

Select Committee on Communications

Corrected oral evidence: The internet: to regulate or not to regulate?

Tuesday 8 May 2018

3.30 pm

[Watch the meeting](#)

Members present: Lord Gilbert of Panteg (Chairman); Baroness Bertin; Baroness Benjamin; Baroness Bonham-Carter of Yarnbury; Lord Gordon of Strathblane; Baroness McIntosh of Hudnall; Baroness Quin.

Evidence Session No. 4

Heard in Public

Questions 28 - 34

Witnesses

I: Rachel Coldicutt, Chief Executive Officer, Doteveryone; Julian Coles, Independent Digital Media Policy Consultant; Dr Konstantinos Komaitis, Director of Policy Development, Internet Society.

USE OF THE TRANSCRIPT

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

Examination of witnesses

Rachel Coldicutt, Julian Coles and Dr Konstantinos Komaitis.

Q28 The Chairman: May I welcome our witnesses to this evidence session of our inquiry into regulation of the internet? Thank you to our witnesses for joining us, agreeing to speak to the Committee and bringing your expertise to inform us.

The meeting is being broadcast online and a transcript will be taken. There is a possibility that our session will be disturbed by votes in the House of Lords Chamber. If that happens, we will adjourn for about 10 minutes and ask you to wait for us to come back.

May I ask our witnesses to briefly introduce themselves and tell us a little about their background? In so doing, so we know where each of you is

coming from, could you tell us what your thoughts are on whether the internet needs further regulation? If it does, from your experience what type of regulation should it be: some form of self-regulation, co-regulation or a more directed form of regulation? When we have done that, members of the Committee will ask a series of further questions.

Rachel Coldicutt: I am the CEO of Doteveryone. We are a think tank that champions responsible technology for the good of everyone in society. We think of responsible technology as that which is good for everyone, that considers its impact, understands its consequences and seeks to mitigate those. The chair is Baroness Lane-Fox, who I am sure many of you know.

We spent two years looking at how technology is changing the world. We are pretty unambiguously in favour of regulation. We have backed it up by asking the British public their opinion and by trying to understand their attitude to technology and their understanding of it. There is a clear appetite among the public for additional regulation, too. From our perspective, it cannot be solved simply by self-regulation, and there is potentially an issue with the Government taking on the whole regulatory matter, not least because the Government's views of the internet probably also need to be subject to regulation.

We are in favour of an independent body that understands how technology works and is supported by public education. It needs to be a body that people can turn to, because we have heard that nobody knows who to ask on standards and safety tests.

I would add two things. The first is that there has been an enormous amount of conversation about social media, but ultimately it is the tip of the iceberg. There are an enormous number of other potential harms. Whether it is to combat addiction, bias, discrimination or democracy hacking, the internet has to be regulated. There are probably three areas to regulate—big tech, emerging tech and public sector tech—and they all need slightly different approaches.

Dr Konstantinos Komaitis: Hello, everyone, and thank you very much for inviting me here. I am the director of policy development for an organisation called the Internet Society, and I will say a few words about it.

We were established in 1992 by the very people who created the internet, Vint Cerf among them. At the time, the reason was to provide the organisational home to a community of engineers called the Internet Engineering Task Force, which was creating the internet and the standards that made the internet grow and evolve. Over the years, the Internet Society has grown to become an organisation that touches on three areas: policy, development and technology. We are a mission-based organisation. However, we have individual members, chapters in most countries around the world as well as organisational members. As an organisation, we work to ensure the internet stays open, interoperable, global and secure.

On the question of regulation, I would like to make two comments. First, there is a misconception that the internet is not regulated or has never been regulated. That is not quite the case. There has always been regulation of the internet, whether we are talking about copyright, intellectual property, data protection and so on. On top of the regulation that exists, people have always come together and collaborated in a form of self-regulation, or even hybrid forms of regulation and in the development of norms. Especially when it comes to the latter, we are seeing more and more communities coming together to produce those norms. A clear example is the Global Commission on the Stability of Cyberspace, which is a multi-stakeholder group consisting of Governments, businesses, civil society and the technical community. They came together to produce norms that hopefully will be adopted when it comes to the security of cyberspace. They have just released a norm calling for the protection of the public core of the internet.

The second issue is that when we use the phrase "internet regulation", it is a little wide and we need to be a little more specific. The infrastructure of the internet, which essentially is the backbone of the underlying communications that facilitates the sending and receiving of packets of data, needs to remain open and free. There is no reason to regulate that, because voluntary standards support this infrastructure. They are market-based, they come out of multi-stakeholder consultative processes and they ensure that the internet grows and the technology grows with it. Regulation at that layer of infrastructure is not optimal and is not advisable.

However, as we move through the layers, this is where we see the behaviour that has generated calls for regulation, especially lately. There the question becomes how we can regulate it in a way that will not prohibit innovation and creativity. We can discuss this later in more detail. The question is not whether we should regulate but where and at what layer. As I said, we see particular behaviours at the application layer, and the question is who or what should be doing that.

