HOUSE OF COMMONS

Petitions Committee

Oral evidence: Tackling Online Abuse, HC 766

Tuesday 23 November 2021

Ordered by the House of Commons to be published on 23 November 2021.

Watch the meeting

Members present: Catherine McKinnell (Chair); Martyn Day; Christina Rees.

Questions 77-130

Witnesses

I: Dr Nicholas Hoggard, Lawyer, and Professor, and Penney Lewis, Commissioner, Law Commission.

II: Theo Bertram, Director of Government Relations and Public Policy for Europe, TikTok, Katy Minshall, Head of UK Public Policy, Twitter, and Rebecca Stimson, UK Head of Public Policy, Meta.

Written evidence from witnesses:

- Katy Minshall, Head of UK Public Policy, Twitter

Examination of witnesses

Witnesses: Dr Nicholas Hoggard and Penney Lewis.

Q77 **Chair:** Thank you very much for coming to talk to us today about tackling online abuse. This is our Committee's third session on this issue in recent weeks, and so far we have heard from people focusing on the regulatory and technological responses to online abuse, but we are very much looking forward to hearing from you today about the work of the Law Commission, how the law is working in terms of reviewing hate crime and online communications, and what role that might play in helping to address the behaviour that we know the petitioners are very concerned about and is the reason why we are undertaking this inquiry. Before we launch into our questions, will you please introduce yourselves and explain briefly what your involvement is in this work? Who would like to go first?

Penney Lewis: I'll start. I am Penney Lewis and I'm the criminal law commissioner. The particular project that we are here talking about today is one of a number of projects that I am leading at the Law Commission in the area of criminal law. Maybe I can just explain the Law Commission very briefly. It is an independent agency that provides recommendations on law reform to Government and Parliament in order to make the law modern, simple, clear and cost-effective. In addition to the project that we are talking about today on modernising the communications offences, we are undertaking work on hate crime, which I understand the Committee may also be interested in. The other project I want to mention is a project on intimate image abuse, some of which, obviously, happens online and is therefore within the scope of your inquiry. I'll hand over to Nick.

Dr Hoggard: Hello, I'm Nicholas Hoggard. I am a lawyer at the Law Commission, and I was the lead lawyer for the online abuse project, which culminated in the report "Modernising Communications Offences".

Chair: Thank you. Are you happy if we refer to you as Penney and Nick?

Penney Lewis: *indicated assent.*

Dr Hoggard: Of course.

Q78 **Chair:** Thank you. Could I come back to you to start, Penney? I will come to you afterwards, Nick. You have very helpfully summarised some of the work that you have been doing, but could you briefly set out what you see as the key points from the work that you have been doing to review criminal offences relating to online communications?

Penney Lewis: Thanks. That is a really useful question. We have taken a look at the existing communications offences, and our concerns are about the vagueness of those offences. They use terms like "grossly offensive", "indecent" and "obscene", and different people understandably have different views on what constitutes offensive material and what the

borderline between offensive and grossly offensive is, for instance. We are also concerned that any criminal offence in this area, including the existing communications offences, is, by definition, an interference with freedom of expression. Some interferences with freedom of expression are justified; it is not an absolute right. But we have taken the view, in focusing our recommendations on harm, that that is a more proportionate interference with freedom of expression. In other words, the kinds of reasons that justify an interference in freedom of expression are generally focused on harm. They include risks to national security or the rights and freedoms of others, for example.

Those are the two main issues that we have tried to deal with: the concern about vagueness or over-breadth; and a renewed focus on harm, whether that is in relation to our general harm-based offence or some of the more specific offences that we have recommended—for example, on cyber-flashing, threatening communications or sending flashing images to a person with epilepsy with intent to provoke a seizure.

Chair: Nick, do you want to add to that?

Dr Hoggard: Only to say that one of the biggest problems with the existing law and the way it works is that too often, because it is proscribing or banning content based on categories—as Penney said, you have to define something as grossly offensive or indecent, for example—there are many types of genuinely harmful content that stakeholders have told us about that just don't fit comfortably into those categories. One of the overriding objectives has been to ensure that where harm actually falls, where there is the real potential for harm, the criminal law is able to engage with it. We are engaging with the real potential for harm on the one hand and, as Penney rightly said, we are trying to ensure protection for freedom of expression on the other.

Q79 **Chair:** Do you have any practical examples to illustrate what you are saying, in terms of where content doesn't fit in?

Dr Hoggard: Of course. It might be helpful to dwell a little on the existing law. Under the Communications Act 2003 and the Malicious Communications Act 1988, in order to engage those offences, the material would need to be grossly offensive or indecent, for example. One of the difficulties with having to constrain the scope of the offence by reference to category is that you miss out on those context-dependent harms, which is to say that where something is harmful only by virtue of the context in which it exists, the criminal law will not bite.

A good example of that is that we have heard a number of examples whereby an abuser has managed to find a survivor—a former partner who had maybe escaped to try to find some form of sanctuary or hiding—and sent that person a picture of their new front door, say, of a street sign or of a family member. It is difficult to describe those as grossly offensive or indecent on the face of it. In some cases, it might even be difficult to describe them as objectively menacing. For example, in the case of a picture of a child, someone's offspring, it is hard to say definitively that

that is menacing. Clearly, however, within a context, it can cause very serious distress.

Another example might be where you take advantage of someone you know to suffer from PTSD. You might send them a sound, for example, which you know will trigger their PTSD. Again, under the current law, it is very difficult to say that that particular sound would fit one of the proscribed categories. That was partly why we thought we needed an offence that was more context-dependent, in order to accommodate those harms.

Q80 **Chair:** Is that part of the proposals you have put forward for what you would create? Will you describe technically how you would change the law to take context into account?

Dr Hoggard: Penney mentioned some of the offences that we are proposing. In this context, it might be helpful to talk you through the more general offence—what we call the harm-based offence. That offence is designed to replace the existing section 127(1) of the Communications Act and to replace the Malicious Communications Act. There are some other offences, which we can talk about, that draw on some of the features of the existing law, such as the existing criminalisation of false communications and of threatening communications. We have proposed new and, we submit, better offences, but the general offence, which I will talk about now, is the harm-based offence.

That would be an either way offence, by which we mean that it could be triable either in the magistrates court or in the Crown court. It is the same as the Malicious Communications Act in that respect, but slightly more serious than the Communications Act. That is partly because we are criminalising at a different threshold from the Communications Act.

The elements of that offence are that it would be an offence if the defendant sent or posted a communication that was likely to cause harm to a likely audience—that is, those who were likely to see or encounter it. I stress here that "likely" does not just mean a mere risk or a remote possibility; it means a real or substantial risk that someone was likely to see it or that harm would ensue—I will come on to the definition of "harm" in a second. That is the first element.

The second element would be that the defendant, the person sending or posting the communication, must have intended to cause harm to those who are likely to see it. That is a big difference, certainly from the Communications Act, which has no intention requirement at all. It would be an offence for you to send something that was grossly offensive or indecent—it does not matter what you intended, so long as you at least intended to send the thing in the first place. So, that is a big change from the current law.

Further, the defendant must be proven by the prosecution to have sent the communication without reasonable excuse. A lot of people think that this is a defence; it is not a defence. This is not something the defendant

has to raise and rely on; the prosecution has to prove to the criminal standards—that is, beyond a reasonable doubt—that the communication was sent without a reasonable excuse. "Reasonable excuse" is quite a broad church, and we do not want to constrain it. The constraining factor is the word "reasonable". However, we do say it is important for the purposes of freedom of expression to have mind to whether the communication either contributed to, or was intended as a contribution to, a matter of public interest. That is drawing straight from article 10 of the ECHR freedom of expression jurisprudence.

I said earlier, and I will pause after this for questions, that we had said that the communication must be likely to cause harm. What we mean by harm is psychological harm amounting to at least serious distress. I want to dispel one thing: this is not a new, slightly recondite thing that the Law Commission is proposing. This is a well-established harm threshold in the criminal law. Many offences rely on the notion of distress—harassment, stalking and the public order offences, for example. Offences also rely on the concept of "serious distress"; coercive and controlling behaviour is one example. The Protection from Harassment Act 1997 in respect of stalking also talks of serious distress. This is not an unusual concept in the criminal law.

What we are trying to do is to place the seriousness threshold within a schema and hierarchy of law. Harassment, which requires a course of conduct, requires proof of alarm or distress, or intention to cause alarm or distress. Similarly, the Malicious Communications Act 1988 includes distress. By necessary implication, we are saying that the threshold is more serious than that. It is not about hurt feelings. It is not about mere emotional distress. It is more serious than that. That is what we mean by harm in this context.

Q81 **Chair:** Thank you. Obviously, there is huge interest in this and how it will work in practice. From the report the Petitions Committee undertook in 2019, there appeared to be significant confusion among both the public and the police about how the law applies to online behaviour, and clearly the purpose of your review is to try to clarify that. Are you confident that the proposals you are making will help to change that? Will we see an increase in the ability to pursue this sort of harm?

Penney Lewis: I think we are confident that these recommendations are better targeted than the existing ones. I don't think it is as simple as saying that there will be more prosecutions for online abuse. At present, there is criminalisation where there should not be: consensual sexting between adults is currently criminalised and should not be. Cyber flashing is sometimes criminalised and should be. We need to focus not just on a numbers game of, "How do we increase the number of prosecutions for online abuse?", but, "How do we make sure the right prosecutions for online abuse take place?"

