



Select Committee on Democracy and Digital Technologies

Corrected oral evidence: Democracy and Digital Technologies

Tuesday 29 October 2019

11.40 am

Watch the meeting

Members present: Lord Puttnam (The Chair); Lord Black of Brentwood; Lord German; Lord Holmes of Richmond; Baroness Kidron; Lord Lipsey; Lord Lucas; Lord Mitchell; Baroness Morris of Yardley; Lord Scriven.

Evidence Session No. 4

Heard in Public

Questions 45 – 56

Witnesses

I: Professor Helen Margetts, Director, Oxford Internet Institute, 2011-2018; Dr Martin Moore, Director, Centre for the Study of Media, Communication and Power, King's College London; Professor Cristian Vaccari, Professor of Political Communication, Loughborough University.

Examination of witnesses

Professor Helen Margetts, Dr Martin Moore and Professor Cristian Vaccari.

Q45 **The Chair:** Good morning and welcome. I am really sorry we are 10 minutes delayed, but we just had a very informative session. Would you first briefly introduce yourselves?

Professor Helen Margetts: I am professor of society and the internet at the Oxford Internet Institute, University of Oxford. I used to be the director of that—the Oxford Internet Institute, not the university. I now direct the Public Policy Programme at the Alan Turing Institute for Data Science and AI.

Dr Martin Moore: I am a senior lecturer in political communication and education at King's College London and Director of the Centre for the Study of Media, Communication and Power.

Professor Cristian Vaccari: I am professor of political communication at Loughborough University and Co-Director of the Centre for Research in Communication and Culture.

The Chair: Thank you very much indeed. I am required to read this rather boring thing into the record; I hope you do not mind. As you will know, this session is open to the public. A webcast of the session goes out live and is subsequently accessible via the parliamentary website. A verbatim transcript will be taken of your evidence and put on the parliamentary website. You will have the opportunity to make minor corrections for the purpose of clarification or accuracy. Thank you very much indeed for being here with us.

Q46 **Lord Scriven:** Good morning. In your view, does digital democracy replicate the power structures of offline political activity or does it represent a radical change in where power lies?

Professor Helen Margetts: I am going to start by saying something positive; do not hate me. There is a sense in which it does not necessarily radically transform power structures but encroaches on them in what I believe is an important way. Social media allows new tiny acts of political participation that were not possible before. Politics used to be very lumpy; it was the preserve of an activist elite, which had the time, money or whatever to do something big, such as joining a political party or tramping the streets at election time.

Just the act of liking, following, sharing or viewing a political issue as portrayed on social media or in a news item draws people into politics. It makes the entry costs much lower. We should not lose sight of that. That is drawing young people into politics in a way that they were not before. Who thought we would see US schoolchildren campaigning for gun controls, or the mass climate change demonstrations? Who thought we would see the viewing figures for Parliament.tv go through the roof?

Although we are going to talk about negative things, we should not lose sight of the fact that social media allows anybody with a mobile phone to fight injustice or campaign for policy change. That is one positive thing.

Dr Martin Moore: I entirely agree that there is a tendency to be either extremely utopian or dystopian, and we are going through a dystopian phase at the moment. Part of the reason is that we are going through a period of radical transformation.

To cite a particular example that I focus on, around elections and campaigning, as Helen has said, the hurdles to participation have reduced significantly. The tools available are now so accessible, and in many cases free, that the gatekeepers, particularly the mainstream media and the main political parties, have essentially lost their monopoly and dominance of the process. We are seeing many more individuals, organisations and new parties emerging that are taking advantage of the opportunities presented by this.

Depending on where you are coming from, that could be a very good thing because it enormously increases participation, although it also enormously increases noise, or a very bad thing, because we are seeing the decline of the parties, serious problems with mainstream media and lots of arguments and polarisation online. Coming back to Helen's point, we are going through this period of radical change, and in that radical change we are seeing an awful lot of challenges to the status quo.

Professor Cristian Vaccari: Echoing what my colleagues said, I have done most of my research over the last few years on the relationship between social media and political participation. We argue that social media both deepens and broadens political participation. It deepens it because, as Helen said, it adds to the activities that people can do. If you are already engaged, you can do even more things than you were able to do in the past.

It also broadens participation. Social media is a two-way street. Sometimes other people bring you in by tagging you, by inviting you to vote for someone, by exposing you to information that you were not necessarily looking for. There is evidence that engaging with politics on social media has a greater effect, in terms of participation, on the people who were less likely to participate in the first place.

About 15 years ago, there was an argument that the internet was just going to make the rich richer, that the people who went on party websites were those who already cared about parties, so there was not going to be any big change. Social media has changed that, because everyone is on social media. Everyone, at least some of the time, will be exposed to political information. When political information and opportunities for engagement reach people who are less engaged, they will benefit more than the people who are already engaged. For those people who are less engaged, every little will help.

Lord Scriven: You have all talked about how people can come on and,

once you come on, it is an equal playing field. Are the platforms equal? Where do businesses and influencers, et cetera, come into this? Are there power structures that people need to be aware of, which this technology can manipulate and amplify? Do you have a fear about that? Are there examples of it that we need to be aware of?

Professor Helen Margetts: There is no such thing as unmoderated public space online. All the platforms have different designs, which now shape the way that people behave. Because so much of our politics takes place on digital platforms, social media platforms are now institutions of democracy, and we should treat them as such and put pressure on them to change, if that is happening.

That is a key thing to remember. Every single design decision on a social media platform will influence the way people behave. For example, you see far more social information on social media platforms about what other people are doing. We know from decades of social science research that social information influences us. We tend to like what other people like, for example. That means there is the possibility to adjust those things or to put pressure on social media companies, for example, maybe to show trending information or not show it, because that will influence how people behave.

