



# Communications and Digital Committee

## Corrected oral evidence: Media Literacy

Tuesday 25 March 2025

3 pm

[Watch the meeting](#)

Members present: Baroness Keeley (The Chair); Viscount Colville of Culross; Lord Dunlop; Baroness Fleet; Baroness Healy of Primrose Hill; The Bishop of Leeds; Lord Mitchell; Baroness Owen of Alderley Edge; Lord Storey; Baroness Wheatcroft.

Evidence Session No. 1

Heard in Public

Questions 1 - 16

### Witnesses

I: Professor Sander van der Linden, Professor of Social Psychology in Society, Department of Psychology, University of Cambridge, and Director, Cambridge Social Decision-Making Lab; Dr Mhairi Aitken, Senior Ethics Fellow, Public Policy Programme, The Alan Turing Institute.

### USE OF THE TRANSCRIPT

This is a corrected transcript of evidence taken in public and webcast on [www.parliamentlive.tv](http://www.parliamentlive.tv).

## Examination of witnesses

Professor Sander van der Linden and Dr Mhairi Aitken.

**Q1 The Chair:** Welcome to this meeting of the Communications and Digital Committee in the House of Lords. My name is Baroness Keeley and I am the Chair of the committee. I would like to welcome our witnesses and thank them for joining us today in the first evidence session of our new inquiry into media literacy. The session is being broadcast live and a transcript will be taken. Our witnesses will have the opportunity to make corrections to that transcript if that is necessary. As this is the first session of a new inquiry, committee members may need to declare relevant interests before they speak. I will say that I do not have anything to declare.

Can I ask our two witnesses to introduce themselves and then I will move into the first question? Dr Aitken, can you introduce yourself and your expertise?

**Dr Mhairi Aitken:** Good afternoon, everyone. I am a senior ethics fellow at the Alan Turing Institute. The Alan Turing Institute is the UK's national institute for AI and data science. In my role as a senior ethics fellow I look broadly at social ethical considerations around advances in AI technologies, and I have a particular interest in public engagement relating to AI.

**Professor Sander van der Linden:** I am a professor of psychology at the University of Cambridge and I study how people are impacted by misinformation, what it is about the brain and psychology that leads people to endorse false beliefs, how it spreads online, and how we can design interventions to help people.

**Q2 The Chair:** Thank you. I will kick off the questions by asking the first one myself. Can you describe for us what the most pressing risks and threats are to the UK's online information environment and to people using information online?

**Dr Mhairi Aitken:** For me, some of the most pressing risks relate to advances in generative AI technologies. Generative AI technologies are a broad category of AI technologies that are used to create new content. That might be images, video, text or audio. The advances in generative AI technologies are making it increasingly difficult to reliably or consistently identify what is real and what is AI generated. As AI-generated content is proliferating online, particularly across social media, it is becoming increasingly difficult to consistently and reliably detect what is real and what is AI generated.

The threat here is not just that people might increasingly see or hear something that is fake and believe that it is real. The deeper threat here is that increasingly, as there is exposure and awareness of AI-generated content, people begin to lose trust in all content online. It is not just that we might see or hear something fake and believe that it is real.

Increasingly, people will see and hear things that are real and the first

reaction might be, "How do I know that is not fake? That could be AI generated". It is a particular concern when what people are confronted with or what they encounter challenges existing ideological positions or political viewpoints that the initial response can be that it could be fake. It is becoming increasingly difficult to have access to reliable information about what is fake and what is AI generated.

That is in combination with the ways that we know that people are increasingly accessing news and information on social media, and AI in other forms has been used by social media for a long time to personalise, filter and tailor content online. Increasingly, people are accessing information in echo chambers and bubbles where it is not so much about the veracity or the authenticity of the information, it is more about who is telling you that information and who you trust on claims about what is real and what is AI generated. That is the perfect breeding ground for conspiracy theories and for a wider erosion of trust in the information ecosystem. For me, that is the biggest and most pressing threat that we have to deal with in relation to media literacy.

**Professor Sander van der Linden:** I agree, and I will add some risks. For me, it is looking at the impacts of misinformation. Last summer we saw that a false news story can erupt into national violence. We sometimes use the term "stochastic terrorism", which is the idea that misinformation cumulatively builds over time and then erupts. The challenge is how you can predict and counter such activity.