As a last point, I would like to say that the internet is a by-product of collaboration. When it emerged, a lot of people came together because they wanted to create something that would allow them to communicate: essentially, to send packets of data from point A to point B. That was all that it was meant to be. Of course, right now we are in a completely different era, but the technologists and engineers working on the internet still view it as this idea that we want to transfer packets from point A to point B.

Where feasible, regulation of the internet needs to be part of a collaborative and informative process. It is true that the internet is a complex ecosystem. The input of the different stakeholders, especially those who have the experience, is very significant, and I would like to congratulate you for opening up this consultation process and seeking those inputs. Thank you.

Julian Coles: Thank you, too, for the invitation to appear here. I am an independent digital media policy consultant. I am sorry for the mouthful there. I have worked as a consultant with DCMS, Ofcom and the European toy industry. I am a trustee of Childnet International and I am on the Ethics Committee of the Internet Watch Foundation. For about 15 years, until three years ago, I looked after the BBC's editorial policy for its online and interactive services.

This is a wonderful open question. If you will allow me to, I would like to jump backwards a little and come from the past to the present. There is a great spectrum between self-regulation, co-regulation and statutory regulation. If we go backwards to where we have come from, I would start the clock running at about 1996 when the Internet Watch Foundation was created. One might call that a broadly self-regulatory initiative and one that in its own terms has been extremely successful at driving down the number of child abuse images hosted on UK servers. That is just one example.

I remember that when the Home Office task force came into existence in 2001, it was a multi-stakeholder approach, sponsored by the Home Office, to look at harm and offence and other things online. We churned out good practice guides on instant messaging, chat and moderation. They were pretty good and well worth having. That turned into the UK Council for Child Internet Safety, and we did the same sort of thing.

The last piece of good practice guidance that I worked on was the social media guide in 2015-16. What struck me about it was that it was designed to help small companies—SMEs, start-ups—and the good practice examples were provided by the great big tech companies. That was a perfectly sensible idea, but after the document was published you could not help but notice that there was no sustained attempt to follow up to see who took it, who used it, what they did with it and, crucially, what happened: what the effect was on the performance of the companies that this was targeted at and to what extent it worked.

At that point, two-plus years ago, I thought our run of useful and valuable self-regulatory initiatives and good practice guidance had got something wrong and we were not doing enough. Particularly when you think if we're looking at harmful and offensive content, there is no independent regulator. For some areas, this is covered—the Information Commissioner's Office and so on. That's particularly true for that area.

Last autumn, along came the *Internet Safety Strategy* Green Paper. The Government had the idea of an annual transparency review. I thought, "This looks really interesting, because here's the missing bit. Here's the bit that says the big companies cannot go on marking their own homework; they must have an independent evaluation". That is when the penny dropped for me. It is only in certain areas. We are talking particularly about social media and child protection.

At about the same time, the idea of the internet commission was born. I think you have a submission from us, or at least an outline of what the

proposal is, but very briefly the thought is to take on the annual transparency review and provide a single common reporting framework. Different companies would provide information into this common reporting framework, so that if you wanted you could benchmark performance one against the other.

The intention as far as possible is to run a comparative analysis. The process would be independent of government, because we think there are some really tricky issues here, such as freedom of expression, which go beyond the safety side, where, yes, government will be an important stakeholder but it will be at one remove from the process.

Secondly, the intention is that where the work happens it should be in a neutral private space, so everything that the companies provide is not automatically published to the world. That is a space for dialogue, disclosure and evaluation, and ultimately an independent assessment and a first report. That first report would say what is working well, what is working badly, and, crucially—coming back to the start of your question; I am sorry it has taken a couple of minutes—what area from our evidence base may need more regulation. That might be statutory legislation, or it might be something else, but at least that would be evidence-based. There would be many months of rigorous debate and discussion, undertaken broadly in private, and out of that would come evidence-based recommendations.

Vicky Nash talked about procedural accountability when she came last week, and that is what we are aiming for and what we are working on. Yes, you look at quantitative figures, process and policy, and you ask, "How inside your companies, particularly in relation to social media companies, are you doing this stuff, because we don't know because you are not telling us. You are not being transparent and it is a bit of a black box".

I will give a couple of examples. In practice, how does the company decide what to remove and what not to remove? It is a very easy question, but the answers are going to be quite complex and subtle. It is then how the company measures and monitors the performance of human moderators and AI inside the company. They have their own metrics and they do quite a lot of this work; we just do not know what it is. In a rather discursive way, that is my response and an introduction to the internet commission.

Q29 Baroness Bertin: Building on that a little, in terms of the industry trying to do more in its own interests to avoid heavy regulation, can I ask a little about design of products? Baroness Kidron, who is not at this session, put down an amendment that you will be aware of on age-appropriate design. Can we have your views on whether that could have quite a big impact if done properly?