You are absolutely right. There are big differences between platforms. Snapchat, for example, does not show much social information. If you are friends with Boris Johnson, you do not know how many other people are friends with Boris Johnson. That may make you feel kind of special—or not. Do you see what I mean? The platforms are very different in the kind of behaviour they encourage. They also have different rules for sanctioning it, as I am sure we will talk about.

Dr Martin Moore: I would make two slightly different points about power and platforms. First, platforms, most notably Facebook and its subsidiaries, have allowed people to navigate around some of the existing norms, conventions and rules, particularly for politics and democracy. As such, they have favoured those who are most willing to go around those rules. People who are willing to breach conventions, rules and norms have benefited most, in the last decade at least. That may change.

Secondly, we have seen that a new source of significant power has emerged, in the collection, analysis and reuse of personal data. That power has accrued both to the platforms themselves and to political parties, organisations and campaigners, which have seen the value of personal data and have learned how to use it well.

Professor Cristian Vaccari: The issue of power and platforms is a very important one, where the reality has completely changed in the last 10 years. There is a new book by Matt Hindman called *The Internet Trap*, where he explains that platforms have accumulated so many resources that they can keep their users glued to them, in a way that makes it nearly impossible for new entrants to compete at the same level.

We are in a situation where, here in Europe, we have either some American giants or some Chinese giants dominating the market, which are going to dominate the market in the future. What these platforms do is crucial for democracy, for the kind of information people will access and for the kind of information people can share. As my colleagues have explained already, their decisions might be different, but there are very few of these platforms that matter. That is a key difference from the previous media environment where, for better or worse, we had a variety of TV broadcasters or newspaper publishers. While they might not all get it right, this guaranteed that they would not all get it wrong.

One concern with the present age is that, because of the way capitalism works with data and the ability to keep people glued to the platforms, a substantial part of political discourse now goes through these intermediaries that are either American or Chinese.

Q47 Lord Holmes of Richmond: Are there good reasons to be concerned about online echo chambers? How does the design of the algorithms used by social media companies impact the content that users see? Have aspects of social media become, in effect, outrage factories?

Professor Helen Margetts: One thing to remember here is that, as human beings, we really like echo chambers. We seek out people like us in the street, in the office, in bars and in clubs. We should not forget that. The perfect echo chamber is just watching CNN, just watching Fox News or just reading the *Daily Express*.

It is very difficult to compare now with a time when there were not any social media, but there is quite a bit of recent research—my colleagues may cite different research—suggesting that the phenomenon of online echo chambers has been exaggerated and that people who use social media look at more news on more news sources. Many of the pathologies of social media are concentrated in older people or more conservative audiences in the US, for example. That does not mean we should not think about it and encourage platforms to cross-fertilise information and not force people down into filter bubbles, but the evidence is more hopeful than people used to think it was.

Dr Martin Moore: In a way, as Helen says, the evidence is mixed, not least because it is an extremely difficult thing to research. Probably a more helpful way of looking at it, for me at least, is to think about the fracturing of the public sphere rather than specific echo chambers.

We have moved from an area where, yes, it was not a single public sphere, but it was a relatively constrained public sphere, to one that is atomised. Necessarily in an atomised public sphere you have some groups that are more extreme, whose extremism is enhanced, particularly by some of the platforms. We have talked about the main ones, but we have not talked about the minor ones. You also have public spheres that are highly cosmopolitan and centrist. In a way, the fracturing of the public sphere is a more important aspect of this.

The later part of your question was around outrage factories. I do not know whether we will come on to this in subsequent questions, but it seems to me that you cannot dissociate that aspect from the economic model and the business model of the platforms. That is probably a longer discussion. I do not know whether that is for this question or perhaps a separate discussion.

Professor Cristian Vaccari: I second what my colleagues say about the exaggeration around echo chambers. In my own research, we conducted surveys in nine different countries around the world. In all these countries, we found that people are more like to say they encounter disagreeing viewpoints than viewpoints they agree with on social media, and there are lots more echo chambers in face-to-face conversation, as people generally tend to talk to people they already agree with.

Although there is partisan media around the world, it is the mass media that exposes people to greater levels of information they disagree with, because most journalistic outlets will provide information on all main parties and not just one. Social media stands in between the really homophilic face-to-face conversations, where we mostly talk to people we already agree with, and the broader and more inclusive public sphere that most mass media provides.

It is also important to think about the connection, especially with regards to the outrage industry part, between mass media and traditional media. The biggest outrage factory in the US is Fox News; it is not Facebook. Fox News is the number one news source shared on Facebook, so there is an interaction there. In research I have done with my colleagues Andrew Chadwick and Ben O'Loughlin, we found that there is an affinity between tabloids and the sharing of misinformation on social media. The people who share misinformation on social media in the UK share a lot of tabloid news stories. Obviously, tabloids are big outrage factories, but they have had a new lease of life thanks to social media, because their content is more likely to go viral than content that is less provoking and exciting.

Q48 Baroness Kidron: Helen touched on this in her opening remarks, but I want to ask how effective campaign groups are at using digital tools to increase their political power. I specifically ask it in this context: we have heard evidence that many pre-existing political campaigns co-ordinate activity to suggest it is a groundswell of grass roots, when it is actually from the centre. I would be interested to know what you think about campaign groups as opposed to political parties specifically.

Professor Helen Margetts: On campaigning—not party campaigning but more general campaigning—I talked about tiny acts, and sometimes these tiny acts scale up to something really dramatic. There is a process by which they do that. Every time you view a YouTube video, you put up the viewing numbers by one and you send a tiny signal of viability to everybody else: “This campaign is going to take off”.

You can try to manipulate those numbers. It is not within the control of large, centralised campaign groups, but there are all sorts of ways they

can try. We have shown that it also introduces some randomness into the process. We do not know why some petitions get 6 million signatures and some get 10. It is really difficult to demonstrate, even though we have near-perfect data on that, so there is a sense in which it is producing unpredictability.

