HOUSE OF LORDS

Communications and Digital Committee

Corrected oral evidence: Freedom of expression online

Thursday 13 May 2021

3 pm

[Watch the meeting](#)

Members present: Lord Stevenson of Balmacara (The Chair); Baroness Bull; Baroness Buscombe; Viscount Colville of Culross; Baroness Featherstone; Lord Gilbert of Panteg; Baroness Grender; Lord Griffiths of Burry Port; Lord Lipsey; Baroness Rebuck; Lord Vaizey of Didcot; The Lord Bishop of Worcester.

Evidence Session No. 30   Virtual proceeding   Questions 243 - 251

Witnesses

I: Henry Turnbull, Head of Public Policy, UK and Nordics, Snap Inc; Elizabeth Kanter, UK Director of Policy and Government Relations, TikTok.

USE OF THE TRANSCRIPT

This is a corrected transcript of evidence taken in public and webcast on www.parliamentlive.tv.

Examination of witnesses

Henry Turnbull and Elizabeth Kanter.

Q243   **The Chair:** Good afternoon. This is a meeting of the Communications and Digital Committee. We are taking evidence today from two colleagues, and we are very grateful to them for joining us. Henry Turnbull is head of public policy for the UK and Nordics at Snap Inc, and Liz Kanter has been director of policy and government relations at TikTok since May 2019. I thank both witnesses for agreeing to speak to the committee. As you probably know, a transcript will be taken and the session is being broadcast live.

Given that we have had to make some slight changes to the order of the questions that we have been thinking about asking—all questions will be asked but not quite in the order that they were first promulgated—it might be helpful if you could both give us a couple of minutes' background about your company and what you do there, as well as any thoughts you might have that would be helpful to the session this afternoon.

Henry Turnbull: Thank you, Lord Stevenson. Thank you for the opportunity to be here today. I will give a little bit of background about Snapchat as a platform and how we look to ensure users' privacy and safety, while also encouraging and promoting self-expression.

At Snap, we put a lot of thought into designing a platform that enables people to express themselves freely with their close friends in a safe and positive environment. We have taken some very conscious design decisions with the aim of enabling that.

At a high level, we use two principles that guide our design process: privacy by design, which focuses on data minimisation and protecting user data; and safety by design, which is about ensuring the safety of our community. Privacy and safety experts are involved in the development of any new products and features at Snap, the idea being to prevent privacy-intrusive or safety-deficient products or features from coming into existence.

The focus on privacy and safety by design is reflected in the build of Snapchat, which is designed to function very differently from traditional social media. Snapchat is, first and foremost, a visual messaging application intended for private communications, either one to one or small groups. The whole app is designed for, and geared towards, communicating with your close friends, rather than connecting to, or broadcasting out to, people you may not know very well in real life.

We have consciously taken design decisions not to include some of the most common features of traditional social media. When you open Snapchat, you do not open a newsfeed that might encourage a user to scroll through or consume content; you open into the camera. That is a design choice that nudges the user to engage with the camera and the world around them. In general, the app is much more geared towards creativity than to consumption.

In other design decisions that we have taken, we do not offer the ability to publicly comment on another user's snaps. We have seen on other services how that can create potential for harmful interactions. You can never see another user's list of friends on Snapchat or see how many friends they have. That is something that adds to social pressure. Default settings are that you cannot message or receive a message from somebody you have not accepted as a friend. All these features are basically designed to ensure that Snapchat is a safe and positive environment for our community.

The committee might also be aware of the first innovation on Snapchat, which is ephemeral or disappearing snaps or messages. The idea behind that was to encourage creative self-expression rather than needing to worry about everything that you say online being recorded indefinitely. We have found that that is particularly valuable for younger people, and it is reflected in a wider focus on privacy and data minimisation across our services.

However, ephemeral messaging does not mean that there is nothing that we can do about illegal or harmful content. That is a bit of a common misconception that we face. Anybody who agrees to use Snapchat agrees to abide by our community guidelines. These are publicly available online. They are very simple and clear about what activity is prohibited on Snapchat, whether it is bullying, harassment, sexually explicit content or threats of violence. We make it very easy for users to report any violations across the platform. We review reports and enforce against them quickly.

The only other distinction I would make, given the basis of the inquiry—we may get to this later—is that, in our view, there is a difference between expressing yourself to your friends and broadcasting content publicly to large groups of people. Snapchat is not, and never has been, designed as an open town-square-style environment in the form of some platforms, where people can publicly broadcast what they like. When it comes to more public content, we have additional guard rails in place.

All public-facing content in Snapchat, whether in our Discover tab for news and entertainment or our Spotlight tab for the community's best snaps, is pre-moderated, prior to being surfaced publicly, against the guidelines that I mentioned. This helps us to ensure that things like hate speech, terrorist content and misinformation are not broadcast in public spaces on Snapchat. Because of those guard rails, I think it has been well recognised that we have been successful in keeping activity that violates our terms from public areas of the platform.

Thanks again for the chance to be here today. I look forward to getting into the discussion.

The Chair: Thank you very much for that. We will certainly pick up a number of those points. Liz, would you like to say a few words?

Elizabeth Kanter: Definitely. Thank you so much for having me this afternoon. Some of you may be familiar with TikTok. We are a newer platform. We are about two and a half years old. Users typically engage

with content on TikTok via what we call the For You feed, which is a unique feed for each individual user; your feed of videos on your For You feed will not be the same as mine.