That is what these recommendations are designed to achieve. The focus on harm ensures that, if someone sends a friend whose sense of humour they know well a grossly offensive joke, where no one is likely to be

harmed as a result, it should not be a criminal offence. On the other hand, if you send someone a grossly offensive joke, knowing that the recipient is likely to be harmed, it should be. "Grossly offensive" is not really doing the right work here. What is important in relation to the criminalisation of communication is harm, and that focus will ensure that those who are likely to be harmed are significantly better protected, while freedom of expression is also better protected.

Q82 **Chair:** We have also seen some examples of individuals who have been convicted for sending abuse to England footballers following the European championship final earlier this year. Overall, do you think that incident helps to illustrate the case for reforms in this area? How do you think those current convictions will be impacted in terms of the proposals you are bringing forward?

Penney Lewis: That is a really good question. It does depend on the content of the communication, but if you take, for example, a racially abusive tweet that contains a racial slur, it is likely to meet both thresholds. It is likely to constitute a criminal offence under the current law, because it is likely to be considered grossly offensive, and it will continue to be criminalised under our recommendations, because it is likely to cause harm to someone who is likely to encounter it—either the footballer or the people who follow them on Twitter. Those people are likely to encounter it if it is directed at that person.

In some ways, I think these recommendations don't necessarily change the way that racial abuse directed at footballers is prosecuted. It is possible, as Nick mentioned, that by having an either way offence, which will have a maximum sentence similar to the Malicious Communications Act offence—two years—some circumstances that can only be prosecuted at the moment under the Communications Act will be folded into the more serious offence with a more serious penalty. It is possible that some abuse that at the moment is prosecuted under section 127 will be treated more seriously by our recommended offences.

Q83 **Chair:** One final question on this: do you see a risk of online abuse becoming a very subjective thing, where, for example, public figures become very accustomed to levels of abuse being directed at them online and those incidents potentially wouldn't fall within the parameters of the proposed legislation, because in many ways people become desensitised because of the amount of abuse that circulates? However, somebody who is not accustomed at all to being abused online might be caused incredible distress and harm from receiving just one abusive message. Is there any concern that it could continue to be normalised? That is, I think, one of the big concerns that many people have in the online sphere—that we are almost normalising a level of abuse that is harmful, but in a very low-level yet pernicious and damaging way, and that it would perhaps not fall within this definition of the new offence that you are describing.

Dr Hoggard: The first thing to say is that we have received quite mixed evidence. Indeed, we have received a lot of evidence to suggest that, far

from people becoming desensitised, actually the ongoing abuse affects them far more. They might be able to brush off one or two messages, but as it increasingly starts to dominate their lives, it has an increasingly adverse effect on their ability to conduct their day-to-day activities as they would wish. We have seen evidence of people no longer being able to engage online and have social media accounts, which is at the more minor end on the impact it can have on their lives. None the less, given how we live, that is not insignificant.

I think that, yes, some people may well be desensitised to it, but not everyone will. I think it would be wrong to assume that this would normalise that sort of behaviour. Fundamentally, what we are saying is that it is no longer the case that you need to think, "Is this going to be deemed by somebody to be grossly offensive?" Instead, the criminal law will be saying that it will be wrong for you to do something whereby you intend to cause somebody harm, so it is shifting the focus towards those who are going out of their way to cause somebody else harm. Because of that, it is sending a very clear message that, far from normalising behaviour that might cause harm, it is saying that we really should not be doing this.

If I can, I want to dwell briefly on the notion of subjectivity, because that is not particularly well understood by many people we have spoken to. The conceit is that the current law is objective and what we are proposing is subjective. I actually don't think that is the case; I think it is the reverse. The current law, especially in the case of gross offensiveness but also of indecency—if you think about cultural norms, what constitutes indecency may well be different—it is very hard to find consensus on what constitutes gross offensiveness. If a comedian makes jokes about a particular community, some people might say that is innocent humour, some people might say that is important public interest provocation and some people might say it is grossly offensive, such that it should be banned, or even criminal. It is not clear that the criminal law is meaningfully able to mediate those different opinions.

Inevitably, what happens is that a jury or a magistrate will have to draw on their own subjective view about what they consider to be grossly offensive or indecent. We are saying that, actually, this is an objective inquiry under our proposed offence, to say, do we think that in this context, given the nature of the person who is likely to see it, harm was likely? Does it mean that if somebody has very thick skin and is not easily harmed, the answer to the question, "Was harm likely?" might be no? Inevitably, yes, but I go back to my earlier point that by focusing on intent and those bad faith actors, and by recognising that repeated harm and pile-on harassment can cause real harm, we think we have the balance right.

Q84 **Chair:** The other way of looking at it is, if you have multiple perpetrators, where many people are abusing you online and the cumulative effect is serious harm, who is the perpetrator? Is it the final one who breaks the camel's back—the final harmful communication that puts you into a very

distressed situation? But many perpetrators have brought you to that point. How would the law deal with that?

Dr Hoggard: That is pile-on harassment, where you have lots of people, potentially unco-ordinated, who just see that something is happening and jump on the bandwagon. The criminal law is not a very good tool for dealing with pile-on harassment generally, simply because the scale is so enormous that to create a single offence of pile-on harassment would not work; we would not be able to constrain it in practical ways. Data of the scale in some cases of abuse directed at public figures suggests that they are receiving hundreds of abusive messages a minute. It is truly astonishing. That is not something that the criminal law is able to deal with by means of an individual offence. But that is not to say that we don't recognise the real harm that can come from pile-on harassment.

However, there are existing offences that can do more. There is a co-ordinated harassment offence under the Protection from Harassment Act, which can deal with the early stages of somebody co-ordinating a pile on. That means that a course of conduct can be attributed to the actions of the person who tries to instigate it. That can help in those early stages. Similarly, there is no reason to suggest or to suppose that the harm-based offence that we propose could not bite at that stage in the right context.

To the specific point of at which point the offence would bite if the serious distress only manifests later on, there is no reason to suppose that the offence should not be used where you are sending a message intending to cause harm and you may know that the person has already received a number of messages, and that forms part of your intention. There is no reason why the offence should not be engaged at that stage.

Chair: Each person sending abuse should be conscious that they might be the person that tips it over the criminal threshold of abuse.

Penney Lewis: I think there is also an important distinction here. Our recommendation focuses on likely harm, not actual harm, so the prosecution will not have to prove that this person—or a particular person—was harmed. Obviously, if someone has been harmed, that will be evidence that might be relevant to, for instance, whether harm was likely. But there are circumstances where someone has actually been harmed when harm was not likely and the defendant cannot be expected to have foreseen the potential for harm. We think it is important not make this an actual harm offence.

Similarly, in the context of pile-on, what the sender of a particular message would have to be conscious of is: is this likely to cause harm? It is not: am I going to be the person who tips this recipient over into the threshold of being seriously distressed? It is: is this message that I am sending right now likely to cause this recipient harm? That is the right inquiry, because this is an offence which is either committed at the point of sending or not committed. What we do not want is to create an offence that relies on actual harm and then, maybe a month after it is sent, suddenly the defendant becomes guilty of the offence because someone

has now been harmed. We need defendants to be in a position to be able to predict: am I going to breach criminal law here?

Q85 **Christina Rees:** Welcome, Penney and Nick. To dig a bit deeper, how would your proposed reforms affect how much abusive content online could be considered illegal?

Dr Hoggard: As Penney mentioned earlier, this is an important question because it is often easy to assume that what we are doing is simply layering on more offences and trying to make more things illegal online. That is not what we are doing. We are trying to make the criminal law work better in terms of criminalising only those communications that have the actual potential for harm and where there was culpability in sending it—in other words, somebody intended to cause harm. At the moment, the criminal law criminalises either where there is no potential for harm or where there is no culpability because somebody did not intend harm or because what they were doing was entirely reasonable.

We are, I think, refocusing the criminal law on the right areas. I think it is fair to say that if the Communications Act 2003 were prosecuted to its fullest extent wherever it cropped up, that would criminalise an extraordinary amount of content online that is neither culpable nor harmful. Sending an indecent image over a public electronic communications network is an offence, so consensual sexting between adults is, on its terms, a criminal offence, but of course there is no public interest in prosecuting that. Even saving an indecent image of yourself to your own online storage—your iCloud, for example—constitutes a criminal offence on the face of the Communications Act. That is, on its terms, a very, very broad offence. So I suppose it would be easy to say that we are narrowing the scope of the criminal law, but what we are doing is more nuanced than that; we are refocusing it to track more closely both harm and culpability.

Christina Rees: Penney?

Penney Lewis: I think Nick has pretty much said it all. I will leave it there.

Q86 **Christina Rees:** Okay. How would you see your new proposed offences—the harm-based communications offence in particular—overlapping or interacting with relevant existing offences?

Penney Lewis: Do you have particular offences in mind? Are you thinking about the hate crime framework or intimate image abuse? I am trying to think what else—

Chair: Maybe the hate crime framework.

Christina Rees: Hate crime.

Penney Lewis: The way in which hate crime law works is that there are, strictly speaking, four different regimes. There are aggravated offences—for example, assault and racially aggravated assault. Those do not exist

for the communications offence at the moment, and that is something that we are looking at in our hate crime project. The regime that currently applies to the communications offences is enhanced sentencing, which means that the judge, or indeed the bench of magistrates, is required in certain circumstances, where hostility towards one of the five protected characteristics—race, religion, sexual orientation, disability and transgender identity—has been proven, to increase the sentence. But there is no higher maximum sentence, so it is not like an aggravated offence, where there is actually a higher maximum sentence in circumstances of that aggravation.