It is not just the riots. Misinformation can be dangerous for people's personal health, but it can also undermine democratic processes and elections. There are indirect effects that have been mentioned already, eroding trust in each other, in the media, in democratic institutions. That is a more slow, long-term project. It intersects with several related challenges. There are the impacts of misinformation but then who is spreading it? Some of it is foreign information manipulation. We are seeing increased activity from potentially hostile nations that try to disrupt the information discourse in the UK.

The question is how effective that is and how resilient the UK population is. If you look at indices of information resilience, the UK often does not rank in the top 10. There are many reasons for that. Some of them are structural and have to do with freedom of the press. Some of them have to do with inequality. However, part of it has to do with how well prepared the population is at the individual level, media literacy, science education and other measures that we can use to help people to identify propaganda.

We are vulnerable in the sense that we do not have a good first line of defence. Look at countries such as Finland. They are educating very young children early on about the techniques of propaganda, how to identify it, and how to neutralise it. That is a real risk for the UK in the long term.

I agree that AI and inauthentic activity is an increasing challenge. Some people say that what we are seeing right now is that misleading information is most prevalent, but if you look at Ofcom data, for

example, 40% of people say that in the preceding month they have seen misinformation in the UK. Out of those 40%, 90% say that they are very concerned about the impacts of misinformation and about 20% say that they have seen deepfakes and fake audio. As technology is improving, as that escalates, thinking about LLMs and ChatGPT, the risk of AI and inauthentic content will increase in the future. The thing is that maybe right now not everyone is flooded with deepfakes, but they might be in the future if we do not do something about it. I think that there is an urgency to this problem and that is part of the risk assessment.

**Q3 The Chair:** Thank you. As the second part to that question, the technological developments are evolving very quickly, but how might those risks, threats and challenges evolve as a result of technological developments? What are you seeing in terms of those risks?

**Dr Mhairi Aitken:** Generative AI technology is continuing to advance at a rapid pace. Equally, there is interest and investment in tools to detect AI-generated content online and, interestingly, on approaches around, say, watermarking and ways of labelling and identifying the content online as either authentic or AI generated.

The problem here is that the advances in generative AI technologies are often outstripping the advances in developments around those measures or those second call responses to address those challenges. The output that these tools create is becoming more and more sophisticated. It is becoming more and more compelling and that is posing additional challenges. In the past we could point to ways of potentially identifying what might be an AI-generated image. There are tips and hacks that you can use and things to look for to potentially identify what might be AI generated. However, this is continually advancing and those tips and tricks are becoming less effective or less reliable.

We really need to move beyond approaches that are focused on just giving individuals the skills or the ability to identify what might be AI generated and to have effective regulatory and policy responses and focus on the responsibilities of platforms to protect users from the spread of misinformation.

**Professor Sander van der Linden:** I agree that the AI challenge is a big one. If you look at problems like micro-targeting—that is where they try to tailor fake news based on people’s digital footprints, what websites they are visiting and what information they are clicking on. They try to aggregate all that data and then they try to target people with messages, sometimes true but it can also be fake. Research shows that this is more persuasive than untargeted messages.

It used to be the case that, let us say, Russian troll farms employed people—hard labour—to do this. This can now all be automated with large language models. They can turn out hundreds of variations on a story based on people’s circumstances to try to optimise the effectiveness of those messages. That is a growing concern. Nano-targeting is the ability to target and locate a single individual who will be persuaded by a particular story. That can be done with AI. There is a

long-term risk in how that is affecting people and what we will do about that.

In terms of how this is evolving, the European Commission, for example, has the Digital Services Act so it has the means of addressing some of these risks. The Online Safety Act currently does not have misinformation and disinformation present, which limits Ofcom's ability to intervene and potentially do something about this problem. I think that is extremely relevant in light of the evolving landscape of most social media companies. We have seen Facebook get rid of its fact-checking programme in the US. It is getting rid of its moderation policies. It is not following its actual policies that hate speech is not allowed. It has rules but it is not enforcing its own rules. Many researchers have shown that. Misinformation is going viral on these platforms now without intervention. While they are rolling back intervention and letting disinformation run free under the banner of free speech, there are no good regulatory ways to deal with this problem now if it is not going to be addressed by, for example, the Online Safety Act. Of course, we have to balance freedom of speech with speech that causes harm, but the UK has a long history and track record of doing so.