Baroness Kidron: As a point of information, do you mean that some petitions on the same subject get 6 million and some get 10?

Professor Helen Margetts: Yes, I could show you several petitions that have hugely varying rates of success. We cannot tell you, with all the analytical and machine-learning tools we throw at it, why some succeed and some do not. I believe that is part of the reason we talk about politics being more volatile and turbulent than it was before. I do not know if that is reassuring.

Baroness Kidron: Not particularly. I am interested to know whether you think the missing bit is what is happening at the platform level. Do you not know, because you do not know what is happening in the engine room behind that experience?

Professor Helen Margetts: I am sure we are going to talk about how none of us has the kind of data we need to do that research. That is a major issue and challenge to any of the possibilities of regulation or pressure that we will talk about, and even to answering some of your questions. We have Twitter data, which is why there is a lot of research on Twitter, and Reddit data, but Facebook data is not available. WhatsApp is encrypted and it belongs to Facebook anyway. So does Instagram. Snapchat data is ephemeral. That is a major issue to researching that question; that is right. To answer one of your concerns, it is the case that some campaigns are pumping out advertisements in a very manipulative way, as you describe. That is very poorly regulated. Our system of regulation is not adapting to regulate that sort of campaigning outside election time. We have seen that happening, particularly in the Brexit debate.

Dr Martin Moore: On Helen's point about the difficulty of assessing this, there are lots of claims made by campaign groups and professional campaigners, who obviously have a self-interest in promoting their success, but also grass-roots campaigners, on how they have succeeded. One thing we have seen from campaigns that have claimed success, where those claims appear to have substance, is that those that are more successful have recognised the affordances of the different platforms. In other words, they are conscious of what different digital tools enable them to do and use them to do that. Central to that is collective action at speed.

One thing many people have misunderstood about social media is the degree to which it is incredibly effective for behavioural response, rather than as a persuasive tool. In other words, it encourages people to click, share, like and comment, and then potentially to go and do something in the real world. It is much more effective at doing that than at trying to

convince someone of a particular political opinion or bring them over to particular political issues. You are essentially trying to mobilise people who might be sympathetic, rather than convert people who otherwise would not be.

Professor Cristian Vaccari: We probably all agree that, as a result of the relatively open public spheres we have now, we have more campaigns going on around the country and around the world than we had in the past. Some are more savoury and others are less savoury. By all means, there is a degree of astroturfing. Twitter bots are a good example. Why would somebody create a Twitter bot and let it automatically retweet stuff, when they have three followers? Well, it is because that increases those retweet counts. It counts the same as the genuine activist who is expressing genuine support. I know from experience that it takes very little money to buy thousands of accounts. Up until recently, Twitter would have peacefully let them go and exercise automatic retweeting and astroturfing. It has cracked down a little on that now, but it has mixed incentives. If it cracks down on inauthentic accounts, the user base goes down, and its shareholders do not particularly like that.

There is one important change that we need to acknowledge, especially for the more grass-roots campaigns that have tried to mobilise organic support, even though they might not have a lot of resources. It has started in the last few years with Google and Facebook, and pretty much all the other platforms have followed suit. If you are a public-facing actor and not just an ordinary user, the way Facebook and Instagram work is pay for play. You do not get much organic reach unless you can mobilise a lot of people to share your content.

Five years ago, I followed Barack Obama on Facebook and my chances of seeing what Barack Obama posted, as a user, were 0.3 or something. Facebook introduced changes in the algorithm under the pretence that it wanted to give people more information from their friends and family. There are many pages run by news organisations, political organisations and civil society groups, rather than user accounts, so they are not your friends and family. That meant my chances of seeing content from the same page of Barack Obama in 2017 went down to 0.05 or something. I am making these numbers up, of course, precisely because we do not have the data. We know from research on drivers of traffic to news organisations, for example, that a lot of news organisations have missed a lot of traffic from Facebook as a result. The same applies to the not-for-profit sector.

The end result of this is that, if you want to reach people on Facebook, Instagram and, to a certain degree, Google's search engine, you need to pay. You need to buy ads. That means that, if you are not an organisation that can mobilise a lot of money, you are at a disadvantage compared to the previous regime, which was arguably more open.

Baroness Kidron: For the record, the next Barack Obama would have to pay a lot more to have the sort of campaign he had.

Professor Cristian Vaccari: Yes.

Q49 **Lord German:** You have answered this question in part, but I want to try to crystallise what I am asking here. It is about the impact on our democracy of each of these different social media channels. Could you tell us whether there are any large differences in the impact on our democracy between YouTube, Facebook, Instagram, WhatsApp and Twitter? Have we left anything off that list, which ought to be put in, because it has a major impact on our democracy?

Professor Helen Margetts: There are a couple missing from there. I mentioned Snapchat. I do not know what the future of Snapchat is, but it is used by a lot of young people.

Lord German: Is it used in a democracy context? That is what we are asking.

Professor Helen Margetts: Quite a bit of the success relative to expectations of Jeremy Corbyn in the 2017 election, you might argue, comes from Snapchat. It is particularly difficult to research, because the data disappears as soon as it is read, or at least that is what they say, so they cannot possibly give it to you. There are circulating YouTube videos and things like that.

Lord German: Are there any differences between the others?

Professor Helen Margetts: They tend to have different users. Twitter has established itself as a forum for politicians, journalists, academics and researchers. There are a lot of the public there, but I see it as much more of a public space and less of a socialising space. Instagram's users tend to be younger than Facebook's, for example. Very few teenagers are active on Facebook, apart from having a fake account, so their parents think they know where they are. There are big differences. LinkedIn is a professional platform. We are doing quite a bit of research on hate speech at the moment, and some platforms are used more for hate than others. Twitter and Reddit particularly are popular for hate, perhaps more so than the others, because of these dynamics and design features, such as the extent to which you are anonymous, the ease of setting up an account and the extent to which you can be public to lots of people. Reddit, for example, is completely unmoderated. They do not take stuff down. If you are into a really hateful discussion, it is a good place to be.