The advanced sentencing framework, assuming that it remains—we consulted on a proposal to retain it—will also apply to any new communications offences. Similarly, hostility—or hate crime, to use the more commonly used term—will be available in sentencing people for the new communications offences in the same way as it is currently. If you think about, for instance, racial abuse of a footballer, at the moment that would probably be prosecuted as a grossly offensive communication, and the sentence might well be enhanced on the basis of hostility towards the footballer's race. Similarly, using the harm-based offence, enhanced sentencing would allow you to increase the sentence if hostility towards the footballer's race is proven.

I don't think we envisage a significant change in way in which the regimes interact. What we have done, though, is consult on whether there should be more aggravated offences, and whether aggravated offences should apply to all protected characteristics—at the moment, they apply only to race and to religion.

The other area that I will briefly mention—I don't think it is really the subject of this inquiry—is stirring up hatred online. Those are very serious offences; they have penalties of seven years' imprisonment attached to them, and they currently require the consent of the Attorney General before they are prosecuted—we consulted on whether that should be the Director of Public Prosecutions. We are looking at whether there should be a harmonising of those offences so that they are all similar. At the moment, it is easier to prosecute someone for stirring up racial hatred than it is for religious hatred, for instance. We also consulted on whether there should be stirring up hatred offences in relation to disability and transgender identity, as well as sex or gender.

Q87 **Chair:** Would you mind also explaining how your proposed new offences would interact with existing offences relating to harassment and threats to the person?

Penney Lewis: Sure. I will start and then maybe Nick can come in. In terms of threats, in addition to the existing communications offences, which criminalise threatening or menacing communications, there is also a threats to kill offence in the Offences against the Person Act 1861. We recommended an offence of sending a threatening communication that would include threats to rape, which is something that we had earlier recommended in our offences against the person project. In that project,

we can really make recommendations only about sending a communication, so not a face-to-face threat to rape, but we have in that previous project recommended an offence that would cover all threats to rape.

In addition to threats to rape, threats to cause serious harm is another recommendation that we have made in our offences against the person project, and which we have reiterated in this project. In this project, we have also included threats to cause serious financial harm. We have recommended a specific threats offence and I think that would sit well with the existing threats to kill offence. Of course, if Parliament were minded to implement or create the threats offence that we recommended in our 2015 "Offences against the Person" report, that recommendation would also cover in-person threats to rape or cause serious harm, which would obviously be more comprehensive than the "sending a communication that is threatening" offence that we have recommended here.

Q88 **Chair:** Would somebody get around it if they said, for example, "I hope you are raped. I hope you get what's coming to you"? Would that get around the proposed offence that you are talking about?

Penney Lewis: It might not be a threat, but it certainly could still fall within the harm-based offence. It is certainly possible that some communications that come close to being a rape threat—I suppose the most infamous example is the threat sent to MPs along the lines of, "I wouldn't even rape you"—might be a threat. There are circumstances where it could be—one way that it might well be is if it was accompanied by a picture of the person's house—but it would almost certainly be something that was likely to, and intended to, cause serious distress. So it would be criminalised, but it is possible that it would not fall within the threats offence.

Q89 **Chair:** Would posting the location of somebody's address online, for example, with a picture of their street and their house—somebody that you are clearly protesting against—fall within the existing legislation as threatening or within your proposed recommendations as threatening or as a criminal act?

Dr Hoggard: It is a good question, which highlights one of the difficulties with the existing law: I do not think that the mere fact of posting personal or private information—in this case, the location—could meaningfully be described as grossly offensive or indecent, and so on. So I think this case demonstrates well the need for a more context-specific offence, such as the harm-based offence. As long as the person whose private information was disclosed was likely to see that it had been—in many of the cases we have heard about in the context of private information, the victims have been well aware of what is being posted about them in relation to their private information online—it is perfectly possible that the harm-based offence would capture that behaviour.

On the issue of whether it would fall within the threatening communications offence, given that the offence is complete at the point of sending, it would need to be understood on its face as a threat or, at least, it needs to be clear that the object of the threat—the person intended to appreciate the threat—would understand it as such. It could then well be in the scope of the offence, so I would not want to say that it definitely would not be, but it would depend on the facts and how that information was introduced. None the less, I think it is perfectly possible that it could be caught within the more general harm-based offence.

Q90 **Chair:** Okay—that is being proposed.

Dr Hoggard: Yes.

Q91 **Martyn Day:** We have heard how a number of specific communities, people from them or those with specific characteristics can be subject to more online abuse than other people. How would your recommendations take that into account?

Penney Lewis: Do you mean in terms of a higher prevalence?

Martyn Day: Yes.

Penney Lewis: It is difficult for the criminal law to respond in a sort of systemic fashion like that. One of the ways in which it does, to a certain extent, is through hate crime law. Hate crime law makes something a more serious version of an existing offence, such as assault or sending a grossly offensive communication, for instance, in circumstances where there is hostility towards a protected characteristic. That doesn't mean that hostility towards other protected characteristics where there's not prevalent offending can't be taken into account by the criminal law. If you take, for example, the racially motivated murder of Stephen Lawrence, the hostility towards his race was a significant factor taken into account in sentencing. Racial hostility is a recognised hate crime characteristic, or race is a recognised hate crime characteristic.

In contrast, if you think about the terrible murder of Sophie Lancaster, who was murdered apparently because she was a goth and because of hostility towards her goth identity, that is not a protected characteristic. There is a significantly smaller number of crimes motivated by hostility towards what are known as members of alternative subcultures, but none the less that hostility was taken into account by the sentencing judge.

So the most prevalent kinds of targeted characteristics and those where there is an evidence base for the causing of additional harm to members of that group—for instance, members of racial minorities are harmed by the knowledge that they may be targeted on the basis of their race—are recognised explicitly in hate crime law, but the law is flexible enough to also recognise forms of hostility that are perhaps less prevalent but are still pernicious. Does that help?

Martyn Day: Most certainly. Do you want to add anything, Nick?

Dr Hoggard: I realise that this is not about how the criminal law responds specifically to individual groups, but I just want to reiterate that one of the driving forces behind the change here was the fact that the criminal law was not working well in certain very specific contexts. The context of violence against women and girls is one example of where the existing criminal law has not been working well. I have spoken already about some of the ways in which messages don't necessarily fall comfortably within the existing categories. We know that domestic abuse and coercive controlling behaviour uses communication. A very high percentage of the way that behaviour manifests itself is through communication, be it online or through text message or other means. Because of that, we need to make sure that the criminal law is working well so that it can assist with the existing range of offences designed to combat domestic abuse and coercive controlling behaviour. Like harassment and like stalking, we also have communications offences that are able to deal with those problems. A number of our stakeholders, such as Refuge, the End Violence Against Women Coalition and Women's Aid, recognise the problems in the existing law and have supported the move to a more context-specific response.

Q92 **Martyn Day:** That leads me on to the next point that I was going to raise. This Committee produced a report in 2019 that raised concerns that the hostility test used for hate crime offences didn't often capture the actual nature of the online abusive content being faced by disabled people. Has your work considered how that type of problem could be dealt with?

Penney Lewis: Absolutely. One of the proposals that we consulted on in the hate crime project, the report for which I hope will be published in the coming weeks, was to amend what is known as the motivation limb of the hostility test in order to better reflect the reality of disability hate crime.

The hostility test has two limbs: either the prosecution proves that the offence was motivated by hostility towards a protected characteristic or the offender must have demonstrated hostility towards a protected characteristic. The most common form is to use the demonstration limb, and usually that is through some form of slur—a racial slur or a homophobic slur, for example.

In the context of disability hate crime, what disability stakeholders have told us—it sounds to me like they have told you as well—is that that fails to capture the kind of disdainful, contemptuous attitudes that characterise offending against people with disabilities. Some of that offending may be based on vulnerability as a motivation, but much of it, while not overtly hostile, is contemptuous in a way that really should fit within the hostility test but does not appear to.

So what we consulted on was to change the motivation limb, so that instead of requiring the prosecution to prove motivation by hostility, it would be motivation by hostility or prejudice. We think the inclusion of prejudice is a good way of capturing that kind of disdainful contempt that we see, for example, in so-called mate crimes, where people might be befriended online and then humiliated in a way that they are not easily able to grasp but is none the less hurtful, and indeed harmful, without

moving toward a much broader approach where anything motivated by a protected characteristic should constitute a hate crime. In other words, if you were to attempt to exploit someone's disability simply because they are an easier target than someone else, that would not necessarily fall within this category of prejudice, but we are also very conscious that much of that kind of exploitative targeting actually falls within this category of disdainful conduct, because what it represents is an analysis that this person is somehow less worthy of respect, less worthy of protection, because they are somehow lesser than the perpetrator because of their disability.

Much of that kind of targeting of disabled people is prejudicial targeting within this hate crime analysis. We have provisionally proposed that expansion of the motivation limb, which we think will not make an enormous difference in relation to other characteristics, which are more clearly characterised by hostile offending, but will make a significant difference to the targeting of disabled people on the basis of that sort of contemptuous attitude.

Q93 **Chair:** Although I am sure we have many questions we would like to put to you, we are out of time. If there is anything you feel that we have not covered or asked you about that would be helpful, now is your chance, or are you happy that you have managed to convey what you think will be helpful to this inquiry?

Penney Lewis: I am happy. Nick?

Dr Hoggard: I am happy too.

Chair: Thank you very much. We are grateful to you for giving the time today to help this inquiry. We will now suspend for a few minutes, as we are taking evidence from the next panel virtually.

Examination of witnesses

Witnesses: Theo Bertram, Katy Minshall and Rebecca Stimson.