Julian Coles: Design is really important, and Baroness Kidron's amendment for the ICO to start looking at age-appropriate design for children is a very interesting idea.

Baroness Bertin: It is clear that children are a huge constituency.

Julian Coles: Yes, absolutely. One-third of humanity is under 18, as John Carr always reminds us.

Baroness Bertin: And the most vulnerable.

Julian Coles: Exactly. The design is shaped substantially by an advertising-funded model, to grab attention, to serve advertising and to collect data, but the design needs to benefit users and be for users and not put users at risk. Working on safety, privacy, transparency, consent and a fair and equitable exchange of value is an interesting idea.

Baroness Bertin: In reality, we are talking about dividing up the profit-making arm and the ethics. Do you think the industry is very receptive to that, or is it going to have to be dragged there?

Julian Coles: Things may have changed a little in the last couple of months when it comes to transparency.

Baroness Bertin: I wonder why.

Julian Coles: Quite. The threat of legislation always concentrates minds, and that is a difficult game of bluff. If you are doing an independent evaluation in the way that we have illustrated, if you have a stronger evidence base you can develop key indicators, which we have not done yet—we are in the middle of wrestling with those issues. Out of that could come, through procedural accountability, a more substantial test of willingness—the willingness to adapt, to change, to be much more open about how they work—so that third parties can measure independently how they are doing, because we do not know.

The Chairman: Rachel Coldicutt, do you have any insight into public opinion and what the public think of these issues?

Rachel Coldicutt: Absolutely. When we were asking people about their confidence in using the internet, it was extraordinary; people were marking themselves anything up to nine out of 10, but when we started to ask them how it worked, they said, “I don’t really know”. A lot of that is because things are easy to use and hard to understand, and that becomes more and more the case as we are looking at AI and voice. It is certainly not a matter of thinking only about children and adults. It is thinking about the cognitive load of, say, asking a person to understand everything that has happened after they have made an Amazon order or everything that happens in instant messaging. It is too much.

It has also come out that people have a deep feeling of ambivalence. There is a feeling that technology is helping them as individuals and it is not helping society as a whole. The conflict is: “I’m connected but everybody is walking around looking at their phones”, which again plays into the addiction.

On design, the question is how to create experience that is transparent but does not overload, and how people can understand what is happening without understanding everything that is happening, because it is enormous. It is a question of how design can be more responsible and focus on individuals' needs and society's needs as opposed to simply driving advertising. An element of regulation at the level of design could lead to better outcomes in good behaviour, too.

Dr Konstantinos Komaitis: I do not have a lot to add. Both my colleagues have used the word transparency, and we should not underestimate the power of transparency and how important it has become. Users, especially after recent events, want to learn more, and they want to engage more and to know exactly what is happening.

Currently, the question that everybody is asking is: would we have signed up to those services had we known exactly what and how our data was used? Initially, we all understood that advertising is part of the business model. Platforms are offering free services, but they are businesses and they need to make money. At the same time, we did not have a very clear understanding of the extent to which how our data was used—or abused, for that matter.

Rachel raised a very interesting point about how we can give information to users but make sure that we do not overwhelm them with that information. Where is the balance of information so that users can make informed choices about whether they want to sign up to those services? The design plays a critical role in the ability to have built-in tools that allow this information to be given to users but in a way that facilitates their interaction on the platform rather than trying to change the platform itself.

Baroness McIntosh of Hudnall: I am very interested in what you are saying, Rachel, about the extent to which people need to understand what they have in their hands when they use their smartphones or whatever. I was thinking about whether, when you say that, you mean that they need to understand the technology, but I rather think you do not mean that.

Rachel Coldicutt: No.

Q30 **Baroness McIntosh of Hudnall:** At the same time, you have all made the point in different ways that people do not understand how the technology works. They do not understand it enough to be able to judge what is going on when they interact with it.

I would take the example of a car. It is a very naive example, but I find it helpful. You may not. Most of us probably know how to drive a car, and we are tested and licensed on our capacity to do so safely, but if I were asked, "How does this car work?" I would not be able to say. Increasingly, as the cars have become more sophisticated, I would be less able than I would have been 30 years ago to say how it works.

I do not know whether that analogy works at all for you, but in trying to

grasp how we regulate how people behave in relation to this technology, which we do when it comes to using sophisticated equipment of other kinds, where do you see our responsibilities as users, not just our rights?

Rachel Coldicutt: I disagree that most people have an interest in learning more. People have busy lives, things are hard and the ease of technology is one of its problems. The driving analogy is close to my heart, because I am currently learning to drive. It is quite interesting, because in doing that I have learned quite a lot about how an engine works. I suppose over the years it becomes sublimated, you do not think about it and it is there in the back of your head. It is not that people need to learn, it is no one's fault, but there is responsibility on companies that are making money out of people's data to be extremely clear about the consequences of providing the data they are asking people for.

