ABI)

OSB15 Microsoft

OSB16 Age Verification Providers Association

OSB17 techUK

OSB18 UK Interactive Entertainment (Ukie), the trade association for the UK's video games industry

OSB19 Centre for Media Monitoring

OSB20 Save Online Speech Coalition

OSB21 Bumble

OSB22 Professor Clare McGlynn

OSB23 Antisemitism Policy Trust