Q94 **Chair:** Thank you for appearing today to answer questions on this important subject. The companies you represent play an enormous role in the online ecosystem, where the kind of behaviour we have been looking at and discussing in this inquiry takes place, so we are really looking forward to hearing how you currently respond to the problem of online abuse, and how the Government's Online Safety Bill will or might affect this. Before we get under way, will you briefly introduce yourselves, please?

Katy Minshall: My name is Katy Minshall, and I am head of public policy for Twitter in the UK.

Theo Bertram: Hi, I'm Theo Bertram and I am responsible for public policy across the region for TikTok.

Rebecca Stimson: Hello, I'm Rebecca Stimson, head of public policy for Meta in the UK.

Q95 **Chair:** For those watching, could you explain who Meta are?

Rebecca Stimson: I appreciate it is a new thing for all of us. Meta is now the company name for what used to be Facebook, which oversees Facebook, WhatsApp and Instagram—the apps that everyone is more familiar with.

Chair: Okay, so you are here to speak on behalf of Facebook, Instagram and WhatsApp.

Rebecca Stimson: Correct.

Chair: Thank you.

Your platforms' community standards already prohibit bullying, hate speech and other abusive behaviour, but many of the witnesses we have heard from have raised concerns that these standards are not well enforced. Would you explain what you are doing to address that? Would you like to go first, Rebecca?

Rebecca Stimson: Sure, having unwisely jumped in. Thank you so much for the question and for giving me the chance to be here today to talk about these issues.

You referenced community standards, which are the rules we set out on what we do and do not allow on Facebook and Instagram. Those are made public; we also publish quarterly transparency reports, which show how well we are enforcing against those rules. If you look at those reports on things like hate speech, bullying and harassment, you will see that we have been continuously improving, quarter on quarter, in how we enforce against those policies. Particularly important is how much of that content we find and remove ourselves, proactively, through our detection systems, rather than relying on people reporting.

One of the figures that we think is most important is around prevalence, because it helps give people a sense of how much of this kind of content there is on our platforms. To give the Committee a sense of scale, in terms of hate speech on Facebook, it is roughly about three pieces of hate speech per 10,000 pieces of content viewed, and on Instagram, for bullying and harassment, it is about 15 per 10,000. Obviously, we continue to invest to improve those figures as much as possible.

Katy Minshall: Thank you for inviting me to give evidence this afternoon. Much like Facebook, the majority of the abusive tweets we detect, we take down proactively, and that is a very different position to the one we were in last time we gave evidence to this Committee, in 2018. Back then, we were still pretty reliant on people reporting these issues to us, which meant that there was a huge burden on victims of abuse. Now, as I say, the majority of the abusive tweets we take down, we detected using technology.

Much like Facebook, we have thought about different ways to measure how effective we are at identifying and removing this kind of violative

content. One of the metrics we have started using more recently is impressions. An impression is the instance of a tweet appearing in someone else's Twitter, and our research shows that impressions on violative tweets account for less than 0.1% of impressions overall globally, so this is a case of a minority of bad actors with abhorrent views trying to undermine the experience for the majority.

However, what I would probably also say in response to your question is that our terms and conditions—our rules—are not the only way in which we assess our effectiveness when it comes to online safety. What we have been trialling much more over the past few years is nudges: behavioural nudges, safety features, controls, and sharing the results of those experiments publicly.

Just to give you a flavour of that—I am sure that Theo will want to come in—one of the trials we have been running over the past year is a prompt that nudges people to consider the language they are using, such that if you tweet something out and I go to reply to you, and it looks like the reply I want to send could be offensive or might break our rules, I might now see a prompt saying, "Are you sure that you want to send this? It looks like it could cause harm." Actually, we have found that in a third of cases, people do not send that tweet, or they change the tweet they were going to send. I think that is quite illustrative: that is potentially a third of replies that could have existed that now don't. That is why, in terms of measuring effectiveness, we consider all of those different aspects as well.

Theo Bertram: Thanks also for having me today on this important issue. I am going to try and avoid repeating what Twitter and Facebook—sorry, Meta; I've got to get used to that—have said. We also have a set of community guidelines; we have effective enforcement. I know that Committees like this have heard from companies like ours before, giving those data, and I can tell you that 94% of that content is removed before a single report, and 87% of it before a single view. You do not have to take our word for it: the European Commission did its own analysis and test of companies when it came to hate speech, and we are pleased that we came out top in that survey, despite being the youngest of the companies being tested.

However, we approach this with a tone of humility. We can talk about community guidelines and we can talk about enforcement, but I think product design is really key here. Twitter has just given an example, and we have a similar thing, which is what we call a kindness prompt. When someone is going to leave a comment, we do that same thing: if we can detect that the comment might have something in it that might be offensive or harmful, we will prompt the user to reconsider that tweet, and in four out of 10 cases, they will choose not to post that tweet—sorry, that comment. Steps like that can be helpful.

We go further in a couple of other areas. For everyone under 16 on our platform, we have made our product private by default. That means that when they are on the platform, the only people that are following their accounts are those that they choose, or no one at all. We have also made

it so that the platform will not allow one person to send another person unsolicited direct messages, so you can only get a direct message from someone who has already agreed to have a direct message with you. In addition to that, we don't allow under-16s to have direct messaging at all.

I think that that combination of community guidelines, which is something you are familiar with, the enforcement process, and constantly getting that better, and product design is a really important and key part of this. Education on the platform is key as well. Campaigns to show that this is not normal behaviour is something that we can all do a better job on.

Q96 **Chair:** I think we can all agree with that. We welcome any of the changes that have been made to improve the experience of people online, which is, from the evidence we have heard, too often a hate-filled and abuse-filled experience. While there have been some improvements made, it has to be said that they haven't come from the platforms themselves without being demanded and called for. So, obviously, we are now—

Theo Bertram: I don't think that's fair. For example, at the start of this year we made it privacy by default. We were under no obligation to have done that.

Q97 **Chair:** No obligation, or nobody called for it, because I said it hasn't come without being demanded and called for?

Theo Bertram: I don't think the other platforms have done it, so I think we have done that because we thought it was the right thing to do.

Q98 **Chair:** Okay. I think most people who experience some of the hatred and abuse online would disagree with you.

The legislation is now coming down the track. The Online Safety Bill is in its draft form, and it would make platforms like yours liable for enforcing your own rules on content that is legal but harmful to adults. I would be interested to know how you expect this to affect your current enforcement practice, in terms of community standards. From what you have seen from the draft, would it bring within it any action beyond what you have just outlined? Would you like to come in, Rebecca?

Rebecca Stimson: Sure. On Theo's last point, we have made Instagram private by default for younger users as well. On the Online Safety Bill, I was so interested in the session before this one with the Law Commission. I thought a lot of what they said showed the complexity of addressing some forms of harm, by focusing on things like intent—did someone do something knowingly?—and citing subjective terms like "distress" and so on. I think a lot of that is true for the proposals that the Government have put forward. I think all of us are looking for more detail and more granularity as this process goes forward.

On the question of liability, I think Theo is right. Meta might have a new name, but we are one of the older companies here. We have had these public standards and enforcement teams and have invested very heavily for a long time without any need for regulation, so I do think Theo is right.

On the impact of the Online Safety Bill, it is hard to say at this very early stage. I would caution that liability is a very complicated issue when talking about a user-to-user service. You are talking about, as you have mentioned, regulating the legal speech of individuals online. I understand that any regime needs teeth behind it.

Enforcement can play an important part in that, and we understand that, but I think you have got to think about the incentives that those very stringent sanctions might drive. For example, as all of the platforms here today will talk about, we innovate and trial different things. There are some very interesting social science behavioural things. We also have a bullying nudge, where if someone is going to post something harmful, we try and steer them another way.

My personal worry is that very strict sanctions might drive us towards more sweeping censorship and erring on the side of caution rather than interesting experimentation and working with expert groups to design better interventions to help behaviours online more generally. So I understand why it is there, but as we go forward in the conversation that is one of the points I am more interested in focusing on.

Q99 **Chair:** Do you have anything to add to that, Theo? And then I will come to Katy.

Theo Bertram: It is helpful if you call us out one by one, just because it is quite hard to read the cues in the room when online. I am happy to go next.

This is really impactful. The Bill will set a big precedent, in the same way as GDPR in the European context really changed the industry when it came to privacy. This will have a global impact, not just one in the UK. I would call out a few things in the Bill that are new and different from what has gone before. One of them that I really like is that this Bill recognises the different developmental stages of the child.

Previously, it is almost as if teens are treated like younger kids when it comes to the concept of child safety. This proposal would have that age-appropriate context—we have to think about risk, the age of the child, the content and the way that the product works. That is a really smart concept and one that younger people will respond better to than a much more blanket approach to younger people.

The Bill will be a precedent for lots of the world in this way also: it will take out of the hands of tech companies and put into the hands of a regulator and, to some degree, the Secretary of State the ability to define what "harm" is.

Q100 **Chair:** Sorry to interrupt—I know it is more tricky when you are online, rather than in person. The question is, how do you expect it to affect your approach to enforcing your own community standards?

Theo Bertram: One of the ways is that the Government—or the Secretary of State and Ofcom—get to determine the priorities. At the moment, all

companies would consider the priorities to be CSAM—child sexual abuse material—and terrorist content, but the Bill allows for the codes to go beyond that, so the UK could set its own priorities as well.

Another way that the Bill does something new is with risk assessment. The way I think of this, it is similar to the data protection impact assessment that is part of the GDPR. It will require companies to carry out a risk assessment as they are doing new products or features, and for that risk assessment to be available to the regulators, so that they can pull it out and look to see, when they were designing the product or the features, whether they were thinking about the harms that we asked them to think about and how that increased or decreased the level of harm on your platform.