Currently, it is very hard for people to say that they do not agree to the terms and conditions of Facebook, partly because they are hard to understand and they are long, and because their lives are wrapped up in it, and there is no choice. Understandability to us is more about mental models and about "more or less" understanding how something works.

There is a chart that shows how advertising tech works. I do not know if anybody here has seen it, but it is enormous and it has 300 squares on it and arrows pointing everywhere. It is not possible for anybody to understand that, but it is possible to understand that your data is being used and personalised and monetised. It is about the level of understanding, I think.

Baroness Quin: To follow on, this may seem like a silly example but it has been worrying me since it happened just a couple of days ago. This sounds rather sad to say, but I was playing a word game on my phone and at some point something like, "Are you happy to share your contact list?" came up. I pressed "yes" instead of "no", and then wondered what on earth the consequences of that were. I was not able to find out and I was not able to go back.

The Chairman: Were all of us on it?

Baroness Quin: I do not think you were. Do not worry. It was rather an old contact list, in fact. Sometimes when you do something like that it will say, "Are you sure?", or, "This is what this means", but it did not say that at all. Is it as basic in some ways as ensuring that people are given more checks so that they know exactly what they are letting themselves in for?

Rachel Coldicutt: Yes, indeed, I would agree.

The Chairman: If that is the solution to Baroness Quin's problem, how do you codify that? If you say that is a solution, how do you get that imposed or adopted?

Rachel Coldicutt: I would say there are probably three elements. There need to be standards. I do not know if anyone has spoken to the Committee about dark patterns. Dark patterns are exactly as you

describe. They are design patterns that encourage you to do things that you would not ordinarily choose to do. It would be relatively easy to have best practice standards to mitigate against those. There is then auditing and refining.

Looking at the research, just over half of people admit to agreeing to terms and conditions without even bothering to look at them because they are overwhelming. I would advocate looking at a number of key journeys and mandating patterns for design.

Q31 Baroness McIntosh of Hudnall: This is moving on slightly and is about the vexed question of the content—sometimes worrying content—hosted by platforms. The question of their liability for the content they host is very important, but they are very highly protected by the laws that were set up at the time the internet began to take hold.

To what extent do you think platforms should be legally liable, given the extent of their reach now, for the content that they host? Secondly, do you think that we need to have a new way of defining them? The distinction between publisher and platform is very crude in the current era. Is there some in-between state that we can identify and nominate that will help us?

The Chairman: We will adjourn the meeting for 10 minutes.

The Committee suspended for a Division in the House.

The Chairman: Baroness McIntosh had asked the question about the liability of platforms and I think Mr Coles was about to start the answers.

Julian Coles: A couple of months ago, I attended a whole day conference at Cambridge on intermediary liability, with the best academics and lawyers. I did not come out with a clear picture of exactly how to take things further. There was some very intelligent discussion, but one of the puzzles is that under the e-commerce directive, for example, we are told that there is no general obligation for platforms to go looking for illegal content. That is explicit and very clear.

Last September, the Commission came out with a communication that said, "Wait a minute. We are now very keen that you should promptly and speedily do precisely that: you should go looking for illegal content on your platforms, you will not lose your immunity and, if you don't do it, we will legislate". I said, "Am I confused?" and they said, "Yes, you are, but you are right to be confused, because no one knows what this means".

You mentioned platforms and publishers, and I suppose people have already come to the conclusion that social media, for example, is a sort of hybrid. I have heard people talk about whether it would be possible to split the host or the intermediary definition into an active host and a passive host. The passive host would be something like a cyberlocker where nothing really happens. There is not the busyness and the moderation and the AI buzzing along that you might have on somewhere

like Facebook. If one were to try that, one would have to be very careful that it worked for the giants and the SMEs and the start-ups.

How do you get people to become unicorns and beyond from very little? It should not penalise platforms that actively manage their content. That is a tricky one. You could perhaps reward those with some sort of diminution in liability when things go wrong. I do not know. You also have to make sure that you do not give perverse rewards for taking down content when the content should not have been taken down. The freedom of expression side of things—the false positives—has to be weighed in the balance as you try to come up with a cleverer, more sophisticated and more focused definition.

Dr Konstantinos Komaitis: I would like to go back to when these exemptions on liability were introduced and why. Back in the day, it was—and still is—very important to ensure that every business that enters the internet and connects itself to the network is offered the same opportunities. That was essentially what the exemptions on intermediary liability were meant to do. It was a pivotal rule for a long time, and even today it has ensured innovation and competition. Imagine an environment where any new entrant could be liable for everything they were hosting. They would not be able to make it into this very highly competitive environment. This is particularly so for small and medium-sized enterprises in that they need to be able to enter the market and to grow. Thus, to an extent, we can argue that intermediary liability provides a level playing field for all entrants.