Dr Martin Moore: Yes, there are substantial differences in their political effect and the way in which they work. The platform of choice for political campaigns for the last six or seven years has been Facebook and its subsidiaries. That is because Facebook has a remarkable amount of personal information about people and allows you to target people based on that personal information and those demographics, with tools that were created for the commercial market, but are incredibly useful for politics, such as custom audiences and lookalike audiences, which I can talk more about. They allow you to do things politically in campaigns that you were never able to do previously. For example, you can do A/B

testing at remarkable scale, putting out 50,000 to 100,000 digital ads a day, which is what they were doing in the 2016 Trump campaign.

By comparison, it seems that Twitter has much more of an influence on the process of politics. You can do targeted advertising on Twitter, but it reaches a much smaller audience than Facebook and it is less representative. Twitter has significantly increased the speed of politics and the journalism around it, and affected the nature of deliberation within politics.

We are going to see an awful lot more influence from YouTube than we have in the past, not least because its growth has been phenomenal, particularly among certain younger groups who watch substantial amounts. We are not talking of 10 or 15 minutes a day; we are talking hours of YouTube a day. On the effects of YouTube, I talked about behavioural influence before. YouTube is more persuasive. In the Brazilian election, for example, there was an awful lot of use of WhatsApp. The highest proportion of URLs shared for a website was for YouTube, because people were sharing political videos and vlogs — video blogs.

The final one I want to mention is Google, which we sometimes forget about, partly because we think it is not a social media company. If you take Google in its entirety, Google AdSense has had a profound influence on politics, because it is the way in which many sites fund themselves. Look at the original eruption of concern about fake news, about the teenagers in Veles, Macedonia: much of their funding was coming from Google AdSense ads. One reason that they were inventing the news was to earn income from Google AdSense. Each of them plays a considerably different part, but they all have their effects on politics.

Professor Cristian Vaccari: When we think about platforms, we definitely need to differentiate between them. We can differentiate them according to two main criteria. One is the audience and kinds of people who use them, the other is the actions they allow and encourage users to perform. Everyone is on Facebook. Young people are there, but they do not use it, as was already said, whereas other platforms have a more selective audience. YouTube is near-universal, because from time to time almost everyone will watch something. WhatsApp's distribution and popularity varies by country, largely as a result of how expensive it is to send text messages, but it is very helpful for interpersonal relationships.

We have done research that shows that, especially for the platforms that are private social media, such as WhatsApp, Snapchat and Facebook Messenger, even though you can have broader conversations in groups and public groups, and Boris Johnson can be your friend, most conversations that happen are between people who know each other and have close ties. On Facebook, whenever you post something, you do not really know who is going to see it. You might have 800 friends, but they are not really your friends; they are just people you had conversations with at some point, or no conversation at all, but you just got connected.

We have found evidence that these private platforms can be helpful for people to express themselves politically, who would normally be shier in expressing themselves on more public platforms. If you are the kind of person who does not like to talk about politics in public, private platforms such as WhatsApp are great for you, because that is a space in which you feel more protected. That can mean good things in terms of self-expression, opening up opportunities for people to exchange information and have a dialogue. It can mean the below-the-radar spread of misinformation and the running of stealth persuasion campaigns.

Finally, one key distinction that cuts across the platforms is whether they are used on mobiles or from a computer. In the last few years, the big change has been that most of us use these platforms, at least some of the time, from our phones. Research is still developing in this area, but it has been found that, for example, when watching a video of news on a mobile phone, a lot of people do not understand it as well as when it is watched from a computer. That is because of the physical limitations of the screen and because you might be watching it in a context where you are distracted, in a queue at a supermarket or while somebody else is telling you something. That means the level of cognitive engagement you have with that content is not as good as it should be to properly understand it.

Lord German: Will these differences accelerate over the coming years, as platforms try to create different niche markets, or are they likely to be sitting on top of each other, competing in the same space?

Professor Cristian Vaccari: In Silicon Valley there is a culture of going with what is popular. A few years ago, Facebook decided it wanted to colonise some of the space that Twitter had conquered, and introduced features such as share, hashtags and trending topics that Twitter had pioneered. When Facebook bought Instagram, the original founders were reportedly dismayed that a lot of the features of Facebook were rolled over, in more or less the same form, to Instagram. There is an emulation culture on these platforms, because they can very quickly test whether something works. If it works, they will very quickly adopt it. If the same things work on different platforms, they will be emulated. At the same time, not everybody can spend all their time on all the platforms. Users themselves will select the platforms that cater more for their needs. There will probably not be one platform that captures or corners the whole market.

Baroness Kidron: I was interested in your list, Martin. I wondered whether you thought that the platforms that push recommends have more power, in terms of how they engage with users. I was interested in what you said about behaviour and then persuasion. With YouTube, we know the recommend pattern means that, once you go down a tunnel, it is a very involved tunnel. Does the persuasion/behaviour bit have to do with how the platform recommends or is it because it is video, which is more sensory?

Dr Martin Moore: I think it is mainly because it is video and it is more sensory, but the recommendation engine on YouTube is phenomenally important and has become more so in the last few years. There used to be a degree of serendipity on YouTube. You used to have options on your homepage and more ways to explore. Now, essentially you subscribe to channels, search for something or are recommended something. If you subscribe to things, you personalise them. The more you subscribe to or watch, the more personalised your page becomes. It therefore becomes a very individualised feed, partly because the quantity of video being put up on YouTube, which is insane, is such that, a bit like TikTok, it has to base it entirely on past precedent, popularity and search relevance.