These are significant differences from legislation that we have seen in other countries: the age-appropriate approach and the risk assessments. Both—

Q101 **Chair:** But in what way will those impact on TikTok's approach? That is what we are asking.

Theo Bertram: With age-appropriate, this Bill encourages us to think about the different stages of child development. We are doing that already—we have no direct messaging for those under 16—

Q102 **Chair:** So it will not change your current approach?

Theo Bertram: It will solidify that and require us to do more. Risk assessments we will have to do as part of this Bill—

Q103 **Chair:** That is not something you do at the moment—

Theo Bertram: Not in the way that the Bill specifies. What I would say is that child safety and online safety are a priority for our company, from the CEO down, so of course we are committed to safety. That principle is not changed by the Bill. However, there are things in the Bill that it requires us to do, which we will do.

Q104 **Chair:** Okay. I will come to you, Katy.

Katy Minshall: It is really challenging to get into the specifics of how the Bill will affect us. I will draw on something that a previous witness to this Committee said, though not in your hearings. Ellen Judson said that the challenge with the Online Safety Bill is that so much depends "on things that don't exist yet", such as secondary legislation and codes of practice.

In previous sessions, we heard from the Law Commission, which highlighted that one of the reasons for its work is the vagueness of the offences. It is critical, especially after so many years of putting together online safety legislation, to give clarity at the outset. If you will allow me, I will give an illustration of what that might look like on Twitter.

At the moment, exemptions are intended for content of democratic importance. That might be journalistic content or content from politicians. Clearly, that is well intended. One of the challenges in practice is that,

over the years, there are plenty of accounts that we have suspended from Twitter for abuse or hateful conduct that would describe themselves as journalists. In practice, the exemption might create a loophole. If someone who has been suspended from Twitter stands for election, or is registered with the Electoral Commission as a political party, does that create a loophole by which they are expected to be allowed back on to social media? Those are really big questions that cannot be ducked and should not be passed to us or the regulator to make. It is important that they are considered up front in the Bill by Parliament.

Q105 **Chair:** Thank you. I want to ask you all about some of the features. I think it was you, Theo, who said that what is really important is product design. Clearly, that has become very apparent, particularly in some the evidence that Parliament has heard more recently about the business model of online platforms. The growth of your audience is obviously very important, but in terms of priorities, sometimes it can override safety and the abuse that some users can experience.

To what extent have you already identified features or user functionalities on your platforms that increase the risk of abuse being amplified or, indeed, allowed. You have touched on some of the changes that you are currently putting in place. What more do you think can be done in this area to minimise abuse? You were asked how you can enforce your own community standards on your platforms. Clearly, it is not job done, so to what extent are you tackling, for example, gendered abuse or racist abuse, and what more can you do on your platforms within the current legislation, and with a view on incoming legislation, to improve those issues? Katy, would you like to go first?

Katy Minshall: I suppose one example of Twitter functionality where we have identified a particular area of risk is conversations. If you send out a tweet and lots of people reply to you, that is a conversation. We have taken a three-pronged approach to try to address it on that specific surface.

The prompt that I have already touched on is about nudging towards better behaviour. The second is control. When I joined the company three and a half years ago, if you sent a tweet, you had very little control over who could reply to you. Now, you are able to limit that to just people who you follow, or just people you mention, or to turn off replies altogether. The third prong of our approach is algorithms. We have experimented over the years with our algorithms to try to reduce the visibility of tweets that look like they could be abusive. Typically, if you are on Twitter and you see a tweet with many replies, a chunk will be automatically hidden at the bottom in the "view more replies" section. The culmination of that is intended to reduce the risk on that surface of Twitter and improve the health of the public conversation overall.

Something that I would touch on with regard to algorithms specifically, if you will allow me, is that we now have a team within the company called META—our machine learning, ethics, transparency and accountability team. META is a combination of researchers and practitioners embedded in

Twitter to look specifically at our algorithms—is there a risk of these algorithms causing harm; is there a risk of bias; what mitigations do we need to put in place now?—and most importantly, to share that research publicly.

The most famous example from this team that you will probably have seen is from a year or six months ago, when they identified that our image-cropping algorithm was creating racial and gender bias. Obviously, we turned it off and we shared the code on that research publicly so that others could learn from it. More recently, they have looked at the interaction of our algorithms with political content on the platform.

To come back to your question, I think part of the solution to those challenges is to look at our terms and conditions and our content moderation, but that cannot be the whole story. Again, to build on what Ellen said last week, you have to be looking at design, algorithms and nudges—how you are encouraging best behaviour—and measuring your progress that way as well.

Q106 **Chair:** To come back to you, Katy, would Twitter welcome Ofcom having a role in assessing the use of algorithms, given that Twitter seems to be comfortable with publishing the findings of its unit that is looking into this? Would Twitter welcome a code of practice that sets out clear expectations of platforms in respect of these issues?

Katy Minshall: We have long been supportive of Ofcom being designated as the regulator. It is a credible, experienced regulator that is well up to the task. On algorithms, there is partly a need for more transparency and clarity around the intention of algorithms, which I have covered. One of the public policy outcomes has to be control of the algorithms that affect your online experience every day. As an example, on the home timeline on the Twitter app, with just a couple of clicks you can turn the algorithm off. So if you just want to see tweets from accounts you follow, in reverse chronological order, that is your choice. At the moment, that is quite binary.

In future, ideally we will get to a place where there will be a lot more variety in the algorithms you could choose to control your experience. In terms of the code of practice and expectations of how to address something like online abuse, that is part of the solution. You might be someone who does not want to see profanity on Twitter or any social media site, to give a basic example. If there was an algorithm that you could choose where you did not have to see that kind of content, that option would be available for you to control your online experience.

Q107 **Chair:** Thank you. The same question to Rebecca and Theo. What are you doing to control this, particularly in terms of algorithms and the amplification of online abuse? Do you welcome the role of Ofcom? Who would like to go first?

Rebecca Stimson: I am happy to go first. My answer is similar to Katy's. We think one of the key things about algorithms is ensuring transparency. We publish our recommendation guidelines so that people can see what

we try to make people's experience online be, what we will not recommend and what we filter out through our AI and enforcement.

We also give people control: there is a feature on Facebook called, "Why am I seeing this?" You can go in and look at the information that we are using to inform what you see in your newsfeed, and you can adjust it. Similar to Twitter, you can have an entirely chronological, un-algorithmed experience if you want. We tend to find that people do not actually enjoy that—it is not the best experience because there is so much content online and having it curated to some extent by algorithm is better.

Looking forward, we have a dedicated responsible AI team to ensure that the algorithm and artificial intelligence are used for best effect to society and people. We are doing a very large, multimillion-pound project on AI and ethics with various universities, including the University of Munich. We are looking forward as well as building tools and transparency right now. Ofcom's role in that can only be helpful. Greater transparency and more common standards around how algorithms are used can only be helpful as the online safety regulator gets into its role.

Q108 **Chair:** Is that feature to remove all algorithms available now?

Rebecca Stimson: Yes.

Q109 **Chair:** Is it easy for a user to access?

Rebecca Stimson: I am sure I am not the only platform here where people always give us feedback that these tools could be easier to access, but yes. It is easy to access and available.

Q110 **Chair:** Is it something that Facebook and Instagram might promote?

Rebecca Stimson: It is all publicly available. We have various help centres and transparency centres. It is all there and it was announced publicly when we rolled that feature out. Hopefully, it is easier to find. I tend to think that most people do not tend to use it, because it ends up not being a particularly great experience for the user, but it is available to remove that completely if you do not want it.

Theo Bertram: I agree. I think Ofcom will be a good regulator: in fact, we are already regulated by Ofcom under AVMS. In terms of gender and race and the abuse we have seen in that area, on race, we were one of the sponsors of the Euros, and ahead of the Euros we ran a campaign on the platform called #SwipeOutHate. That was viewed more than 80 million times. I think we underestimate the power of positive campaigns to change social attitudes, at least on the platform, and I think it did have a positive effect: we saw less abuse on our platform than I think others had been victim to.

We are currently just about to launch a campaign around 16 Days. As you will know, Thursday is the International Day for the Elimination of Violence against Women, and starting then and going through to Human Rights Day on 10 December, we have a campaign we have been working on with Women's Aid and Refuge on violence against women, to promote

awareness of the issue and issues like consent on the platform. We think that is quite an effective way of trying to help guide people towards doing the right thing, but of course, we also have the same enforcement and policies as the others.

Then, on algorithms, we are slightly different in the way that our app works than the other apps. The difference is the notion of a content graph, rather than a social graph. Just to explain what that means, if you are on Twitter, the content that you see in your timeline is shared into your feed by what other people have posted: who you follow determines what is in your feed, so the contact book is what you see. On our platform, the videos you see will not be from people you know or people you follow; it will be based on your interaction with videos that you see, so it works in a different way. That means that the challenges we have are different in some ways, but we are always trying to make sure that we diversify the content we have in our algorithm. I am happy to talk more about that; I do not want to go on for too long.

Chair: Okay. Martyn, have you got some questions you want to ask?

Q111 **Martyn Day:** Yes. This one is for Theo: the Bill puts a stronger duty on protecting children than it does for adults, and obviously your platform is quite popular with the younger audience. How does that translate into stronger protections for under-18s on your platform?

Chair: From bullying and abusive content.

Theo Bertram: The Bill requires us to give more protections the younger people are and the more risk you have, so if you are a porn platform you have to give a lot stronger protections, but for us, our content is all aimed at everyone who is over 13. Obviously, we already have a duty to keep under-13s off the platform. We are currently the only platform to publish the number of under-13s that we remove, and we removed 7 million in the first quarter of this year and 11 million in the second quarter of this year.