**Q4** **Viscount Colville of Culross:** Good afternoon. Dr Aitken, you talked about a concern that by endlessly telling people that they are exposed to this misinformation and disinformation you will build up an over-scepticism and people would not trust anything at all. How do you balance that so that you make sure that people are aware that there is misinformation and disinformation but they do not become sceptical about everything?

**Dr Mhairi Aitken:** That is the key challenge here. Ensuring that there is awareness of the use of generative AI and how that can be used to mislead or to create AI-generated content online and things that may not be real, although they can be very convincing, is very important. It is important not just in terms of media or news but also in how it impacts individuals. We have deepfakes being created, or sexualised images particularly being created of women and girls, so this awareness of generative AI technologies is important beyond just the media landscape. It also impacts individuals, so people know that they can contest these things when they affect their lives and they can have recourse to the ways it impacts their lives.

What is important is that when we are thinking about this in the media landscape we do not place the burden on individuals to have to scrutinise all information or to have to be constantly questioning whether this is real or fake, source checking and fact checking where this information comes from. That responsibility should belong with platforms where this information is being shared and it should be the responsibility of regulators to address that. I think that it is important to get that balance right. It should not be a burden that is on individual users to have to constantly be scrutinising that landscape.

It is also important to note that, even when we advocate for approaches to increase understanding of how you might be able to identify what is

an AI-generated image—there are things like looking at the details in the background, which are often a bit blurry, looking at text and images because it is often garbled. In the past it has always been fingers on hands, because they tend to not be quite right. The reality of how images are shared on social media is that they tend to be lower resolution and they tend to be quite small. People are scrolling past on their phone, they are not stopping and spending a long time scrutinising every image. We have to be aware of the context in which people are interacting with this content and, again, not placing the burden or the expectation that people will be taking the time to scrutinise everything, to fact check everything and to verify everything. There need to be additional mechanisms in place to provide the necessary safeguards in that context.

**Q5** **Viscount Colville of Culross:** I was involved in the Online Safety Bill and I was one of the people who was worried about putting misinformation into the Bill. I can see that with disinformation it is quite possible that you work out who the authors of that were and then you rule against them, but with misinformation it is very difficult to be able to decide what is right and what is absolutely wrong in many different cases. Do you really think that we could have a regulator who would be able to arbitrate on what is true and what is not true in some of these very blurred areas of news, for instance?

**Dr Mhairi Aitken:** You are right, it is not always a clear-cut question of what is true or not true. There should be ways of being able to reliably identify what is AI generated and what is real, and that should be identified and labelled so that that is up front and visible.

I suppose the gap in the Online Safety Act is that there was a new offence of false communications, but that only relates to the deliberate sharing of false information with the intention of causing non-trivial harm. That is a very specific and important development, but a lot of the harms we are talking about here today on misinformation are more the cumulative impact of the volumes of misinformation that are shared and the impacts due to the accumulation of that and the repeated exposure to misinformation on social media. We do not yet have an effective mechanism to address that significant impact. That is the responsibility of platforms to take accountability for the impact that that is having through the volume of misinformation being shared online.

I agree with the point that it is a significant challenge, and it should not be the role of an arbiter as to what is true and false, but being able to identify what is AI generated and what is authentic is something that we really should be working towards.

**Q6** **Lord Mitchell:** I would like to carry on with the point that was raised about scepticism, or perhaps another word we could use is gullibility. In your studies, have you noticed that there are differences between age and sex of the population, the countries they come from and their experience, whether they are democracies, whether they are not? If a fact is told to a person in China, are they more likely to accept it than

somebody here?

**Dr Mhairi Aitken:** This is probably more a question for Professor van der Linden.

**Professor Sander van der Linden:** In the work we have looked at that. There are some differences. Some of the demographic differences are quite small. The biggest one that stands out is the generational difference in susceptibility to misinformation: younger audiences are more susceptible. It also works differently cross-culturally. If you compare interventions in western Europe, they all pretty much work similarly, but when we go to rural India, for example, or Brazil, then people have a very different understanding of the media context. We do some work in active conflict areas as well—let us say Ukraine and Russia or Israel and Hamas. Things require more adaptation in those contexts because people think about consuming media differently. They use social media differently, so it does need some adaptation.