Q50 The Chair: To ask a very relevant question, since we have been sitting here, we now know there is going to be a general election. There could be a lot of marginal seats contested and a lot of microtargeting. Which of the platforms do you think we should be looking at and what sort of activities should we be alert to, in that scenario? Would you like to write?

Dr Martin Moore: Helen has already mentioned the problems. We have not seen any substantive reform of electoral law since 2015-16. Yet, at the same time, we have seen a significant increase in the sophistication of certain technologies. Obviously, everyone now knows about some of the techniques that were used. While we have a little more transparency, thanks to things such as the political ad archives that platforms now have, people can still do most of the things they could do previously. They involve bypassing lots of the rules that we currently have, the most obvious one being around spending. People can focus huge attention on specific constituencies and spend quite significant amounts on them, far beyond the £10,000 to £15,000 threshold that people are supposed to spend. That would be one of the most major things to happen and it would be incredibly difficult to control for or even audit subsequently. In terms of a level playing field, we will see an awful lot of groups outspending others in this election.

Professor Helen Margetts: On the regulation point that I mentioned before, our party system has broken down. Parties have become much more ephemeral organisations, more likely to break up, splinter off and so on. We are seeing parties turn into campaigns and campaigns turn into parties. Take the unofficial leave campaign. They have been pumping out ads on Facebook, unscrutinised and unregulated, for three years now. They have segued effortlessly into being the Brexit Party. Presumably, although I do not have any insider information, they will have the same databases as the unofficial leave campaign and the same likes. None of this happens in isolation. Advertising is putting things in front of people; then some people like it and they will always see that, regardless of whether it buys ads. That is an important point. It will happen on campaigns of all persuasions, but this transition from campaign to political party will be important.

Q51 Baroness Morris of Yardley: Helen mentioned hate speech earlier, in answer to another question. What is the evidence on this? How much do you understand about the level of hate speech and how it affects political

debate? It is another of those things, like the echo chamber, where you made me think that the evidence might not be as strong as I thought. What do you think?

Professor Helen Margetts: We submitted a consultation response from the Turing hate speech project, if you want more details, but I will summarise. I am afraid we know astonishingly little about hate speech. Part of the reason for that is that we do not try to measure or monitor it. This change needs to happen. We have two conflicting sources of evidence about this. We have some official statistics from the Home Office on hate speech that was so bad it was prosecuted—actual criminal cases of hate speech, of which there are about 1,700 a year. Those are all the official statistics that we have on hate speech, and they do not include hate speech against women, for example, because that is not a crime.

At the other end is the hate speech that people actually perceive or feel that they receive or see. Between 30 per cent and 40 per cent of people report seeing hate speech online, and between 10 per cent and 20 per cent report receiving or being a target of it, so there is this incredible difference in what we know. It is very difficult to measure, as well. As we have said, the platforms do not release data. You can scrape data, but you have to do it quickly, before it is taken down by the platform. If you take a chunk of social media data and analyse it, you will not get enough hate. I am sorry; once you start researching hate speech, you start talking about it like it is a good thing. Obviously, I know it is not. It is a very difficult thing to research and we need to get much better at it and at monitoring the hateful environment.

There is no doubt that it is a problem. Some people believe it could deter a whole generation of women from public life, for example, not just because of the hate speech they receive, but what they see. If you see women MPs — and all public figures who are female, black, Muslim or whatever — being endlessly bombarded with hate, how does that make you feel, as somebody considering entering the public sphere? It is really important. We do need more research. We need more research on the simple measuring issue, more than for other things, because we have got better at measuring misinformation and a lot of the other bad things we have been talking about. I could talk for a long time, but the challenges are discussed in our response to the consultation.

It is difficult. When Mark Zuckerberg says it is difficult, he is not kidding. Facebook used to say that it could automatically catch between 38 per cent and 48 per cent of hate speech. Now it says 50 per cent. It can automatically catch around half, but that still leaves a half to be tackled by human moderators, so it is difficult to do. There are all sorts of issues such as context. When Arron Banks tweets, "Freak yachting accidents do happen in August", with a picture of Greta Thunberg on a solar-powered yacht, is that hate? I get completely different responses to that from different audiences, so context is difficult. Language is a big problem. So much of the research here is just on English and almost all is on text.

What about videos? There are a lot of challenges and we need to gear up to those.

Baroness Morris of Yardley: Are you trying to measure both the amount of it and the impact? I can see that both are very difficult, but are you thinking of trying to do both those things?

Professor Helen Margetts: For all the things we have been talking about, impact is really hard. With echo chambers, misinformation and particularly hate speech, the effect of seeing it is so difficult to measure without the data. You could do it with the data, in response to your point about being able to see behind the platforms. You would need a data-sharing agreement with one of the platforms that allowed you to look at somebody's behaviour before and after.

Baroness Morris of Yardley: You would need that level of information.

Professor Helen Margetts: I do not want to cause complete alarm and despondency. There are things we know and are getting better at. For example, we know it is very time-sensitive. When there is some kind of terrorist attack, for example, there will be a peak in hate speech. That is important because, if we can understand more about the time-critical aspects of hate speech, we might be able to put pressure on the platforms to be more careful about what they take down during those times, to have more blocking or to make it less visible. There are things platforms can do, but there is a way to go.

Dr Martin Moore: Briefly, because I know that Helen has done a lot of work on this, some good work has also been done by Kalina Bontcheva and her team at Sheffield University, who looked particularly at the abuse of public figures on Twitter. They have seen, in both absolute and proportionate terms, an increase since 2015 in the amount of abuse of public figures.

My research has been more focused on particular services, such as 4chan. Some research I did recently brought out two other aspects that make this research particularly difficult. One is the evolution of language. We found it was hard to do the quantitative research, because many of the terms used were deliberately obfuscating. When you read things in long form, it was quite clearly hate speech, but quantitative tools would not have picked it up.