We inform the content that is on the For You feed by looking at two different categories. They are the categories that you have indicated you are interested in when you joined the platform—it might be sports, fashion or beauty—and we look at the way you engage with the content of the platform. If you look at a new video and you watch the whole video, or if you say that you are not interested in a video, or you like a video or comment on a video, that informs how our algorithm delivers content to you. Our ultimate goal with the For You feed is to deliver diverse content to our users.

Moving out more broadly to the way we approach trust and safety in our community guidelines, we have a hub approach to our trust and safety model at TikTok. We have a hub in Dublin, we have a hub in California, and we have a hub in Singapore. Our global head of trust and safety, who sits in Dublin, is our colleague Cormac Keenan, who comes to TikTok with lots of experience from other tech companies.

What I want to talk about with our community guidelines is what Henry alluded to with Snap's community guidelines. Similarly for TikTok there are guard rails. We have set out very clearly in our community guidelines what is and what is not allowed on the platform. We think this approach has been very transparent about what we expect from our users. It really creates a place where people can freely express themselves.

So we have guard rails that are set in place. We allow unpopular opinions on the platform, of course, and controversial opinions, but we do not allow hate speech, harassment, bullying and other types of content that we may be discussing today from a content perspective. We have a whole range of controls that we put in the hands of the users so that they can create an experience that they are comfortable with on the platform.

From a privacy perspective, we recently announced what we think is an industry-leading decision to make all accounts for under-16 year-olds on TikTok private by default. They cannot change that setting. If you want to use TikTok and you are under 16, your account will be private. Additionally, under-16s are not allowed to send a direct message on TikTok. Importantly, in Europe, anyone who wants to send a direct message has to be connected to that, so we do not allow unsolicited messages. I could go through a whole range of features that we have put in place to protect our users, both from a community guidelines perspective and from a user control perspective, but we may get into those features later.

One other thing that I want to mention briefly is that we always know that our work on trust and safety is never done. It is really important that we engage with experts and NGOs. We do that a lot in the UK. We spend a lot of time with the Holocaust Educational Trust, with Tell MAMA, and with other organisations such as Beat to talk about eating disorder

content. Having those conversations helps to inform how we set out our conduct moderation guidelines.

We are a short-form video platform. Videos are between 15 seconds to a minute. Our whole goal is to create positivity and joy. At some point, we called ourselves the last sunny corner of the internet, because we think that people come to TikTok to have fun and express themselves. Our utmost priority is to provide a safe experience for our users, where they feel free to express themselves in a way that is to their liking, with the privacy settings and the whole range of other features I have gone through.

I could probably go on for quite a long time about the things we are doing, but I will leave it there.

The Chair: We will come back to a number of the points, but thank you very much, both of you, for the very helpful introductions.

Q244 **Lord Griffiths of Burry Port:** I am delighted that connectivity and joy may be the outcome of the hour that we spend together. I will do my best to contribute to that.

The balance between privacy and freedom of expression is at the heart of all that we are trying to do in the various avenues of communication that we have been questioning people about. I must therefore start with you, Liz, because the former Children's Commissioner and the ICO are both involved in asking you some pretty critical and demanding questions about your use of people's data. It is a use of data that might lead to some people fearing that their data was going to be used inappropriately, and that would diminish their instantaneous and spontaneous desire to express themselves freely.

Can you help us to see how you have drawn that line? How have you given rise to the feeling that you may be commodifying, as they say, people's data and using it to create the fear or reluctance to express themselves freely? It would be very helpful if you could do that.

Elizabeth Kanter: Absolutely. It is a really important topic. I noted earlier that the committee has done inquiries on this issue over the years, so it is a very topical issue that we have been talking about for a long time.

As a platform, we take the privacy of our users and the control of their data very seriously. It is our utmost priority to be transparent with our users. We provide a privacy policy that is published on our website. We have also published a privacy policy for under-18s so that they can more clearly understand how we use our users' data. We are very open about that.

On giving our users control, if you go into the app, we have a particular setting where users can indicate what they are comfortable with us using. For example, a user can turn personalised advertising off in the UK if they so desire and they do not want that type of advertising. We would respect that user's choice. In addition, a user can download any of the data that we have been using for them through the app. We provide

a range of granular controls that users can take advantage of through the app.

More importantly, we publish a transparency report. It talks about government requests and gives a very detailed breakdown of the requests we have had from Governments for user data. We talk about the numbers in the UK and in a variety of other countries where we operate.

Finally, we have recently opened what we are calling a transparency and accountability centre. We are inviting regulators, legislators, the media and other stakeholders to come into the transparency and accountability centre to understand in a more detailed way how our algorithm works, how we handle data, how we handle security and how we handle privacy.

Broadly speaking, we look to be very transparent with our users about how we use their data but, importantly, giving them control about the use of their data.

Lord Griffiths of Burry Port: Henry, from the way you presented yourself you have been looking at precisely these issues and have come up with a formula that seems to me at any rate to be slightly differently centred. Are there any comments that you wish to add to what Liz said?

Henry Turnbull: Thank you, Lord Griffiths. I mentioned our first innovation on Snapchat, which was ephemeral messaging. In keeping with, and following on from, that first-ever innovation, ensuring the privacy of our users is fundamental to our approach at Snap. We practise data minimisation, which basically means that all our products and features are designed to collect and store as little user data as possible. That is something we take very seriously. We do not compile detailed social graphs of our users. We do not micro-target ads at them.

I mentioned our privacy and safety by design processes. Essentially, privacy and safety experts are involved in the development of any new products and features at Snap. All new products and features go through a dual review process, which involves teams from across the company, including legal and including trust and safety, which is the team that reviews user reports. They join reviews of products and features very early in their lifecycle, and consider the potential impact of features on users' privacy and on their safety to ensure that there are no unintended consequences of new releases that we have not thought through.