Of course, once you enter a market and begin to grow, this is when the questions start to become more interesting. Your behaviour begins to change as you move from being a platform that hosts content to a platform that does so much more with content. We need to think about the threshold where suddenly you are no longer just a platform but you are becoming something completely different and bigger, and it is not just hosting content; it is curating content, publishing content or managing content in general.

I know that I am not answering your question categorically, but it would be very dangerous to create a blanket rule where intermediary liability was scrapped off for every company, or every platform was treated the same when it comes to liability issues, because we would see fewer and fewer businesses emerging, which would mean less and less innovation, and ultimately it could affect competition.

Rachel Coldicutt: I would add that content is a bit of a distraction. The content that is published online by people and shown on social media is a symptom of the original problem, which is the underlying business model, and how the services are designed to encourage people to behave. While there is the issue of coping with the fire hose of content that is created and shared every day, we should be looking at the business models underneath that are encouraging people to create content that is pernicious, bullying and those other things. If there is an issue of liability

there, I would say that it is about changing the models and the ways in which people are encouraged to share content and return to it over time.

Q32 Baroness Benjamin: One of the Government's Digital Charter's key guiding principles is that people should understand the rules that apply to them when they are online. We know that young people, for instance, exchange material that is unsuitable and they do not quite understand that that is an offence. They are not quite conscious of their actions.

The charter commits the Government to protecting people from harmful content and behaviour and working with industry to encourage the development of technological solutions. What role should users play in establishing and maintaining online community standards for content and behaviour? What initiatives should be undertaken to ensure that all users are digitally literate and aware of their online rights and, more importantly, their responsibilities?

Dr Konstantinos Komaitis: In determining the role that users are playing in this ecosystem, we need first to make a distinction. It really depends on the type of platform. That is the starting point. There are some closed platforms where users are much more active and they are able to set those rules and self-regulate and, ultimately, to self-govern and create those best practices that can move on and evolve. Wikipedia is a very clear example. Users are self-governing and there are very specific rules about entries and mistaken entries and about when a Wikipedia participant gets excommunicated from the platform depending on how many times they have failed to adhere to the rules.

There are also more open communities such as Facebook where the user does not necessarily have the ability to self-govern. There we see the increased responsibility of platforms kicking in in making sure that users are engaged, bringing them in as much as possible and making them active participants. We saw this recently with the flagging of illegal content. We see users being more and more active in flagging up illegal content and asking for it to be removed and so on.

There needs to be greater effort to involve users, but before we do that there are a couple of things that need to be done. First, we need to re-instil trust. The idea of trust is more important when it comes to the internet. Recent events, even events before that, have made us question the trust that we place in the internet, in the platforms operating and in everybody involved. This comes back to transparency. We will come back to that word a lot, because it is so very important and it is becoming more and more important.

The second point is how you empower users so that they want to participate and be part of that. We need to look at accountability and how we can create accountable processes that will carry users with them, which at the same time users know will be at their disposal to be able to address those issues.

Baroness Benjamin: Where does WhatsApp sit in your thinking? The other platforms that you have mentioned are quite open, but WhatsApp

seems to be a very closed area with content that is not really suitable for people to be creating.

Dr Konstantinos Komaitis: WhatsApp in particular is an instant messaging communication medium in which you are creating content that you are exchanging between a group or one on one. It is not similar to Facebook, for example, which is an open platform. I think you are referring to the secretive communication that is happening through WhatsApp, if I understand you correctly.

Baroness Benjamin: My feeling is that you are driving people away to another area that is almost underground. Should we be focusing on that? A lot of people use WhatsApp, and racism, sexism and all kinds of issues have been happening on that platform. Have you thought through how we can embrace platforms such as WhatsApp as we move forward?

Dr Konstantinos Komaitis: I am not sure I fully understand your question, because WhatsApp is an instant communications medium between people, and the conversations are secret because they are encrypted. That end-to-end encryption is important to safeguard a lot of people who want to communicate secretly—and, yes, you are absolutely right that, as on any other platform, both online and offline, not very nice things sometimes happen within those apps.

This goes to your second question about how we get to a place where users are digitally literate and can flag this up and address those concerns. The other panellists might be better able to answer that, but there needs to be a lot of education.

The Chairman: Rachel Coldicutt, could you look at this burden of additional responsibility from the point of view of the user?

Rachel Coldicutt: First, the problem in communities that moderate themselves is that privilege tends to be afforded to people who are already privileged—people who have apparent authority in their offline life. Wikipedia has a problem in that it is not very diverse. It does not represent a large number of women or people of colour.

When it comes to platforms such as WhatsApp, we have heard that people do not know who to turn to. At the moment, if something goes wrong in their online life, people do not know who to go to. Is it a police matter, do they go to the platform, or is it even a legal problem? There is an issue there about transparency of escalation and a duty of care to the people. Equally, though, there is the problem of people not wanting to admit that things have happened, so there is a vulnerability there, too.