People do think about media. In some regimes people are more sceptical of the media because it is state controlled. In other countries, people might accept it as given because they are worried about speaking out. For researchers it is difficult to know because, when you survey people, some people may feel that they cannot freely speak their mind when you are in, for example, an authoritarian regime. It is difficult sometimes to get good insights about the differences in people's psychology based on the regime that they are in.

I would say the differences between males and females are small, but if you look at some of the work on conspiracy theorists, for example, in the West at least they tend to be predominantly male. There are certainly females who also endorse conspiracy theories and it is different in different countries, but on our side of the world a lot of the conspiracy theories are consumed by men online, for example. There are different groups. If you look at Ofcom's data, it suggests that some groups—the LGBT community and minorities—are being targeted with more misinformation so they might be more vulnerable. Yes, there are differences but sometimes they are fairly subtle.

I do recognise the point about scepticism. It is an interesting point because you want people to be sceptical, I suppose, but not cynical. We do not want people to dismiss all media, but we want them to be sceptical enough to be good discerners. If you just tell people, "This is a big problem and there is misinformation everywhere", then people will just be sceptical because you have not given them the ability to discern between the two. That is often what we forget. In order to get discernment and healthy scepticism, you need to give people the tools and the ability to discern because the warnings in themselves just make people sceptical and distrustful. That is an important distinction.

I do think that we can do that in the Online Safety Act. Climate change is real, the earth is not flat; people say all sorts of things online, and we just know that some of it is misinformation and some of it is factual information. For a large segment of content, we can confidently say what is science based and what is not. There are situations where it is tricky,

where science may be emerging or where there is some uncertainty. I do not think that it is Ofcom's role or remit to tell individual users what is misinformation or not, but it can consult experts. It can talk to social media companies and say, "Provide some more context for people to make an informed decision. Regardless of veracity, there is some manipulation going on here—inauthentic activity. Maybe this is AI. Maybe this is lacking crucial context". It can help social media companies to do better in delivering that under the umbrella of dealing with misinformation. I do not think that it should be the case that Ofcom will tell people what is misinformation or not, but certainly for a large amount of content we can verify veracity. That is what we have fact-checkers for. We do not want to throw the baby out with the bath water, so to speak.

**Q7 Baroness Fleet:** Dr Aitken, I was interested in what you were saying about identifying and labelling AI content. Is that possible and how does that work? Is anyone doing that? In the old days, it would be a bit like the difference in a newspaper between an advertisement and an editorial. You would be clear which was which. How can that be done and who does it? Will the platforms do it? What is the process for that?

**Dr Mhairi Aitken:** It is a very active area of research. There are tools that have developed to detect AI-generated content. The problem at the moment is that most of these tools will be trained on particular models. You can have a tool that might quite reliably detect wherever, say, an image was generated by one particular generative AI model—one particular image generator—and it would not detect if that image had been generated by a different model. As these tools are trained to identify outputs generated by particular AI models, once they are available and we are using them, new models come along and images might be developed by a subsequent generative AI model that does not get picked up. There are real challenges at the moment in creating something that is robust and reliable and can detect outputs by any new model that has been developed and is being advanced.

There are also approaches to labelling and identifying AI-generated content, such as digital watermarks. These can be watermarks that are embedded within digital content, so not necessarily visible to the eye, but they would be picked up by an AI tool—a detector. Again, these rely on willing compliance. If you are creating an AI-generated image or AI-generated content, you might voluntarily put a watermark on that. That might become something that we require for compliance, but malicious actors can fairly easily evade that. As I said, most people are consuming low resolution images or content that is shared on social media. It might be that when an image is cropped or otherwise edited or shared in low resolution, the indicators that it was AI generated are lost in the way that it has been altered or shared. It is very much imperfect but these are active areas of development. It is important that we continue to advance these technologies and the tools and systems that can do this.

It is important to note that we should not necessarily be expecting members of the public to be using these tools to check systems. Rather,



we need trusted fact-checking bodies and trusted organisations that can provide that and can be a source of authenticity and credibility in this space.