Importantly, those processes result in an approach and actual product features that protect the privacy and safety of our community, and that work to prevent abuse from the outset. I have mentioned a couple already. One example is our Snap Map feature, which allows users to see what their friends are up to on the map. Location-sharing in Snap Map is always off by default, so users can never share their location with anybody who is not their friend. The default settings are always off, so even if you have been sharing your location for a while with only a few friends, you will periodically get notifications reminding you that you are sharing your location and asking if you want to do that.

Essentially, this is something that we think about very deeply and is an area where it is totally important to us to protect the privacy of our users.

Lord Griffiths of Burry Port: Both of you have said things that would fill me with pleasure if I could ask supplementary questions, but I must yield the floor to others. Thank you both.

The Chair: In fact, that leads very neatly into the next question from Baroness Buscombe about what sort of assessment is made of new products.

Q245 **Baroness Buscombe:** Thank you both very much indeed. In fact, your intros make me feel as if you have the perfect site, but let us delve a bit further. It is very encouraging to hear about the guard rails, which you both referenced, that you have put in place.

I want to ask about human rights impact assessments. Several witnesses that we have heard from so far feel that all platforms should be required to go through an impact assessment of all their new products—you have just been talking about your new products—preferably at the design and development stage. We heard from one witness, who said that there was really only one platform, Microsoft, that actually carried out such an impact assessment and had a proper human rights programme. As that person put it, it was run by a senior or with a senior in charge who knew what they were doing, and that is it. Could you explain to me whether you have one, and what form it takes. Henry, would you like to go first?

Henry Turnbull: Sure thing. Thank you, Baroness Buscombe. Speaking in the context of human rights, I will talk about two areas. One is how we use those considerations to guide decision-making on content, and then I will refer to the processes that we have in place on designing new features.

When it comes to decisions on content, our community guidelines, which I mentioned, form the basis for any actions that our team takes in respect of illegal and harmful content. The guidelines are developed by our internal policy team in collaboration with teams across the company. They are designed to encourage people to express themselves freely with their friends, while also ensuring that Snapchatters can use the app safely. The guidelines are designed to ensure that Snapchatters' rights are protected and that they are not bullied, threatened or discriminated against on the basis of their race, ethnicity, nationality, sexual orientation, et cetera. It may not be framed in the exact language of human rights, but the guidelines are built with those things in mind and we think about them when we are ensuring the safety of our community.

Your second point as about the design process. I will not reiterate too much what I just said about privacy and safety by design, and the kind of processes that are applied with teams from across the company thinking very carefully about the development of new products and features through a safety and privacy lens. Certainly, that kind of collaborative exercise is designed to ensure that there are no unintended discriminatory consequences of the design decisions that we take, or of

our products and features. Again, it is possibly not framed in the exact language of human rights, but these are all things that we are thinking about as products are designed and as we design our policies.

Baroness Buscombe: That is very interesting. You said that it was framed in the terminology of human rights, which might mean something very different to one person from another person. I think that is a good point to make. Thank you. Liz, over to you.

Elizabeth Kanter: Absolutely. I was looking at a previous session you held. Seyi from Glitch had a really great principle that we think is one of the principles that we use: the freedom to express yourself does not equate with the freedom to abuse online. I thought that was a really powerful summary of what we try to achieve with our approach to the way we draft our community guidelines.

You mentioned that Microsoft has some senior bodies at the company who work on human rights. We have colleagues working across the company who have a very long history of working in human rights. We, of course, engage with experts outside the company. We have recently joined an organisation called Business for Social Responsibility—BSR. We sit on its human rights working group, and this is a way for us to engage with experts on the subject of human rights.

The way that we come at this for the platform is not unlike what Henry was talking about. We look at it from the platform side and then how we empower our users from a policies and processes side. As I said before, we have our community guidelines, which set out what we expect on the platform. We have very clear rules about not tolerating hate speech, harassment and bullying. We know that with harassment and bullying in particular there are certain communities that feel they need to be protected by platforms. By having a really strict approach to harassment and bullying, we think that people will be freer to express themselves and to express their views on TikTok.

We also have a range of features that we can put in place in an age-appropriate way. When we look to roll out features that are age appropriate, we look at things like the UN Convention on the Rights of the Child to inform our approach to those features, such as our move to make under-16 accounts private by default.

From the user point of view, we think that the controls that we have given our users really help them to express themselves and take a human rights-based approach to their experience on TikTok. I will cite a couple of examples. We have a comment filter that we offer our users. From the platform side, we have an ever-evolving list of predefined key words that we use to eliminate negative comment from the platform. We also give our users the ability to have a list of 20 key words that they do not want to appear on their TikTok account. We think that is an important way of empowering our users to create an experience that is to their personal preference.

In addition, we have in-app reporting. If a user sees something they do not like, they can report that content. Importantly, in our processes we

not only have the in-app reporting but provide our users with the ability to report content. Any type of content on TikTok can be reported, and our human moderators look at that content. If the user is not satisfied with the way we have addressed their report, they can appeal. We take that very seriously and our human moderators will look at those appeals.

I mention that, because I think it is a really important right for individuals to be able to express themselves, but also to know that they have the protections in place from the platform to raise concerns about content that they might find objectionable. Between our policies and our processes, we take a human rights-led approach to deal with the way our content appears on the platform.

Q246 **Baroness Buscombe:** I will be quite brief with my supplementary, because we have a number of questions and I do not want to hog the limelight talking to you both.