On the question of research into people's appetite to learn more, we have found, as I said, that people are pretty overwhelmed and that there is certainly a case for public health, as we are terming it. Rather than thinking about digital literacy and turning over the responsibility to every individual to look after themselves, it is more about a level of social awareness of how technology works and how to behave.

Public Health England's mission is interesting. It says, "We exist to protect and improve the nation's health and well-being, and reduce health inequalities". A digital alternative to that, which looks after people's health and well-being and gives them a place to turn to, seems like an attractive model. It should not be all about putting the onus on every individual to understand everything and to be doing the job of the platforms; they should understand who to turn to.

Lastly, in the questions we asked the public about regulating the internet, there was a huge appetite for a consumer body. Over 60% of people thought that there ought to be one. There was a lot of faith in the idea of, say, Which? being a body they could turn to. The Government were felt to have a role but perhaps not the capability.

Baroness Benjamin: A couple of weeks ago, Jeremy Hunt came up with a question about organisations such as Facebook taking responsibility. Do you think the Government should be taking that up or that, really and truly, it should be the various platforms taking up that responsibility?

Rachel Coldicutt: I would say that the Government need to create a culture of responsibility. The change that would have to be made in each of the platforms to move to that is enormous, so there need to be big regulatory incentives.

Julian Coles: Clearly, when you are talking about users and how they can make a contribution and take some control, there are user reports, trusted flaggers and feedback through academic research. End users have a vital role when it comes to the transparency reporting process. You need to look not at what happened in the creation of the report but at the outcomes, and how that creates sensible discussion and debate about the lessons we learn from what we have discovered as part of that process.

In terms of digital literacy, if you are looking at children, there are some excellent programmes—the Childnet Digital Leaders Programme, Parent Zone, the NSPCC—about creativity and critical thinking. It is not just to do with online, with the internet or with social media; a young person growing up in this world has to be equipped with critical thinking. That is fundamental and it runs right the way across the piece.

That does not let industry off the hook and is a really important point. We come up with great media literacy programmes, and companies do their bit and they provide money, but the companies need to take on more responsibility than they do at the moment. Ofcom, in its latest media literacy report for adults, reports that something like 19% of adults in the UK believe that if there is a search return on Google, and it is produced and revealed and comes up on your browser, Google have, to some extent, checked the accuracy of that piece of content. Last year it was 21% so it has gone down a little, but that is a really worrying statistic when one is talking about news, current affairs, elections and all that stuff.

It would be very easy to make it possible for ordinary users to click and get a sense not of the algorithms or the black box but of the key principles of a search and how the ranking is created. It is exactly the same for Facebook and social media news feeds; more and more young people are getting their news from social media news feeds, and some simple explanation of the principles behind the working of the algorithm would be really useful. That is equipping people with some essential tools when weird things happen, conspiracy theories rise to the top on a search engine and so on. I think that might help.

Baroness Benjamin: Who should do that?

Julian Coles: The companies that are running these critical services that we all use every day without thinking too much about should be under an obligation to make it really easy for the ordinary users to find out with one click how these things work.

The Chairman: Is it a moral or statutory obligation?

Julian Coles: I would hesitate to say that it should be a statutory obligation. This is one of the many things you could feed into the transparency review. It is about transparency and about informing the public. They can then decide whether they want to take the risk and what the implications of that might be.

Rachel Coldicutt: The likelihood of anybody looking at that is incredibly low. It would be cosmetic and would cover the company, without giving any further information to anyone who needed it.

The Chairman: Because people just would not engage with it. It is back to the issue of engagement with these tools and opportunities when they are out there.

Rachel Coldicutt: Because everybody expects speed and ease.

The Chairman: At the moment, when you are using the internet or you are on a platform, you are just concentrating on using it and you are not thinking about these things. When you are not using it is when you engage with the wider issues.

Rachel Coldicutt: Yes.

The Chairman: Thank you.

Julian Coles: I would add one thought. A lot of people I talk to—looking at academics, who are really across this, and people doing PhDs on cyberbullying and so on—have not found it easy to access the information on the key principles of how the algorithms work. The majority of the population will not necessarily pick that up, but if you are not making it feasible for academics and policy experts to get their hands on this stuff and have a sensible debate about it, perhaps there is some way to go.

The Chairman: That is a very interesting issue that we have discussed

here. We move on now to Baroness Quin.

Q33 Baroness Quin: I am not quite sure what level of regulation is best and whether it should be national, European or even some kind of international regulation, if such a thing were possible.

My question is specifically about the European Commission code of conduct on countering illegal hate speech online. What thoughts do you have about this? Is it fair, effective or transparent, or none of those?