The other, specifically for those kinds of platforms, is the degree of false positives. One thing we were looking for was around radicalisation and the degree to which these platforms were leading to extremism. You have an awful lot of calls for action, threats and things like that, most of which never come about, but some of which do. There is a particular law on the internet, called Poe's law, that it is virtually impossible to distinguish between satire and seriousness in a textual context. From looking simply at text without seeing the broader context, it is almost impossible to discern the two.

Professor Cristian Vaccari: Besides what my colleagues have said, one area in which I would like us to know more is not only the effects on the target of hate speech, but the effects on the bystanders to hate speech. If I see somebody abusing someone else on social media, does it make that behaviour more legitimate? Does it change my perception that this is not something you are supposed to do? If I see that behaviour happening and social media not taking that content down, does that create the perception that anything goes? Sometimes you see on social media people reporting that they were targeted with hateful abuse; they reported it to the platform and the platform said it did not violate their standards. Obviously there are issues with enforcement, but there are issues with the fact that this is visible to people. Does that make it part of our reality so, sooner or later, you might also do it and propagate the chain?

There is research by colleagues in America who have identified a set of attitudes that they call the need for chaos. These are small groups of people—they have done research in the US and Denmark—within the general population, but they have surprisingly destructive attitudes. They do not like the system. They want to tear it down. It is not distrust; one of the questions that they used to measure it is this: “I just want to watch the world burn”. Some people answered yes to that. The concern is that this group of people could mobilise hate speech and, by doing that, bring in others and slowly change the culture around it.

- Q52 **Lord Lucas:** Who should be allowed to define what is and is not hate speech? There seems to be a lot of calling-out of hate speech, when it is just the butting together of difficult arguments, but academia is getting infected and there are various areas where academics are calling for other academics to be sacked, just because they have taken a different view on a difficult argument. If we are going to allow hate speech as something on which we take action, somebody has to police that boundary. Who should that be?

Professor Helen Margetts: The research community has work to do here. There is a sense in which hate speech research has become a little like telemedicine in the 1990s: there are lots of tools being built to classify one sort of hate, at one time, targeted at one person, on one platform. Write the paper, chuck it over the wall and move on. There is a pile of papers. We need to do more work on synthesising, which we are starting to do at the Turing, and on standards. That is normative work on what hate is.

Remember, a lot of this work is being done by engineers, who are not used to thinking about things such as what hate is. It comes much more from the realm of the social sciences, arts and humanities. That is something we are starting to do, but it needs that multidisciplinary research with engineers and social scientists. That is what we are trying to do. It applies to the platforms, too, which have very imprecise definitions of what hate is and very little transparency of what they do. They have got better but, even so, there is a tendency for Twitter in particular just to take something down because it is causing problems,

rather than for any systemic reason. There is a lot of synthesising work to do there; you are absolutely right. There is a lot of imprecision about the definition, and that is a normative task that has to be part of the research.

Dr Martin Moore: To add to your concern, from looking at these particular communities such as 4chan online, there is little coherent ideology beyond the destructive aspects. But one of the few threads running through most of the conversations is what I call a fundamentalist approach to free speech. I call them free extremists. One of the main critiques they have of liberal society is the degree to which it is starting to censor and close down speech. That seems to be particularly mobilising and is one of the few consistent ideologies running through not just 4chan, but Gab.ai, previously 8chan and others.

Professor Cristian Vaccari: I share your concerns. We need to be careful before we say that things should be taken down systematically. One important distinction that people are beginning to make is between free speech and free reach. You can say certain things. You probably cannot say things that incite hatred and violence, but you can say things that are distasteful and violate certain norms on social media. The question is whether there are the same chances that those statements spread virally as for those that do not violate such norms. Platforms, particularly Twitter these days, are beginning to think, more in the context of misinformation than toxicity in speech, of how we maintain the ability for people to say things that are inconvenient, while at least not giving them an edge.

This goes back to the question about the outrage industry. Platforms are beginning to realise that they need to help societies tone things down sometimes, without censoring. That will be an important distinction, moving forward.

Professor Helen Margetts: it is important not to think about it in a binary way, as we have in the past, as hate or not hate. We have to get better at measuring the whole spectrum, going all the way from offensive to death threats, because it is a continuum. That is challenging, but we have to tackle that task and not see it as a binary thing.

Baroness Kidron: What role do you think community rules have in this conversation? You do not need an absolute definition of what is hate and what is not; you just say, "In my house, we do not do this, because I like my conversation this way".

Professor Helen Margetts: Facebook, for example, has extensive community standards. Until recently, they were published as a black and white pdf on the site, which could be found with great difficulty. Facebook has tried to make that more public now, but it is difficult, because people do not look at the community standards. Facebook in particular wants to think of itself as a community.

Baroness Kidron: I was thinking about it upholding them, rather than

people reading them.

Professor Helen Margetts: It tries, but lots of these challenges remain.

Q53 **Lord Lipsey:** I should declare an interest as a founder trustee of Full Fact, the fact-checking charity. I want raise a question of misinformation — or lies, as it might be better told — in conventional and new media. After all, there is plenty of misinformation in old media. Indeed, the present Prime Minister was a pioneer of making stuff up when he was a Brussels correspondent and it was bloody difficult to get it corrected, too. Newspapers, and I was once a journalist, went to infinite lengths to fend off anybody who sought to complain about it. I wonder if you can compare and contrast the tasks of combating misinformation in conventional and new media.

Professor Cristian Vaccari: This is one of the key questions of our time. We have done research at the Online Civic Culture Centre at Loughborough University, which we submitted as evidence for this inquiry. This shows that 57 per cent of the British public who use social media believe they have seen information online that is misleading or completely false; 42 per cent admit sharing it; and only three-quarters of that 42 per cent report that somebody corrected them. The glass could be half empty or half full; make what you want of that, but one-quarter of people who report sharing misinformation do not report that anyone has corrected them for it.