Liz, you talked at the beginning about TikTok perhaps being the last sunny corner of the internet. You also said that you were not afraid to allow controversial content. Clearly, you are pro freedom of expression and so on, but, for example, TikTok has been praised by a Russian state official, who stated that it "actively cooperated with us, which cannot be said about others". I also found quite disturbing that in March 2020, as reported in the *Guardian*, TikTok "tried to filter out videos from ugly, poor or disabled users". That does not sound like a sunny corner of the internet to me.

Would you like to explain just those two examples? I have several in front of me, but time does not allow us to spend too long on this.

Elizabeth Kanter: I would be delighted to. On the second point about the *Guardian* article, those were content moderation guidelines that have not been in place at the company for a very long time. In the early days of TikTok's development we took a blunt instrument to some of the content. The drafters of those community guidelines thought that they were protecting those communities and safeguarding them by not allowing that content. It is absolutely allowed on the platform now. Those are outdated community guidelines that were referenced in that *Guardian* article.

On the question about Russia, I will try to be as fast as I can in explaining to you the origin of the quote that you read. As everyone knows, there were protests when Alexei Navalny returned to Russia earlier this year, and videos were posted on platforms, including TikTok. Those protests were held during Covid restrictions and without proper legal permission, so they were deemed illegal by the Russian Government. Pro-Navalny content and pro-Putin content is not illegal, but the protests in particular were illegal according to the Russian Government.

We received a legal request to take down the content from the Russian Government. As we do with any kind of content we receive from any Government across the world, we evaluate legal requests for taking down content. With regard to this particular content, we took down some

content. We were ultimately fined by the Russian Government for not moving at a pace that they felt was fast enough. They also did not feel that we took down enough content, so we received a fine. After the fine was issued, we were asked to meet with Russian parliamentarians, in the same way that we are meeting with you today, to talk about how we approach this type of content. We attended that meeting, and a press release was issued afterwards to thank us for participating.

We want to make it very clear that we did not do anything different for the Russian Government from what we would do for any other Government around the world. We did not change our policies to accommodate the Russian Government. We looked at the valid legal requests, evaluated them and the content that they flagged, and, as I say, we took down some content. We certainly did not do anything special for the Russian Government.

Baroness Buscombe: That is very helpful. Of course, it helps to show that these issues can be quite complex and nuanced. Just suggesting that something is not quite right is not always straightforward, so thank you both for that.

The Chair: Unfortunately, perhaps, you are the first evidence givers we have had who have seen the Government's proposals, now finally published, for the Online Safety Bill. I am sure you have spent most of the last 24 hours reading and inwardly mastering every conceivable nuance, comma and question. I do not think we want you to display that expertise at length to us today, but the context is right, because the next two questions are really about where our Government want to steer. We would like your reactions to that, so I ask the Bishop of Worcester to put his point.

Q247 **The Lord Bishop of Worcester:** Thank you, Lord Stevenson, and thank you both witnesses for being with us this afternoon and for what you have said. I echo Baroness Buscombe's suggestion that it has generally been very encouraging.

As Lord Stevenson said, I want to press you on the balance of freedom of expression. Liz, you suggested that it is freedom of expression on the one hand and abuse on the other, but what we have been grappling with is freedom of expression on the one hand and harm on the other, which is perhaps more difficult to assess.

The Bill has finally been published. In the White Paper the Government defined harm as content that "gives rise to a reasonably foreseeable risk of a significant adverse physical or psychological impact on individuals", but Clause 46 of the Bill as published adopts an altered definition, which might be thought to increase concern about freedom of expression. The provisions that now apply are that "the provider of the service has reasonable grounds to believe that the nature of the content is such that there is a material risk of the content having, or indirectly having, a significant adverse physical or psychological impact on a child of ordinary sensibilities".

Some of the terms that appear in that provision are not properly

identified. Among our witnesses there has been general concern about what was expressed in the White Paper. It was thought by some that regulating harmful but legal content should be affected by statute; in other words, if something is sufficiently harmful, it should be transparently and democratically deemed illegal.

There is also the question of how difficult it will be for platforms to fulfil obligations to delete illegal content while protecting freedom of expression. There is the equal and opposite danger of too much material being brought down. This is something that we have been grappling with.

I would really value your thoughts from the perspective of your own platform, and platforms in general, as to whether the proposed requirements on illegal and legal but harmful content, while protecting freedom of expression, are reasonable. Would it be possible for platforms to satisfy the requirements? Can I ask you first, Liz?

Elizabeth Kanter: Of course. I have not quite managed to read the 144 pages of the Bill just yet, but we are of course aware of the proposals on legal but harmful. The way the Government have proposed approaching child sexual abuse material and terrorist content is through the two codes of conduct. We have seen the draft codes and we think they look to be sensible, because they give a lot of clarity and detail about the expectations that the Government have for platforms to uphold the obligations to remove illegal content.

I think you have heard from other witnesses on legal but harmful. It is much more difficult to legislate on those types of harms. If the Government and Ofcom want us to act with the same efficacy as we do with removing illegal content on legal but harmful, we are going to need very clear guidance on what the expectations are, so that we can meet our regulatory obligations.

As a platform, we are already looking at legal but harmful. We have published in our transparency report categories of content: harassment and bullying, minor safety and illegal goods. We publish figures on the legal but harmful categories that we look at. If the Government want to regulate for legal but harmful, we need a lot more clarity for our moderation and the way we approach that content. As you say, we need to get the balance right between allowing that kind of content and allowing for freedom of expression. It is a difficult balance. Ultimately, we need legal clarity from Parliament and from Ofcom so that we can achieve what they want us to achieve through the Bill.