**Q8 Lord Dunlop:** Can I ask about metrics and how you measure the effectiveness of policy interventions? Professor, you mentioned information resilience indices. Can you tell us how they are compiled and what factors are taken into account in compiling them?

**Professor Sander van der Linden:** That is a great question. There are different outcome measures and it depends on which ones we value and prioritise. The most obvious measure is whether people can now discern between manipulative and non-manipulative content. Of course, that does not mean that people will necessarily change their behaviour or the way they vote, but maybe that should not be the goal either. It depends on how you measure it.

The first measure is whether people can distinguish between manipulative and non-manipulative, fake/true content, if that is the dimension you are using. How confident are people in their abilities? Are they calibrating it correctly? Are they overconfident? To what extent are they sharing this material less with other people? People sometimes share things and they do not necessarily have to believe it. Sharing has different psychological motivations than believing, so sharing is another outcome measure that often relates also to the information that people are exposed to. Sharing is the metric of propagation in social media networks.

Then there are the more difficult measures that we often do not have access to because of issues with social media transparency. If we do an intervention on, say, YouTube and we find that indeed the ad helps people to identify manipulative and non-manipulative content, are people watching fewer extremist videos after this? We do not have access to viewing data; only YouTube does. What videos are people watching, for example, or what websites are people browsing afterwards? Often we do not have the crucial behavioural data. We can do it in schools with students. We can do those experiments, but not when it comes to social media. They are the only ones who know what people are clicking on and what they are sharing. We have co-operated with them in the past. You have to sign non-disclosure agreements, and they may or may not do something about the results of the experiments. There is a huge issue around transparency.

I would say there are the behavioural outcomes, then there are the more psychological outcomes, and then you can think about population-level outcomes as well. In terms of self-reports, do people feel that their skills have improved? Sometimes what people feel is not always aligned with what is true. Social media companies like to know whether this works without hurting engagement on the platform and whether people still trust them and so on. We are more interested in objective measures of efficacy: does this help confer skills and change behaviour?

Here is an important point, though, and this is why we do games and videos and use humour. You can have the best intervention in the world that maximises effectiveness but if nobody engages with it, you will not achieve anything at population level. It has to be engaging, and that is the world we live in. It is about how to make it engaging without sacrificing accuracy. It has to be both effective and engaging, and that is the tricky balance to keep in mind.

**Q9** **Baroness Owen of Alderley Edge:** I will start by declaring my interest as a guest of Google at its Future Forum.

We spoke about the motives for sharing and I would like to dig into that a little bit more. In some cases, somebody will obviously not believe the content that has been put in front of them but will share it anyway. I would like to dig a little more into what those motives are. Is it for fun or because they deliberately want to spread the misinformation?

**Professor Sander van der Linden:** It is a great question. There are various motives. We ran an experiment once where we asked people about headlines and whether they were true or false. Mostly people did not do so well. In another condition, we paid them to be accurate. You would get a small financial bonus to give the correct answer, and all of a sudden people know the right answer.

Sometimes it is not about whether they know it or not; they do know but they are trying to signal something. They are sharing because they want to express their identity: "I belong to this organisation; I belong to this group; this is what we believe and that is why I am sharing the information, to bolster the identity of the group and not necessarily because I personally believe it". People who are part of extremist groups might share content that they do not personally believe but they want to bolster the reputation of the group, for example. Bad actors sometimes intentionally share misinformation because they know it is false but they do so for political or financial reasons. They are being paid or they gain political currency from sharing false information. Financial, political and social—those are the three key motivations.

**Q10** **Baroness Wheatcroft:** Both of you have left us in no doubt that there is a real problem and that as far as you are concerned platforms should do a great deal more to protect the public. In an ideal world, yes, but given how prevalent AI is becoming it is asking a lot for them to signal everything that is based on some AI. What else should we be doing? Professor van der Linden, your presentation earlier was fascinating and clearly you can inoculate up to a point.

**Professor Sander van der Linden:** To a point, yes.

**Baroness Wheatcroft:** What age should we start with? How should we start? Should laws be very formalised in doing it? Given that Gen Z is already quite vulnerable, how do you both think that we should reach them?