The question is absolutely right. A lot of responsibility for misinformation in our society lies with professional journalism, partisan politicians and partisan organisations, which bend the truth in different ways. Social media has introduced some additions to that problem. One is scale, another is speed and a third is potentially the stealth by which these activities occur.

We know from research that correcting misinformation is possible. A few years ago, there was a tendency to see a backfire effect, so somebody corrects your false beliefs and you react by sticking to your guns and confirming those beliefs, rather than acknowledging that the information on which they were based is false. More recent research suggests that people take on the new information that is provided to them, but it is difficult to make sure that information reaches them when it is needed. Another important problem to address as a society is that, while overall 42 per cent of people said they shared misinformation, as many as 17 per cent said they intentionally shared information that they knew to be false.

There are two problems here. One is that people are sometimes sloppy and just share stuff that seems funny, interesting or true, even though they have not done due diligence or checked whether it is true. But importantly, a substantial amount of people on social media are doing it intentionally, whether it is because it advances their political agendas, because they want to watch the world burn, as I mentioned earlier, or because they just want to see what people have to say. As a society, we need to live with some degree of misinformation and some grey areas

between truth and falsehood, but we need to make sure there are spaces where that truth has a chance to prevail.

The platforms are trying to do a lot of work and spending a lot of money on that, but, arguably, problems of scale and speed are difficult to crack. They have also taken very different positions in regard to what politicians say. Facebook says it will not fact-check ads by politicians, whereas Twitter has said something slightly different. It makes a more fine-grained distinction: "We will let them speak if it is not true but, if we find it is not true, we will make it very difficult for that to go viral". That again is probably a direction of travel that we want to take, but it has to come from a collaboration between traditional media and the journalistic profession, the platforms and the political class.

Professor Helen Margetts: It is true that a very centralised way of disseminating fake news is by journalists making things up. Actually, the sharing of misinformation on social media seems to be more concentrated than we used to think. Things do not go viral that often. Recent research published in *Science* suggests that 0.1 per cent of Facebook users share 80 per cent of misinformation, and there is other quite congruent research about Twitter. It is not so easy to spread misinformation as perhaps it seems.

Dr Martin Moore: The research that I found most helpful on this was published by MIT last year, looking at the spread of false information on Twitter. It found that false information is 70 per cent more likely to spread than true information. When they tried to correlate it with various different factors about why false information was more likely to spread, they went through a series of possible options and found that the reason that was most correlated was newness. I suppose that should not be surprising: people had not seen the fake news before, because it was invented. Therefore, it was fresh and they were more likely to share it. That was the finding of that research.

To add to Cristian's point, one thing we have to remember here is the quantity question, which is simply that there is almost certainly an awful lot more disinformation. There is also an awful lot more true information. There are many more people making cogent arguments, and many more people making hyperpartisan arguments. We sometimes forget, when we focus specifically on certain types of disinformation, that there is just a huge amount more information out there.

Q54 **Lord Lucas:** How effective is the UK's current regulatory system at ensuring that online activity does not harm democracy? Helen, you referred to that earlier. What do we need to do to make things better?

Professor Helen Margetts: We definitely need to do something. It is a challenge. All regulators are challenged by digital platforms based on huge quantities of data and machine-learning technologies. It is not just elections that we need to worry about, but it is true that the Electoral Commission needs changes in the law and to tool up here. There is a particular problem with political advertising, because everybody is saying,

"It is not me". You have the Advertising Standards Authority, Ofcom, the Electoral Commission and the Information Commissioner's Office, and they are all slightly saying, "That is not exactly us; it is more you". That needs to be confronted head on.

The Information Commissioner's Office is an interesting example of a regulator that has tried to take the lead here. It has set up an artificial intelligence regulators' working group to think about this and put together regulators that are confronting similar issues. We need to value our democracy in the same way we value everything, tool up our regulators and change the law where needed. Some regulators have been good at campaigning for a change in the law, and we need to do the same for elections and to do it very soon—like tomorrow.

Dr Martin Moore: I will point quickly to five specific ways in which the current digital information environment threatens our existing electoral rules and laws. The first is protection of the secrecy of the vote. The second is protecting voters from undue influence. The third is maintaining a level playing field. The fourth is preventing elections being bought. The fifth is enabling people to verify political information. I can go into much more detail on each of those, if it would be helpful.

Professor Cristian Vaccari: My colleagues have already said a lot of useful things. I want to make a slightly different point. Whatever regulation we are talking about, it will affect freedom of speech, one way or the other. We need to be careful in a democratic society, whenever we touch that, that we are really trying to protect goals that are worth protecting, even at the price of limiting some forms of speech. We do not have enough evidence in many areas, besides legitimate concerns, that the things we have been discussing today have a real effect. How big is the effect? How many people does the effect pertain to? To know that, we need research and evidence. As we have said multiple times, we do not have evidence and data yet, because the platforms are keeping it to themselves.

For the last year and a half, Helen and I have been part of an initiative called Social Science One. It was created mainly in the United States, but with committees in all the other continents in the world. It was supposed to facilitate the sharing of Facebook data, and potentially other platforms' data, with the academic community, in a privacy-compliant environment, where the companies would feel that nobody was going to infringe on their patents, damage privacy or damage users. It is not happening. After a year and a half, we have limited access to limited data, which does not begin to crack the problems we need to crack.