The Lord Bishop of Worcester: Before I ask a supplementary, with the Chair's permission, do you have any thoughts to add, Henry?

Henry Turnbull: Absolutely. I should also say that, unfortunately, I have not had a chance to read through the entire Bill and the Explanatory Notes just yet.

Overall, we support the aims of the Government's strategy to protect people from illegal and harmful content. We have been in constructive dialogue with the Government for a long time now about how best to

produce effective regulation that improves user safety but is also proportional and practical for the vast variety of different online platforms and services that exist.

On the Government's strategy as set out in their full response to the White Paper, we support the key planks of the policy outlined there too, particularly the idea of principles-based regulation, based on the statutory duty of care and enforced by an independent regulator in the form of Ofcom.

On your question about legal but harmful content, as I understand it, most of the categories of harm that have been outlined by the Government so far, whether that is false information, self-harm or suicide content, are already prohibited on Snapchat under our community guidelines and we already take action against them. To an extent, it is a bit of a moot point for Snap.

More generally, our concern would be the difficulty for Ofcom in deciding when a piece of content is legal but harmful, particularly if the category is vague or relatively subjective. On many occasions, that might be a fairly subjective decision for Ofcom. I think context can play an important part in any decision, as this committee would recognise. This is one of the reasons why the European Commission considered, but ultimately rejected, the inclusion of legal harms in its Digital Services Act proposal.

An interesting parallel is the approach that they are taking in Ireland, as part of the implementation of the audiovisual media services directive, through their own online safety Bill. The Irish Government have essentially attempted to provide clarity for platforms by looking to define clearly what they mean by harmful online content in legislation.

As I understand it, from a very brief look through the Bill and the memo over the last day or so, the Government are proposing to do that via secondary legislation in due course. Providing that kind of clarity, as Liz said, would be helpful for platforms as we develop our own policies and processes.

The Lord Bishop of Worcester: Thank you both very much indeed. With your permission, Chair, I will ask a brief supplementary, and I ask the witnesses to be brief in their responses.

There are two points. First, witnesses have agreed that the scale of online user-generated content requires platforms to use algorithms for content moderation. Of course, they are unable to identify context, nuance or irony. Are you confident that your monitoring will be sufficient?

Secondly, you both mentioned Ofcom, but the onus will be on the platforms to make sure that none of their content will directly or indirectly have "a significant adverse physical or psychological impact on an adult of ordinary sensibilities". That is straightforward in some of the examples you have mentioned, but let us take the example of trans rights. What would be the way forward on that? It is a highly contested issue. JK Rowling got herself into an awful lot of problems for expressing

one side of an argument. Trans people could arguably say that they were subject to significant psychological impact as a result of that. This is where we get into the really tricky area of legal but harmful. Do you have any thoughts to offer on that?

Henry Turnbull: The main point I would make, my Lord Bishop, is that not all platforms are the same, and not all platforms are open town-square-style platforms that rely very heavily on automated moderation. Snapchat, as I have explained, is predominantly a private messaging service, and to the extent that we have a public space for content it is curated and moderated before it can surface publicly.

The core determinant for us is whether something violates our community guidelines, which are really clear and simple about what kind of activity is prohibited—if it is violent content, if it is false or misleading, et cetera. I appreciate your issue on trans rights is an edge case, but it does not instinctively seem to me to be something that would violate our guidelines.

Of course, there are some complexities around legal but harmful, but as long as the framework enables platforms to enforce their guidelines in a clear and transparent way that people can understand, that is the most important thing. It seems that the framework will allow for that, but we need to go into the detail.

The Lord Bishop of Worcester: Thank you. Liz?

Elizabeth Kanter: As two specific examples of what you are speaking about on legal but harmful, we use human moderators, and we use machines. We have over 10,000 moderators globally, and they think about context and local nuances as they moderate content.

As an example of the grey area, we did a lot of work over the past year with the eating disorder charity, Beat. We talked with people who have lived experience of eating disorders, so that we can look at how we craft our community guidelines and moderate the content related to eating disorders in a way that allows people to express themselves, but does not get into content that might trigger someone who has an eating disorder. We are very respectful of our community and some of the challenges they might find.

On your very specific example of a transgender issue, as a platform we have recently updated our hate speech policy. Rather than having a hate speech policy, we have a hateful ideologies policy. One of the facets of that policy is that we do not allow the deadnaming of transgender people. From our perspective, there is no need for that kind of content on our platform. We have made it a violation of our community guidelines. We think that we can go further on hate speech to protect our users on our platform.

The area of legal but harmful is a conversation we are looking forward to having with government, so that we can try to craft regulations that we can ultimately implement. As you say, the challenge is making the words on paper come to life in what the content looks like when it is actually on the platform.

The Chair: Can I ask Baroness Featherstone to come in now?

Q248 **Baroness Featherstone**: A number of witnesses have brought to our attention that the more regulations that are laid on social media platforms, the harder it will be for new entrants to enter the market. If competition is a good thing, and we think it is, I would like to know your opinions on that.

Elizabeth Kanter: One thing I can say is that I think the growth of TikTok in the UK over the past couple of years shows that competition is possible. One of the ways that we think we have had a competitive advantage over some of our colleagues in the industry is that we have taken the safety by design approach. We hope that the Online Safety Bill will ultimately drive others to see that as a competitive advantage going forward.

As regards the potential impact on new entrants, I agree that regulation can put a burden on smaller companies. Some of the requirements on transparency of reporting are very binary in what the Government expect from platforms and may be difficult for smaller platforms to uphold.