Then, for those users that are on the platform, we are already designing the platform to be safer for younger users. If you are under 16 when you come on the platform, you will only be given the option to create videos that you can share with people you know, but not with everyone, so that is a kind of first protection. We also do not have direct messaging for those under 16, and a number of other features are not available at different age groups until you get older, so only when you are 18 do you have the full suite. I think there is some protection in the design there, and I think what this Bill enshrines is that all companies now have to design with that principle in mind.

Q112 **Martyn Day:** This question is to all participants, but could you start off again, Theo? How confident are you that your methods for verifying young people's ages are robust enough to identify which ones need these additional protections?

Theo Bertram: The numbers I gave you—7 million in the first quarter, 11 million in the second quarter—were removing more people than are actually under 13, because they then have to appeal, and a number of those are coming back on the platform because it turns out that they were over 13. This is an issue that is really important, and I know it is one that has been debated a lot. I find that I am asking myself questions, but you can tell me if this is the right one: should we have age verification for those under 13? Give me a nod if that is the direction you want me to go, or shake your head if I am asking my own questions.

Q113 **Martyn Day:** How do you know if someone is under 13? I suppose that is the gist of the question.

Theo Bertram: Whenever anyone reports a video on our platform, there will be a human moderator who looks to see whether that individual is over 13 or under 13. We have also made it very easy for users to report other users as being over 13 or under 13. Those are ways that we can remove under-13s, and obviously we remove everyone under 13. If you have a formal form of age verification, there are a couple of problems with that. If I, as a parent, think about every time my son goes on to a website or app that requires me to provide documentary evidence of his age being over 13, that will first be a big privacy risk in terms of the amount of data that I have to hand out. Secondly, it will mean that parents like me do that less often, and that will ultimately benefit the Metas and Googles of this universe rather than the smaller players that are going to compete. From both privacy and competition angles, forcing people to have ID is a challenge.

However, if we need to have age verification, there are some bottlenecks in the system already, and that would be at the app store. I have only one app store on my phone. If that verified my age, all the apps on the phone could just ping the app store and check whether the person is over the right age. That would be a better solution than saying every app and every website must verify age, but at the moment we have an industry standard: when you come on the app, you have to declare your age. We don't prompt you for what that is. When you are on the app, we remain vigilant. As I say, we removed 7 million users in the first quarter and 11 million in the second quarter of this year.

Q114 **Chair:** May I come in with an additional question? Theo, you have described how TikTok has made it private by default. Obviously, it makes efforts to ensure that users of TikTok are over 13 years of age. It does not allow direct messaging and has some controls over what videos can be posted online. However, what you have not answered is what TikTok is doing to control the amount of abuse and hatred that young people who have TikTok accounts see, even when they are over 13—there is no age verification issue there. Looking at the recent example of the teacher videos that went viral, they have obviously caused immense distress to teachers, but it is also highly abusive content for young people to see, with young people being the main audience of TikTok videos. What is TikTok doing to prevent something like that from happening and to stop millions of young people seeing their teachers being abused and all the

impacts of that?

Theo Bertram: Yes. I think the way we have designed the product—some of those things do reduce abuse, but obviously we have a huge enforcement team moderating the platform and removing abuse, and we have zero tolerance of that.

On the pattern of abuse towards teachers that we have seen recently, I can talk a bit about that in detail, if that is helpful, but tell me if I am going on too long. We saw this start about a month ago—a relatively small pattern on the platform of teenagers leaving comments, judging their teachers and things like that. It is something that—

Q115 **Chair:** Sorry, Theo, but we don't have a huge amount of time. My question was: what is TikTok doing to prevent it?

Theo Bertram: We have been working with the National Association of Head Teachers, the Association of School and College Leaders, DFE and DCMS. We have been listening to their complaints. We saw this pattern on the platform. After media coverage, we saw a spike in that. We developed an emergency response to that. We listened to the trade unions, the teachers and the Department to identify the trending hashtags and the types of content that we saw on the platform.

We have removed the vast majority of that content and suspended accounts. We have also given funding to POSH—the professionals online safety helpline—which supports teachers and enables them to report to us more quickly and more easily; and we are in the process of writing to every headteacher in the country to give them guidance on how we can tackle this. So we have put in place a mechanism to tackle that content; it is coming down off the platform and we have seen it reduce greatly in the past couple of weeks.

Q116 **Chair:** Okay. Obviously it is very concerning for the teachers affected, but a huge number of young people are also impacted by this situation—people who have not posted videos, but are seeing a climate of abuse on a very familiar platform that has been a daily part of most of their life. The particular subject of teachers is a one-off, but the abuse—the daily abuse—is not. It is a symptom, so my question is what is being done to treat the cause of an abusive atmosphere online? Martyn has more questions for you, but if you want to wrap your response to that in your answer to his next question, that would be good.

Theo Bertram: Through improved product design and our community guidelines and reporting process, we are tackling that abusive content. As I said, 94% is taken down before a single report, 80% before a single view. The overwhelmingly experience of the vast majority of users of our platform is positive and joyful. There are undoubtedly abusive comments on our platform, as there are on all platforms, and that is always regrettable, but we are constantly trying to improve our service to make sure that we can take more down. I do think the overwhelming experience is a positive one, particularly for young people.

Q117 **Martyn Day:** To put the question to the other panel members, what are your platforms doing to identify the younger people who need those extra protection?

Katy Minshall: Like the other platforms here, we have a minimum age of 13. Part of our approach is that we are not a service that is targeting a youth audience. Third-party independent research tends to confirm that the majority of Twitter users are over 21. Ofcom does its own research looking at the popular apps among young people, but this is an industry-wide challenge and I do not think anyone has solved the problem. The Children's Commissioner is looking at this right now; we have regular meetings with the ICO, which has looked at this through a different lens. In the meantime, our priority is to ensure that Twitter is safe for everyone, including any young people who may come on to the service, through things like having, on sign-up, safe search on by default and sensitive media settings on by default.

Rebecca Stimson: Like others, we rely at the moment on stated age. When people sign up for our services, they have to give their age; if they give the wrong age, we do not allow them to continue. We do not tell them that is the reason why they cannot continue so that they do not try to game it, and we do not allow them to make multiple attempts at entering their date of birth, trying to get it right. Once younger people are on the platform, we try to find them and remove them.

Like others here, we publish statistics—for example, in the last quarter, 600,000 underage accounts were removed from Instagram—but there is still a long way to go. As others have said, relying on stated age is clearly not a foolproof way doing this, and I think that as the Bill goes forward, there will be a very interesting conversation about what exactly is meant by age assurance and what companies will have to try to deliver under that Bill. One of the things that we think has promise and gets around some of the challenges that Theo mentioned about requiring some form of ID from everyone on the internet, which carries a lot of risks and concerns about privacy and which could be regarded as disproportionate, involves using AI to spot younger users. We have been working on this for quite some time. You could use various signals and different forms of technology to identify with a high degree of accuracy users who are too young and to remove them in an automated way. That would help us to do this at scale, similar to the way we automate the removal of harmful content. That technology is not there yet—it is not sufficiently accurate—but we think it is probably one of the more positive avenues to talk about in the next few years as the online safety regime is stood up.

Q118 **Martyn Day:** Thank you. You have mentioned the steps you are taking to proactively remove content, but are you doing enough to strike the balance with allowing users to block what they see themselves? How do you feel that is going?

Rebecca Stimson: On a given day on Twitter, we will see about half a billion tweets, so some level of automation is essential to content moderation. We have in place a range of algorithms trying to detect any

content that breaks our rules. To give you a sense of the scale, last year alone we sent about 300 million anti-spam challenges, we suspended about 2 million accounts and we removed 6 million pieces of content. Content moderation is an important tool, but as I said earlier, it not the whole solution. We have really doubled down over the past few years in thinking more about where the risk services in Twitter are and what fundamentals we need to change.

We are experimenting. You mentioned block and how to get the balance right. At the moment, too much of the burden is on the individual to go through and block every single account that is trying to interact with them in a harmful way, especially if you are someone in public life. Blocking itself can be quite confrontational and some people may not want to take the step of saying, "I am going to block this user." We are therefore experimenting with a tool called safety mode, which essentially says that if you want to go into safety mode, Twitter will step in and automatically block, on your behalf for seven days, accounts that look they are trying to interact with you in a harmful, offensive and spammy way.

The tool is in trial at the moment because, as you can imagine, there is a risk if we were to over-enforce—it might prohibit your constituents from interacting with you as Members of Parliament—but if we were to say it is safety mode and under-enforced, that is a problem too. We need to trial it to get it right, but I think that is a good illustration of the approach we are trying to take, which is that you should always feel safe expressing your point of view on Twitter, and controlling your Twitter experience and the experience you want to have is a key part of that.

Rebecca Stimson: I don't think any of us would claim that we are doing enough, even if we were to show some really good progress and really positive statistics, particularly on proactive detection and removal. Adding to what Katy said, I think a lot of this is about prevention. Detecting, finding and removing harmful content online is extremely important and something we are all focusing on doing to the best of our abilities, but we are also doing a lot about preventing harm from happening in the first place. That can be about ensuring that we build in considerations about safety from the very beginning. We have a dedicated team in Meta where our goal is to think about that from the very start and to surface the kind of problems that you may see in a different product or feature from the very beginning. That is greatly preferable to reacting to harmful content that is online.

Other prevention measures we take are around the tools we give people to enable them to control their own experience, to block people and prevent people from contacting them through their private messages—things like that.