I was at a meeting in Brussels 10 days ago, when a lot of policymakers discussed this. It is increasingly being argued in the European Union that the law has to force companies to give over some of their data, in a safe haven or privacy-compliant environment, so we can figure out these answers before we regulate them. It is very clear, after the experience of Social Science One, that the only way in which this will happen is if digital platforms have to share some of their data with researchers in order to

operate legally in a country. You could call it a data tax so, just like companies pay taxes on their profits, companies that process data should share some of it, for us to answer questions about what it does to society. If it does not happen that way, it will not happen in a voluntary way on the part of the platforms to the extent that we need to answer these questions.

- Q55 **Lord Mitchell:** There are lots of democracies in the world, not as many as there should be, and they presumably all face the same issue. They are probably having similar discussions to us. I would be interested in your judgment on how we rate on this, in this country. Everybody is saying we should do a lot more, the ICO should have more funding and so on, but how do we rate on the league table?

Professor Helen Margetts: It is very difficult to say, because we have some unnatural advantages: we are a relatively strong state and we speak English. All the actions of the platforms, in defending us from hate speech and the other things we have been talking about, are geared at English-speaking strong states that will make a fuss if they get it wrong. In that sense, we rate well, but I could not say we rate well for electoral regulation.

Dr Martin Moore: There are different things. I think electoral regulation is seriously problematic. The ICO has been significantly strengthened, so we have a lot more knowledge and expertise when it comes to personal data, for example, and the use of personal data. We also have a very different media environment from the US, for example, as we have a strong public service broadcaster, which makes a significant difference when it comes to the nature of the environment, the types of information that are shared and the degree to which the public sphere is fractured. There are peculiarities to the UK. As Helen says, many of them are advantages, but there are certainly areas where we are behind.

Professor Cristian Vaccari: It is hard to have a comprehensive view of what is happening. Looking at the future, there will probably be some action on the part of the European Union in the next Commission. We are leaving the European Union, but it will be important to keep an eye, at least in those areas of legislation, on what the European Union does. Although we are a powerful country, Mark Zuckerberg did not come here to testify. The platforms are much more afraid of regulation coming from a market of 500 million than regulation, or even just complaints and bad publicity, coming from a country of 60 million people, albeit speaking English and very important. Whatever people think about Brexit and the other ways in which Britain might diverge from European Union regulation, this is one area where it is important that Europe displays a unified front.

Lord Lucas: If you have an example that would help us get our heads around how we might compel Facebook to give you data, would you share it with us, please? You can do it in writing, if you want to give us something to look at.

Professor Cristian Vaccari: I would be very happy. There are conversations happening now, as a result of this meeting in Brussels, which are trying to do exactly that. For example, you could have a data tax in which you say that, to operate legally in this country, you need to release some data. Obviously, what kinds of data is very sensitive and needs to be worked out with engineers, social scientists and computer scientists. It is not something we can write on a piece of paper right now. The people who have access to the data need to operate in trusted academic institutions for the public good, without any trace of other agendas. Just as companies have to pay taxes, they can be asked to provide data. It is just a question of how to do it, which is a very technical question that will take time to figure out. My prediction is that, if the principle is not enshrined and enforced, it will take too long for us to get the data that we need just by virtue of co-operation.

Professor Helen Margetts: Some regulators have the power to demand data in some areas and this could be one of them. That could also be an avenue to explore.

Q56 **The Chair:** I have one last question and a couple of troubling observations. Before you arrived, we talked about Estonia and the fact that its geopolitical position forces it to take certain actions. For the most part, they seem to have been effective. We do not seem to be in that situation of crisis but, post the election, we could end up with an incumbent Government who were not that crazy about changing anything and might quite like the status quo. That is an issue for us as a Committee.

It is sort of history, but I studied the communications industry a lot. If you look back at the 1952 debates in both Houses of Parliament on what then became the Television Act, there was a realisation that this changed a lot. You were introducing advertising into the information sphere. The debates were sensational, because all the right questions were being asked and the resulting legislation was really thoughtful. The creation of commercial television on a regional basis, with public service obligations, was quite brilliant and that debate, for example on the online harms Bill, is not being held in this country. That is a real worry. This is what I go to bed at night worrying about.

The last question is really an exam question: if the Government could choose one thing to improve the internet's effect on public debate, what in your view would it be?

Professor Helen Margetts: It always has to be a multilayered response, so it cannot be one thing. I go back to my monitoring and measuring point. We have finally started to worry about the environment and climate change, and to collect systematic data on it. We have to care enough about democracy to do the same thing, so I will go for that for my one. I talked about that with particular respect to hate speech, but it applies to anything. We need to monitor our democratic health and stop resting on our laurels, because we have some advantages and a very old democracy. It is not a binary thing; there are all sorts of indicators. This

is where democracy happens now and there is no going back. We cannot get rid of all this. It is where democracy happens. If we want to make democracy better, this is what we have to tackle.

Dr Martin Moore: I would say vision. I have yet to hear someone sketch out a view of where we might be in 10 or 15 years, whether we will be highly reliant on a small number of west coast US platforms, whether we will have our own equivalent platforms, whether there will be many platforms or few platforms and whether there will be public interventions in the way that the BBC was a very large public intervention in the 20th century. I would like to see people sketching out visions. I am not saying we will eventually reach them but, at the moment, everything seems quite knee-jerk and reactive, in response to particular crises that regularly erupt around the platforms and digital information economy. Once we have at least a vision, we can start to figure out what paths might get us there.

Professor Cristian Vaccari: We have not mentioned today the issue of digital skills and literacy. This is one area where any Government have a responsibility. It might not solve the problems today, but it might begin to make sure that the citizens of 10 or 20 years from now are better equipped to deal with these problems. All research on digital literacy shows that there are some appalling gaps in the ability of students in schools to parse information, to guard against misinformation and to guard against stealth campaigns. One thing that the Government do well in this country is to provide education, not just to children, but to people all through their lives. We need to invest in that area.

The Chair: Thank you very much. I personally get nul points for timekeeping, but thank you very much indeed for your patience. It is much appreciated.