We have seen a lot of the one-size-fits-all approach to regulating online safety and regulating in the tech space. TikTok is a video-sharing platform. We cannot be regulated in the same way as a platform that is more image based. We would encourage the Government, as Henry mentioned, to take an outcomes-based, evidence-based, proportionate approach to their regulation, really focusing on the end result and the end goal of regulation, and leaving it to the companies to be empowered to determine how they will achieve the objectives of the regulation.

In brief, I think there is potential for regulation to have an impact on competition. From what I have seen, as I have scanned through the Bill so far, it seems that the Government are aware of that and are looking to take a proportionate approach to allow for new entrants to come into the market.

Baroness Featherstone: Thank you for that. Henry?

Henry Turnbull: Thanks, Baroness Featherstone. I echo Liz. This is why the style of regulation that is developed and implemented is really important. We all accept the need for online regulation. At Snap, we support the Government's overall strategy, but there are potential risks to the competitiveness of digital markets in the UK if we do not get it right. From our perspective, regulation in this area is most effective when it focuses on the principles and the outcomes that companies should deliver, setting out what objectives are to be achieved without being too prescriptive about how companies should achieve them.

As I said before, there is such variety in the size, resources and service models of different online platforms that a one-size-fits-all approach, as Liz termed it, will never work. Under a more flexible principles-based model, a regulator can examine a platform's compliance with the outcomes as relevant and as appropriate for their services. Ultimately, the companies best served by overly prescriptive, complex regulation are

the largest firms with the largest compliance teams that can easily deal with the bureaucracy involved, which smaller companies—I am thinking in particular of start-ups and similar sizes—would struggle to comply with.

The good news is that the Government have committed to proportionate principles-based regulation. Obviously, we need to look at the draft Bill in detail. Ofcom has committed to that as well. It is about making sure that that happens in practice. The Government have so far released two interim codes of practice, last December. One of those alone is 50 pages long. That is one code for one harm. A smaller company's chances of digesting that and complying with it, as a package of potentially one of several codes, is really limited. That is the kind of regulation that can lead to market competitiveness issues.

We are hopeful that Ofcom will aim for as much simplification of different codes as possible when it takes on its regulatory role, particularly thinking about the impact on smaller companies. Having one principles-based master code that applies to all platforms is probably the best model.

Baroness Featherstone: You are saying that, in principle, everyone should adhere to that principle, regardless of size.

Henry Turnbull: If you have principles-based, proportionate regulation that is focused on overall outcomes which the platforms should achieve as relevant to their service models and resources, and you give Ofcom a bit of leeway to examine platforms according to their risk profile, I think that model works. It is just about how it works in practice.

Baroness Featherstone: That is very helpful. Thank you.

Q249 **The Chair:** My question is a follow-on from where you have just ended with Baroness Featherstone. I presume that the problem you will face is not just that you will one regulator, Ofcom, which we now know is the regulator, and that is a good thing, but the fact that you will also be subject to regulatory approaches from three or four others, which are not always the same and of course have different cultures.

One of the new ones is the digital markets unit, which seems to come with a heavy brand of economic regulation. It is not quite what we were talking about before, with principles and allowing companies to respond to those and then judging how well they respond to them. That does not quite fit with the regulatory structures that are in place in the CMA and that presumably will be in the digital markets unit.

I have two questions within one. You presumably have had a look at the DMU and its work. Do you have any thoughts about that and the impact it might have? More specifically, given that we are in a rather strange market, in the sense that there are very few large companies, and not very many smaller ones, how do we find competition acting as an aid to consumers? Do you have any views on that?

Henry Turnbull: We have had a number of reviews in this area in the last couple of years in the UK, including the 2018 Furman review and the

CMA's resulting market study of digital markets in the UK. The ultimate conclusions from those reviews—that the evolution of tech markets in the UK is exposing some shortcomings in current competition rules and that there should be a refresh of those—are clear and widely accepted. From our perspective, it is a good thing that the DMU is being established. In response to your second question, it should create a better environment for consumers over time.

To reiterate, our view is that it is a good thing that the DMU is being established. The key objectives that the CMA originally identified for the DMU were the right ones, particularly establishing a new legally binding and enforceable code of conduct for firms with strategic market status, and exploring measures to restore competitiveness to the market. Mainly, it is important that the DMU is given the resources that it needs to fulfil its new role and to get up and running quickly. As the committee knows, digital markets move and evolve very quickly.

Elizabeth Kanter: I echo what Henry said about the DMU. Of course, we looked at it. We think that the pro-competitive regime it is looking to establish is very welcome in the UK. Obviously, we have not seen the legislation that sets out its statutory powers, but we will be watching it with interest and hoping that it has a positive impact on competition in the UK.

The Chair: May I press you a little on that? We are able to be positive about it, because we do not yet know what its powers will be, but we can imagine what some of them might be; they might, for instance, include some form of attempt at interoperability. How would you view that sort of approach, if that was the way they were thinking?

Elizabeth Kanter: You are right: we know some of the proposals. We have seen the work of the CMA. We are looking at something like data interoperability and considering whether it is something that we would support. We are genuinely supportive of anything that will inject competition into the market, and if that is the best way to do it we would support it.

The Chair: That is useful to know. Henry?

Henry Turnbull: These are a couple of the proposals that we have a few more concerns about, Lord Stevenson. I mentioned that we practise data minimisation at Snap. We are extremely careful about our approach to data. All our products are features designed to collect and store as little data as possible, so when it comes to proposals on sharing data we are always very cautious. That is not something that we have any intention of doing now.