The last thing I want to mention in the prevention space is some of the behavioural nudges that we have commented on a couple of times already today. Trying to dissuade people from posting harmful content, from sharing something that is false—that kind of behaviour. You try to prevent harmful content from spreading or reaching the intended recipient from

the beginning. That is additional to the work we are all doing collectively to remove it when it does make it on to the platforms.

Q119 **Chair:** May I ask you a supplementary question, Rebecca? I know that, during the US elections, Facebook was able to temporarily put in measures—safety systems—to reduce misinformation. My understanding is that they were switched off after the election. Is that the kind of broader safety approach that Facebook is looking at? If it was successful or deemed to be successful during the US election, why isn't it still in place? We are all dealing with a lot of covid and vaccine misinformation, for example. Wouldn't Facebook have a role on that with it being a current challenge that's on a similar scale to the US election?

Rebecca Stimson: Yes, I wish we were able to switch off misinformation during the election—our lives would be much easier. We switched off a number of different things in the last few days of that election, including around political advertising. What we were trying to do there was reflect the extraordinary circumstances that we saw in the US around the election and then around the inauguration day. So there were some things that we switched off. I think the Law Commission were quite interesting when they were talking in their evidence about reflecting the context in which posts and online conversations are happening. That was why we did it; it was reflecting that circumstance on the ground.

On your point about covid, I absolutely agree. We have been working very hard since the beginning of the pandemic at the start of last year to support public health efforts and Government efforts. To respond directly to your question, one of the things that we do is that we no longer recommend health groups. Health groups can be really important and valuable ways for people with certain conditions to find a community, connect online and share information. But one of the things we learnt during the pandemic was that actually, in that particular context, it is more important that people get their information from authoritative sources and public health bodies, so we built a covid information centre. We had it at the top of everyone's news feed. We directed tens of millions of people from the UK towards that information. We didn't remove those health groups, because, as I said, by and large they do play a very important role for some people to discuss their health concerns, but we no longer recommend them to anyone, because at the moment, we think it's more important to direct people towards those formal sources of information.

Q120 **Martyn Day:** I have a final point I want to ask you about. Looking at your automatic methods of detecting abuse and deleting it, we have heard some concerns from users that that may disproportionately affect certain languages, such as Urdu or Arabic. Is that a fair criticism, do you think, and how do you address that problem, if it is accurate? Katy, do you want to go first?

Katy Minshall: That specific challenge is not one I am familiar with, but I am very happy to come back to you specifically on that.

In general terms, we have a number of risk assessment and mitigation processes for all aspects of Twitter's product, algorithms, terms of service etc. One of the most important aspects of our system is our Trust and Safety Council, which is a global network of dozens of organisations from around the world that represent communities who have been disproportionately affected by online harms and represent organisations with a specific expertise in online safety. Typically, for any challenge, including a concern like the one you have just raised, it means it's not just Twitter trying to investigate this or trying themselves to identify the potential range of online harms around the world; we are able to leverage expertise externally to bring these issues to our attention and advise on our response to them.

Martyn Day: Thank you. Rebecca, would you like to comment on that as well?

Rebecca Stimson: We have moderators around the world: in 20 different locations, there are about 40,000 people working on safety and security, and around half of those are moderators. And we operate in about, I think, 70 different languages. So I totally recognise the point you are trying to make. The way we try to tackle it is to have teams like the one I'm in—I'm the UK policy head—to give extra context and extra local knowledge, for example of expressions or terms or certain languages that other people may be less familiar with in these global companies. I couldn't say that we get absolutely all of that correct. I'm sure there's a lot more to do. But I have given you the stats, which will hopefully give you the sense that we are quite a long way down the track to try to ensure that languages and local context are not a hindrance to us enforcing our rules.

Martyn Day: Thank you. Theo, would you like to comment?

Theo Bertram: On your earlier question, I think generally there is progress in the tech industry. It is similar to the way the car industry has transformed itself over the last 30 years in terms of safety. You now have all this innovation in cars that warns you about lane changing or how far you're behind the car ahead. And each time you buy a new car, that's the expectation of where safety will be. I think innovation from within the industry is part of it, and I think you can hear from all the companies today that investment and work are going into innovation.

To give another example from our platform, there was a case recently about a boy with epilepsy who was being bullied, and the idea was that flashing images were being used as a way to provoke him. What we designed on our platform was a filter so that if you are creating a video that has this kind of flashing imagery, we will give you a prompt to say, "Don't post this video or it will trigger this setting", and then if you are a viewer who sees this video, when you see it, it will say, "This video has flashing imagery that might be photosensitive, so do you want to skip this video or do you want to watch it?" Users will tend to skip that, so I think those kinds of evolutionary changes we are seeing in the industry will make a difference.

However, just as with the car industry, regulation has a role to play, and so does how we change social behaviour. Listening to the witnesses before our session, I was thinking about this focus on how we criminalise those who are actually making the abuse in the first place. Hopefully, that will have the same sort of effect that criminalising drink-driving did, and demonstrate that there is a real change in the way that society should look at this kind of behaviour.

On the issue of over-moderation of certain types of language or content, when it comes to harassment and bullying or hateful behaviour, these are more nuanced issues than things like, "Is this a nude image?" or "Is this spam?", so detection rates for that tend to be lower because they require more nuance to look at. However, across all the different areas of risk that we are looking at, one thing that we do consistently is that we are always looking at, "How effective is this algorithm?" We look at that in terms of whether it is removing the harmful content, but also whether it is going too far, and I think that is common practice in the industry. You do not have an algorithm detecting or removing content on its own: you always have that human oversight checking to see whether you are getting that balance right, but that is a difficult task.

Q121 **Chair:** Okay. You have touched on one of the big challenges with algorithms, Theo: "Does this algorithm remove harmful content?" One of the biggest challenges that is faced—the most prominent examples have generally been on Facebook, but there are also some examples from TikTok—is that of algorithms actually promoting harmful content. That is a challenge that has not really been properly addressed today.

Another one of the big challenges we have heard about as part of this inquiry and the evidence we have taken is in relation to anonymity online. We have heard evidence that speaks for anonymity online and what it brings to the online community, and also about people who are vulnerable being able to have a voice online, but we have also heard evidence about people who are vulnerable not having a voice online because of the level of abuse that many would attribute to anonymity.

Obviously, the purpose of this inquiry is to come up with some conclusions on this question. What would be each of your respective platform's position on anonymity and whether it is something you want to keep protected? Is it something you would link to some of the abuse that is circulating, and what would you suggest should be done in terms of a platform's approach to anonymity? Do you think that is part of the armoury in trying to tackle online abuse? I can see that you are all nodding, so I am sure you all have an answer, but I will come to you first, Katy.

Katy Minshall: Thank you. I think you characterised it well: there are pros and cons. Anonymity is also described as a risk factor when it comes to online safety, but on Twitter, it is absolutely also an online safety tool. We are a global platform, and there are activists, whistleblowers and political dissidents around the world who could not be on Twitter without the use of a pseudonym. Closer to home, if you are a young person

exploring their sexual orientation or you are a victim of domestic violence, being on social media is potentially a lifeline for you and one that would not be possible without the use of a pseudonym, so I think it is important, first off, to acknowledge it as an online safety tool in the broader context of considering the risks.

When you come to Twitter to sign up for an account, we ask for your full name, your date of birth, and your email address or phone number, one of which you have to verify in order to get on to the service. When you are on the service, there is all sorts of other information that you might share with us in the course of the tenure of your account. What that means in practice is that, when on Twitter, whether you are @katyminshall or @pseudonym123 is no shield from our rules and no shield from criminal liability. We want to work with the police on these issues, particularly considering the evidence that the Law Commission gave in the session before this.

In terms of how you square the circle, in the hearing that you had last week one witness talked about whether a way of achieving this was giving people more control over their accounts on social media. I don't want to mischaracterise him, but I think he said something along the lines of: "If you don't want to engage with anonymous accounts, you can control who should be able to reply to you, interact with you, etc." Aspects of that are, to some extent, the way forward. On Twitter, you can control where you get notifications, so if you do not want to receive notifications from accounts that are new, that have not changed their default profile photo or that have not verified an email address or phone number, you can do that.

As I said earlier, as of last year you can control who can reply to your tweets, so if you only want people who you follow—who you yourself have verified—to be able to reply, that is a choice too. Safety mode, which I talked about earlier in this session, is supposed to take that a step further, in that it is trying to automatically block accounts, if you so choose, whether they are anonymous or not. The way forward will be as much to think about the law, and the work of the Law Commission, and terms and conditions as well, as to think about product features and product safety, and making sure that you are having the experience that you want to have on social media.

Q122 **Chair:** Do you want to answer that as well, Rebecca? I know that Facebook's terms and conditions require users to use their real name. What assessment has Facebook made of how much that improves safety and interaction on the site?

Rebecca Stimson: I would agree with your opening question. I think anonymity does play an extremely important role for lots of people online. I guess when we think about what, if anything, to do about that there are two key questions. One is: where is the evidence that anonymity is driving harm online? I am going to quote one of Katy's statistics that I know: after the Euro final, they found that there was no evidence at all that the abuse that football players faced was from anonymous accounts; in fact, quite the opposite. We certainly do not have any evidence to support the

assertion that anonymous accounts are leading disproportionately to harm online.

The other question is what the use is for the ID that we might ask for from people for accounts. As Katy said, an anonymous account does not prevent us from taking the full suite of enforcement action. It is still subject to all the same rules. One of the other misconceptions is that it somehow prevents us from working with law enforcement. In fact, it doesn't. There is an enormous amount of information that we will hold on an individual, and the most useful information that law enforcement often ask from us is things like traffic data and basic subscriber information—the kinds of things that we will have regardless of what the account may be called on the face of it. We can work with law enforcement, we can still enforce our rules, and as far as I am aware—though that is not to say that some could not emerge at some point—there is no evidence linking anonymity and harm online.