Dr Konstantinos Komaitis: I did some research on that before coming here, because I wanted to get some numbers. My research at least shows that things are getting better, perhaps not as fast as many would like, but things are getting better by the year.

There have been three reviews: one in 2016, one in 2017 and one at the beginning of the year. The results of the third evaluation show important progress. Some 70% of notified illegal hate speech is removed by IT platforms compared with 59% in the second evaluation and with 28% only two years ago. The agreed timeframe for reviewing notifications—24 hours—is respected in the majority of cases. The research shows 81.7%, which is twice as much as in 2016 when it was down at 40%.

The other interesting thing that seems to be coming out is that reporting systems, transparency of reviewers, and co-operation with civil society organisations have been ameliorated. Concerning the transparency towards users, a positive trend has also been identified in respect of the fact that, in 68.9% of cases, feedback is given to the notifying users. I think we are getting there.

Co-operation and collaboration help. This is a very clear example of where you see collaboration between users and platforms. Also, you see the platforms coming back and addressing some of those challenges. There is always room for improvement, but I think we are getting to a stage where both users and platforms understand the critical issues and how critical it is to find solutions to them. However, it is certainly taking time.

Baroness Quin: Do Julian or Rachel want to add anything based on their awareness of this?

Rachel Coldicutt: I would only add to your point about regulation at different levels that there is an issue of different cultural expectations in Europe and the US. One of the defining differences is the approach to freedom of speech. That means that a common framework between here and the States is harder to achieve than one between here and Europe.

Baroness Quin: That is interesting. Do you have any thoughts, Julian?

Julian Coles: Just one aside, which is that I have heard Facebook say twice in the last few months that it is struggling with a definition of hate speech. It was asking other people, "Please come and help us. We don't

think we have the right answer. It's a really difficult problem. We could do with a hand", which I thought was an interesting example.

Baroness Quin: Is that part of the issue I was also going to raise about trying to get the balance between safety on the one hand and protection of the rights of freedom of expression and freedom of information on the other?

Julian Coles: That is really important. We know there is great pressure on the platforms to remove more harmful and illegal content and remove it more quickly, but at the same time how are we going to check whether content should be removed or should have been removed? Is there the right to complain? Is there the right of appeal?

When you ask the final question, which is how you check the impact of AI on the removal of content, we know from YouTube and the recent transparency review, that last June 45%¹ of violent extremist content was spotted automatically by AI. AI was the spotter. In December, it was up to 98% spotted automatically by AI. This is YouTube's own figure.

That sounds very clever, but the essential question I would ask is whether a human being then checks that this is correct or not. When does a human being check and when do they not? That suggests to me that there could be real value in some kind of independent batch testing/observation. We do not want them to be policeman, judge and jury. Obviously, they are not going to get it 100% right, but we could match the risks and the difficulties with what we know about AI, which is imperfect and not as great at context as it might be. YouTube, very transparently, has just said, "With AI, we are going to move beyond violent extremism to start looking at child safety and other areas as well".

Our commission has started to look very hard at procedural accountability and to think about how we can dial into this, not in trying to force the companies to spew their secret algorithms to the world but in trying to find out, "How are you monitoring this stuff? Why should we trust you?"

The Chairman: Rachel Coldicutt, do you have any evidence? You talk to users about the balance of their fears between the harms and freedom of expression. Do users talk and worry about freedom of expression?

Rachel Coldicutt: Not really. People are more concerned about their level of connection to their friends and family. We have found that people feel very uncomfortable about bumping up against things and that there is a feeling of helplessness. You have a neighbour who you have known for 20 years and it is only when you are friends on Facebook that you are

¹ The witness later wrote to the Committee: "I misquoted YouTube's figure for June 2017. It is 40% not 45%". See Google, 'YouTube Community Guidelines enforcement': <https://transparencyreport.google.com/youtube-policy/overview> [accessed 23 July 2018]

aware of their interests, which you perhaps are not all that interested in. People feel they cannot totally control the things that they see.

Slightly more concerning to me is algorithmic moderating, because an algorithm is only as good as the terms it is set. Enormous bias can be encoded within that. If Facebook or others are determining their own meanings for hate crime, the problem is that there is no transparency in the algorithm. Those are the kinds of things that ought to be taken out of the business and agreed separately.

Lord Gordon of Strathblane: I am sure the answer to that is a human court of appeal, as it were, which must act fairly immediately. The argument is whether you exclude on an algorithmic basis and attempt to win on human appeal or whether you allow it through and use the human appeal to get it back.

Rachel Coldicutt: There is a school of thought that there ought to always be a human in the loop in algorithms and every algorithm ought to have a named individual who is ultimately responsible for the parameters.

The Chairman: That is interesting.

Baroness Bertin: Because of the volume of traffic going through, in reality the amount of resource that will be needed to be redirected to police this, for want of a better word, is pretty big. I would like to know your views on that.