Some of the ideas I have heard are about introducing common data standards or mandating platforms to share data. Perversely, mandating these new standards could make actually it more difficult rather than easier for smaller companies to enter the market, as well as reducing innovation in new business models. That is because mandated baseline standards could theoretically become the upper limit, and there would be no incentive for companies that offer different or innovative approaches.

That could therefore reduce innovation, reduce investment and actually reduce competition in the sector in favour of the largest incumbents.

I do not want to give it a firm no, or say that we are fully against it, but we have some concerns about it, which we have registered with the CMA and others.

The Chair: But you would, wouldn't you? I can quite understand where you are coming from, and there is no criticism of you for setting it out so clearly. Given that there is a regulator, it sees the problem that we have both identified as the possibility of consumer detriment, not just economically but perhaps in freedom of information circulation or human rights issues as well. Will there not be a tendency to begin to want to move on companies in order to make them less obstructive to the idea of new incumbents and new ideas and innovation?

Henry Turnbull: I absolutely support the establishment of the DMU. I think it is necessary and that there are a lot of measures that it could pursue, whether related to merger control or other things that could have a real and significant impact. I just think that, when it comes to mandating common standards for how all platforms should hold and store data, which might risk affecting innovation, and mandating the sharing of data, which involves a whole range of risks, that is an area we would want to explore carefully. I would personally counsel looking at some of the measures that have been applied in other countries, rather than necessarily looking at interoperability.

The Chair: Thank you very much.

Q250 **Baroness Bull:** I want to turn to the question of digital citizenship online. You both spoke in your opening remarks about features on your sites that are designed to make a safe environment to ensure that people can find and hold their voice. I do not want to invite you to repeat yourselves, but perhaps I could ask you whether there is anything inherent in the design of your platforms or your business models that you worry discourages good citizen behaviour online.

Henry Turnbull: It is a very interesting question, Baroness Bull. There is a misconception about Snapchat that the fact that messages are ephemeral—that, generally speaking, they are not viewable after you receive them the first time, and the same is true with snaps—can encourage harmful or illegal behaviour. That is something that we have put a lot of effort into addressing. It is very easy to report that kind of activity when it happens, if it happens. It is definitely not the case that, just because messages are ephemeral, there is nothing we can do about harmful or illegal content. We have worked quite hard to address the misconception about that.

Perhaps I might give the example of a design decision that we have taken in a related field that looks to nudge people away from harmful activity. I mentioned that snaps are ephemeral. Most smartphones have a screenshot capability; it is embedded in the smartphone rather than in an app. That is not something we have any control over at Snapchat, but

we absolutely do not want users taking screenshots of other users' snaps without their permission.

Unfortunately, we cannot stop them doing that, but if somebody takes a screenshot of another person's chat, snap or story, we notify the user that a screenshot has been taken. It will also be made clear to the person who has taken the screenshot that the other person has been notified. That kind of notification, and knowing that it can happen, can guard people against that sort of behaviour.

I think there are some misconceptions about the service that we look to address, and we look to design additional guard rails and put protections in place as a result.

Baroness Bull: But do you know that that sort of post hoc wrist slapping, or threat of post hoc wrist slapping, discourages people from bad behaviour? Do you have any evidence of that?

Henry Turnbull: We conduct very regular research and analysis of our users, how they are engaging with the app and their attitudes towards it. That certainly informs the development of these features. The screenshot notification feature has been in place for years. If it was not having an impact, it is something that we would have discontinued. I think that is really valued by the community, because if somebody has taken a picture of their chat without their knowledge, they now know about it. I think that is something that certainly is having an impact.

At the start of the session, I mentioned some of the other design decisions that we have taken, so I will not elaborate on them. We are very thoughtful about this stuff, and it is embedded into the design of the app right the way through.

Baroness Bull: Thanks, Henry. Liz, thinking about other social media platforms where we know there is better behaviour—the ultimate would be LinkedIn, where everybody is incredibly polite—what do you think are the features that encourage that? How do they compare with the features in the design of your platform and your business model?

Elizabeth Kanter: It is interesting to look at LinkedIn. You are absolutely right; that is a place where people communicate about professional achievements and things like that. With TikTok, we take very seriously the idea of digital etiquette. I know that has come up in committee hearings previously.

We introduce a few different nudges or friction into the user journey that we think are really helpful in encouraging people to have more respect for one another. One example is our work on misinformation. We have introduced a label that we put on a video when we cannot verify whether or not it contains factual information: "This contains unverifiable information". We have seen that one in four people see that label and do not share that video. They do not engage with that content. That shows the efficacy of the creation of friction.

We have another feature that you could liken almost to a personal moderation cue. Individuals can turn on a control feature or comments

feature whereby they can pre-approve every comment that appears on their TikTok video. If they do not like what that comment says, the comment will not appear in comments on their video.

We have really tried to put those controls in the users' hands to complement the work that we are doing in our community guidelines on hate speech, harassment and bullying. I love the expression "digital etiquette" that has been used in the committee. It is something that we take a lot of pride in. We will continue to innovate in this area, but those are two of the features that we are using to create friction to promote good digital citizenship on the platform.

Baroness Bull: You mentioned the word "verify". In 2020, Ofcom found that 42% of eight to 12 year-olds used TikTok despite the minimum age requirement. Perhaps I could ask you—and, if we have time, Henry—how you verify your users' ages.

Elizabeth Kanter: In TikTok, the first time they want to use the app, we ask the user to indicate their age. We do not nudge them to say that they are 13. If a user indicates that they are under 12, they will not be allowed to create an account. If a user is not honest about their age, our moderators are trained to identify underage users. We flag underage users and remove them from the platform and ban their accounts.