The last point is almost back to the identification question. As you will know, the Law Commission here in the UK said that 3.5 million UK residents do not have any available form of Government ID, so requesting that kind of documentation before opening an account on Facebook or Instagram would really exclude a very large percentage of perfectly innocent, law-abiding people from using these services. I understand why people are concerned about it, and as we move forward all I would do is encourage it to be a very evidence-based conversation as we think about whether or not this is a problem that needs solving.

Q123 **Chair:** Theo, you are invited to respond to the initial question, but the other thing I would like to know from all your platforms—this is the subject of one of the petitions—is, how do you stop trolls, which many would deem to be harmful even if it is not illegal? If it does get to the point where they are banned from using any of your platforms, what is to stop them just popping up under a different anonymous identity? This is one of the issues that is raised with us repeatedly and that I do not think people are confident that the platforms are dealing with. Do you want to come in first, Theo?

Theo Bertram: Sure. I echo what Katy and Rebecca put well: there is no evidence that anonymity is the driver of abuse. The thing we need to tackle is the abuse; I think anonymity is a separate issue. There is then the question of accountability: if you post abuse, can you be held accountable? All the platforms have explained how we have ways to help the police to identify who uses our platforms, even if they are seemingly anonymous on them.

Q124 **Chair:** Sorry, but that is for illegal abuse or illegal harm. What many people are concerned about is the platforms' ability to deal with harmful content that isn't illegal. So it is not about law enforcement; it is about your community standards being enforced.

Theo Bertram: Yes, you're right—I still had in mind the discussion that was taking place just before this one. But I also don't think that the way

that that type of harm occurs is because of anonymity. What we see is a situation, whether this is in a playground, in a football stadium or online, where the people you are associating with are pushing the boundaries of what is acceptable, so that you start to think that that is acceptable. It is where that behaviour occurs in groups that we tend to see the most abusive content being generated. People think, "It's okay to do this." That is a different thing from anonymity and there is evidence that that is something that we need to tackle.

Your second question was about trolls and the idea of recidivism, where we have kicked them off the platform. In the second quarter of this year, we removed nearly 15 million accounts from our platform. How do we stop those people getting back on the platform? I will not go into the detail of how we do that, but even if they are using a different account or a different name, we will be able to detect that those users are coming back on to the platform.

Q125 **Chair:** How?

Theo Bertram: I don't want to give you the detail of that publicly but I am happy to put it in writing and send it to the Committee.

Q126 **Chair:** Okay, thank you. Let me put that to the other platforms: how do you prevent people from phoenix-ing on your platforms under a different anonymous identity?

Katy Minshall: First, there are our rules. If you are permanently suspended from Twitter, you of course have the opportunity to appeal, but assuming that the appeal is unsuccessful, that suspension is permanent. If you try and get around that suspension by setting up a new account, that is against our rules.

In terms of enforcement, as with all our rules, it is a mixture of proactive and reactive. On the proactive side, like Theo, I would not want to go into the specifics in a public setting, because it might give someone a better sense of how to get around our defensive measures, but in general terms it is done through a range of digital forensics, and enforcement, in many cases, is near instantaneous upon detection. But we do not catch everything and someone who is really determined to get back on to a social media site is going to try and find a way. We do receive reports from partners who will tell us, "We think this individual is coming back. Please look into it," and if they are what we call a ban evader, we will suspend the account again.

I think that illustrates that there is a role for the Law Commission and the law to some extent. In general terms, the strongest sanction that we have is permanently suspending someone from Twitter. That is not a decision that we take lightly, but if the question we are thinking about is driving cultural, societal and behavioural change, I think there is a consideration around real-world consequences—legal sanctions—when it is getting to the point of illegality.

Q127 **Chair:** Do you have anything you want to add, Rebecca?

Rebecca Stimson: Very similar answers. We also have abilities to block people. Obviously, none of us want to talk about it in great detail, because it helps people game our systems but, for example, one of the ultimate sanctions we can take is blocking people at their device, so they would need to go and buy another phone in order to continue to open the account. This is a last-resort measure, but things like that are possible.

One other thing I would add, in addition to what others have said, is that we have a feature where the user can block somebody. That person may not have reached the threshold for us to permanently remove them, but we allow users to block whoever they like in certain ways. There is a feature where you can pre-emptively block new accounts that that individual may go ahead and open in the future, to try and mitigate that reoffending point that you were making.

Q128 **Chair:** I am going to bring in Christina because she has one final question, but we have to go soon as we have multiple votes coming up. I want to put to each of you, if anonymity is not the issue and if requiring verifiable ID isn't the solution, then what is the solution if we still have social media platforms constantly plagued by abuse? That is the reality, so what is the solution?

Katy Minshall: I can start. First, we can't work in silos. This has to be a combination of steps social media services are taking. Then, there is education—making sure everyone knows how to behave acceptably and appropriately online—and criminal deterrence where necessary. The combination of those three strategies is important.

I am sure all of us would say that a key theme of this session is product design and thinking about the fundamentals of your platform—where the risks are and mitigating those risks to try and reduce abuse or other online harms in general.

Something that has not been touched upon so far in this discussion is the importance of transparency and making data available externally. The way that we will be able to measure progress and bring expertise beyond Twitter to bear on these challenges is going to be by making data available. We have made our APIs—the technical means by which you plug into the livestream of Twitter data—available for a number of years.

Probably tens of thousands of researchers have accessed Twitter data to ask these key questions: how can we respond to online abuse, and what are the issues? The challenges we see on Twitter are incredibly well documented, but frankly I think we are all the better for it, because it is far more straightforward for us to work externally with people outside the company, with whom we can discuss the challenges we are seeing and what might be the most effective way of responding to things like online abuse.

Chair: I don't know if you have anything to add to that, Rebecca or Theo, but I am conscious of time so I will go to Christina. If you have something to add about that question, please incorporate it into your response.

Q129 **Christina Rees:** Thank you, Chair. We have heard from previous witnesses that offences relating to online abuse, including online hate crime, remain under-prosecuted. Are you confident that your companies are doing everything they can to help to rectify that, in particular when it comes to providing relevant data to law enforcement bodies? I will go to Katy first, because you have alluded to it already, then Rebecca and then Theo.

Katy Minshall: We want to work with the police on these issues. That is the reason why we have made a specific portal for the police to make requests to us for account information, to remove a tweet that is illegal or to preserve data about an account where they may know they will need to investigate but they aren't ready to request the data yet. That is also the reason why we run regular training for the police so that they have all the information they will need to make requests to Twitter.

There is always more we can do. Thinking about my answer earlier, there is definitely a role for thinking about offences and how to make sure the systems are set up, so social media services and police are interacting in a seamless process. The key thing is that we have those relationships set up, not just in the UK but around the world, with law enforcement, who are our core partners in this issue.

Rebecca Stimson: We have a dedicated law enforcement team based in London, which works closely with UK law enforcement. We have a number of reporting channels where they are able to request information from us to support their work. We also have a number of crisis response protocols that we stand up, for example, when something like a terrorist incident may be happening, or a specific live incident. Meta is part of the GIFCT forum of companies designed to work together to tackle global terrorism. We are able to respond very quickly and work closely with law enforcement to give them information about what we are seeing and support their investigations to the best of our ability.

Theo Bertram: Similar to the other two companies, we have a law enforcement response team that responds quickly to UK law enforcement. We also have an emergency protocol if there is life at risk, which has an extremely quick response time. There is a rapid response in place for law enforcement.

Q130 **Christina Rees:** If the Government adopt the Law Commission's proposed reforms to hate crime and online communications offences, are you clear on how that would affect content on your platforms that would be considered illegal? We will go in reverse order this time, so let's start with Theo, and then Rebecca and Katy.

Theo Bertram: The Law Commission is focused on criminalising the individuals who create the abusive content. You heard in your previous session about the nuance and complexity of doing that. I do not think it would be possible for a platform to prosecute some of those cases. The vast majority of abusive behaviour is obvious and we all remove it, but there will always be content that requires nuance and context. One of the

previous witnesses talked about how sending an image of a front door on its own might constitute a crime under this notion of what is highly contextual, if it was sent with intention and without a reasonable excuse. It is difficult for a platform to know that without having law enforcement provide that guidance and advice. But if law enforcement come to us and say, "Look, we have the context of this and here is the evidence," it is very easy for us to act. But the earlier witnesses were expressing how that will require context and nuance to define.

Chair: Just to warn you, the vote is imminent, so if you can be as quick as you can, that would be really welcome.

Rebecca Stimson: Theo is right. What you do not want is for those powers to end up putting private companies in the position of working out what is crime and what is not. That could potentially lead to some very complicated outcomes when we look at content on our platforms. There is a long way to go. Overall, the Law Commission's recommendations are really interesting and quite helpful for clarifying existing law. There is a way to go in knowing what impact they would have. The one I would flag, which is similar to this, is if you think about knowingly sharing false information, it is very challenging at this stage to know how a platform would work out the motive of a user. Were they knowingly posting that or did they do it because they genuinely thought it was true and made a mistake? That is the kind of issue to unpack as we work out how those recommendations will be incorporated into the Online Safety Bill.

Katy Minshall: I don't have anything to add, in the interests of brevity.

Chair: Thank you very much for your evidence and for answering all our questions. We are out of time. We will conclude our inquiry next week with evidence from the Government.