Julian Coles: I have been wondering and worrying about this. The algorithm is a response to the scale and speed problem. Very rapid progress has been made over the last six months in this one area, from what YouTube has told us. People are starting to push the boundaries all across. We have to keep a very careful eye on it. That is why I thought something such as batch testing could be used. You cannot appeal every decision, but batch testing is done independently, so you get in there and you say, "Okay, we will look at a hundred decisions or a thousand decisions out of a million and unpack them". They may be doing this themselves. I do not know.

Baroness Bertin: I have one very quick point. I know that this is not where we are at yet, but do you think there could ever be a product where the algorithm stops the image going up in the first place?

Julian Coles: That is already happening with hash testing.

Baroness Bertin: Presumably not enough though.

Julian Coles: It is most used where there is a known child abuse image. It goes into the hash bank and there is a lot of collaboration across different companies to say, "Right, we have this bank, we are sharing the bank", and it does not go up a second or third time. In a way, that is easier than origination.

Baroness Bertin: Yes, of extremist views.

The Chairman: To go back to Rachel Coldicutt's suggestion, you were suggesting, I think, that there should not be constant human intervention in what an algorithm is doing but that every algorithm should have a master, a human person responsible for what it does and accountable for that individual piece of programming.

Rachel Coldicutt: Yes.

The Chairman: Thank you. That is an interesting suggestion.

Q34 **Baroness Bonham-Carter of Yarnbury:** Picking up on that, we keep being told that AI is going to mean there are no jobs for humans, so I think we have created one here. Being serious, you were talking about resources, but this may well be a solution that is a good balance.

I am going to ask about something slightly different, which goes back to Baroness Quin's point when I first walked in—and apologies for being late—about inadvertently sharing her contact list. What information should online platforms provide to users about the use of their personal data? To what degree does the GDPR provide sufficient protection for individuals?

I also have a slightly tangential question, which I have asked before but I would be very interested in what you three think, about the misuse of a person's reputation, which is slightly different, to falsely sell things online. There are a couple of different questions in there. I know, Rachel, that your organisation has shown that 83% of people surveyed are extraordinarily ignorant, if that is the right word, about how their information is collected and what is done with it.

Rachel Coldicutt: Yes. It is clear that over 90% of people say it is important to know their data is secure, but only 40% of people claim to understand how their data is used, so there is a huge act of faith there. People are interested in knowing, but more than half of them cannot find out how it is being used. Some 70% of the people understand that data about them and research is collected and they understand that the website they have looked at has collected it, so there is quite a high understanding of active use, but there is a really low understanding of passive data, such as location and other things they are asked about the interaction we have with others. There is a very low understanding of that and how to control it.

The potential problem with GDPR is that it is not very actionable to me as an individual. If terms and conditions are not written in an understandable way, the likelihood of reading them remains low, to be honest. If my data is available, how do they get it? Is it my data? Is it the data of others I have spoken to? How do I look after it? How do I share it? There are lots of practical problems that come out of it, and given that most of the people do not really understand data as a term—it tends to be thought of as the package on your phone as opposed to the information about you—the likelihood of people knowledgeably and

intentionally asking to have their data looked after securely and being able to fathom it is low.

Dr Konstantinos Komaitis: A main issue in the current state of affairs, especially when it comes to platforms, is its centralised nature, where the data from and about the user is literally transferred to the platform provider, which results in a loss of control by us. We do not know what is happening, and on top of that there is the sense that it is an all or nothing situation. If you want to be part of a service, you need to do certain things, and there is no room for negotiating or even making it bespoke to your own needs.

Users are starting to demand to understand where their data is going and how it is being used. I mentioned already that at the beginning of social networking we all gave our information and we knew to an extent that that was part of the bargain: it was going to be used by advertisers, but we were part of the pool. Recent events have demonstrated that there are so many dimensions to how our data is potentially being abused. It is crucial that information is presented in a simple and comprehensible manner, because currently that is not the case.

You have heard all of us saying that the terms of reference are these long lists of never-ending pages that you do not understand. I have a legal background and I have to admit that occasionally when I read them I have no idea what the hell they are talking about. I can imagine people who are not digitally literate or who do not have a law degree will be even more confused. That is the first point.

The other point is that the information needs to be accessible and easily obtainable. We saw, especially in the beginning—again, it is getting better—that information was not so readily available. We often found ourselves trying to navigate a very complex ecosystem to get to that information, and even if we found out where it was we were not getting it. I would be a little conservative in my attitude to GDPR. It is coming into force in the next couple of weeks, on 25 May. At a fundamental level, GDPR tries to give users back control.

The Chairman: We will adjourn the meeting and, sadly, that probably means that we will not have an opportunity to come back, because we have already well extended the time we promised to keep you. We have one further question and we ask the clerk to write to you and ask you to complete your answers on this question and respond to one further question. We will close the meeting with thanks to you for your evidence, which has been very useful.