Yesterday, or this morning, we announced that we will be the first platform to publish the number of underage users that we take off our platform. In our next transparency report we will show that we take this very seriously. We do not necessarily want underage users on our platform. We do not think that our platform is appropriate for under-13s, so we are taking really strong steps to take them off the platform. As I say, we will be publishing the numbers on removal in our next transparency report, which will come out in the next couple of months.

Baroness Bull: Thank you. Henry, the same question to you on age verification, please.

Henry Turnbull: Similarly, I should be clear that we do not want under-13s using Snapchat. If we become aware that someone using Snapchat is aged under 13, we will delete that user's account. There are also other measures that we can take, such as blocking their device.

The way it works on Snapchat is that, in order to create an account, a user needs to enter their date of birth. If they enter a date of birth that is below the age of 13, they will get knocked back. We do not say, "This is because your entered date of birth means you are younger than 13". However, I accept that somebody could reasonably infer that they are getting knocked back because of the age they have entered, and they could try a different age.

We are committed to continuing to work with the Government and a whole range of other stakeholders to explore approaches to age assurance, and have been for several years, but there are real challenges. ID-based age verification requires the collection and retention of things like passports and drivers licences. The ICO has highlighted the risks of that from a data perspective. There is new

technology in development, but, to be honest, its accuracy and scale are unproven. It has some issues.

There are some existing pinch points in the short term that could have a faster impact than asking individual apps and websites to develop their own proprietary solutions. For example, everybody has to go through the current app stores, so there is an existing pinch point for installing apps on their phones. Introducing an age gate at sign-up to those would be an effective, scalable tool to ensure that children only access apps that are age appropriate, rather than asking all individual apps and services to go about their own regime.

I think there is more that we can achieve in the longer term. There are conversations going on, but that is just an idea to think about. We have mentioned it in other forums.

The Chair: We need to move on. Thank you for that.

Q251 **Lord Vaizey of Didcot:** It is very encouraging to see that you are taking age verification seriously. It is an incredibly important issue. Obviously, industry working with government on age verification tools would be very good. I know that TikTok's age verification algorithm is incredibly effective.

I am afraid that I have only three minutes to raise the most important issue, which has been left to me, and that is to talk about the Government's work on digital citizenship. I have to confess that I find the whole phrase "digital citizenship" a complete and utter turn-off and totally meaningless. I understand why Liz is attracted by "digital etiquette", because it is more engaging.

We are, nevertheless, waiting for a government proposal on digital citizenship, obviously particularly focused on schools but potentially on adults. Henry, you have worked in government, and, Liz, you are extremely experienced in working in technology. If I appointed each of you as the Minister, what would be your proposals for government-led digital citizenship?

Elizabeth Kanter: Thank you very much, Lord Vaizey. That is a huge topic, and, as you say, two minutes are not really enough to discuss it. We are eagerly awaiting the Government's media literacy strategy, but if I was Minister for the day, I would probably replicate what TikTok recently introduced in its own media literacy campaign, launched just a couple of weeks ago.

The first focus we have is on critical thinking, to try to encourage people not just to sit back and consume content but to think critically about what they are engaging with and seeing on the platform. That is something we introduced. The engagement levels are really high for that content.

Over the coming months, we will be looking at the Covid vaccine. We will be doing a partnership with the Government on that. We are also looking at news literacy on a whole range of themes. I would be very happy to write to you and share more with you about that. But that is what we

would do: we would take a very broad-brush approach to media literacy in schools and in the public.

Lord Vaizey of Didcot: Henry?

Henry Turnbull: Thanks, Lord Vaizey. I am sorry to hear that you do not like the term "digital citizenship". It is one that we are actually quite a big fan of at Snap. I think that—

Lord Vaizey of Didcot: Come off it, Henry. You are a fan. I cannot believe that anyone else at Snap is a fan of that. I cannot believe that people at the Snap boardroom table are happy to talk about digital citizenship.

Henry Turnbull: I think it is important to have people reflect on their impact on other people online. Better online education and promoting some form of responsibility for people online is a really important part of improving the online experience and the online safety of people in the UK. It is an uncomfortable truth that the vast majority of what we would call online harms are ultimately perpetrated not by hardened criminals but by ordinary citizens. I think we all have a bit of a duty to reflect on that.

Super briefly, I think there are a couple of elements. One is helping young people to consider what it means to be a citizen in the digital world and to understand their responsibilities and the consequences of their actions on other people online. The other is being aware of the tools and resources that are out there, what to do and how to act effectively if you come across abusive or illegal content. Research unfortunately shows that most people are often very reluctant to report harmful content online, regardless of the platform that they are on. They often think that nothing will happen as a result, which is not the case.

This is something that we have been working on at Snap. We have launched a dedicated safety channel that aims to raise awareness of the issues and challenge misconceptions, but there is a role for education as well. I understand that this is now a mandatory part of the curriculum. If those two areas can be brought into schools, I think that will be really important.

Finally, at Snap we work really closely with some excellent UK-based charities, including South West Grid for Learning, and Childnet, which both do a great job educating young people and parents about life online, the challenges and the risks. I hope that those expert organisations are being consulted and can play a role in this kind of education going forward.

Lord Vaizey of Didcot: Are you going—

The Chair: We have to draw things to an end. I am sorry for cutting you off, Ed.

Lord Vaizey of Didcot: I cannot believe it.

The Chair: Ed, you had such a good question that we had to end there; we could not beat it.

Thank you very much indeed to both our witnesses. It has been a very lively and very interesting session. We have enjoyed it very much. I am afraid that we have run out of time, so I am just going to say goodbye and thank you very much. The meeting is now concluded.