**Professor Sander van der Linden:** We should start early. Sometimes people talk about being in an information war and the only solution to

that is to empower our citizens from the earliest age possible. Of course, we need to do that in a non-political way. Some of the manipulation techniques that I have discussed everyone usually agrees are bad and undesirable and they are affecting young people online. We talk to kids all the time and they are embroiled in wellness grifts—pseudoscience. There is all sorts of stuff out there that leads people astray, and they are digital natives, so clearly that is not enough.

I think that we need to have a national educational curriculum where we implement and trial at population level some of the best-known interventions that research has produced and implement those in schools. We talk about booster shots, to further the vaccine analogy, but one dose will not do it. We need to do this once a year, maybe every year of school, repeat it to make sure that we have critical and healthy scepticism and consumers who can make up their own minds free from manipulation. That also to some extent, I hope, would reduce the need to police anything, but at the same time we need top-down regulation to make that happen as well. Other countries that have done that well are some of the Nordic countries like Finland and Sweden, which also score very high on freedom of the press, for example. That is part of what needs to be done.

As I said, we have worked with Meta and with Google. We did great things and then when their incentives changed they stopped doing them. We cannot rely on actors whose interests are not in the public interest in the long term. That is why I think that we need to perhaps look at education and regulation.

**Q11** **Baroness Wheatcroft:** We are looking at doing international comparisons in later sessions and we are very interested to see what other countries do. Clearly, Finland will be worth studying. How does a Government or an industry reach the older generations without it simply being perceived as propaganda?

**Dr Mhairi Aitken:** It is an important area, and it is important to recognise that across society everybody is being impacted in this area. Media literacy is relevant and important for all groups within society. Within any demographic group that we focus on, there will be a wide range of experiences, levels of exposure, levels of interest and levels of understanding. It is important to recognise that.

What is also important here is that we are talking about trust. It is not just trust in the information that people access online but also trust in the claims that are made about that information and the messages that we are communicating in relation to media literacy. It is important that that engagement comes through trusted organisations and from trusted sources.

At the Alan Turing Institute one of the projects that I am working on is around public voices in AI. This is a project that is funded by RAI UK. In that work, we are looking at existing approaches to involving the public in conversations around AI and engaging the public around AI. What has come out very clearly from that work is that a growing number of

community groups and community organisations are engaging the public in these areas, involving the public in conversations around AI, but these are not necessarily groups where AI is the central focus or the central topic. They might be community groups that are focused on education or healthcare or welfare or housing or local community issues, but through those established relationships and those established community groups they are engaging in topics relating to AI because that becomes relevant to the central topics they are interested in. That is also true in the case of media literacy. Media literacy is something that is relevant to all these different areas, whether it is about health information or information about housing, welfare, rights or whatever it might be.

When we are thinking about approaches to media literacy, it is important to recognise that there are these existing networks of community groups and trusted organisations that have the relationships needed to take on that role and to have those conversations about media literacy or AI or generative AI. What they need is funding and resourcing. They need sustainable funding to enable those relationships to endure and continue. Those are the networks through which this information can be communicated within the context of relationships of trust. It is also an opportunity to discuss it in ways that are meaningful and relevant to different community groups. I think that is an important pathway forward.

**Q12** **Baroness Wheatcroft:** I would like to ask this question to you both before allowing others to come in. Propaganda has always been around. I wonder whether your research has shown you that parts of our society for all sorts of reasons are more vulnerable and more gullible towards propaganda than previous generations have been. I should declare my interest, though not in relation to that particular question, as chair of the *Financial Times's* appointments and complaints committee.

**Dr Mhairi Aitken:** It might be more for Professor van der Linden on the point about vulnerability to propaganda. I would say that it is important to look at different groups' experiences.

On the point about children and young people, I lead a programme of work at the Turing on children and AI. Over the past four years we have been working particularly with children between the ages of eight and 12. Earlier this year we held the Children's AI Summit, which brought together 150 children between eight and 18 to discuss their experiences around AI and their visions for how they wanted to shape the future of AI. A headline finding from that is that they are very competent, capable and enthusiastic to engage in these discussions, but children and young people's experiences with these technologies can be very different from adult experiences. That will be the same with different groups that we look at. There are particular unique needs, experiences and perspectives.

With children and young people in particular, we often make assumptions about the ways they might interact with these systems and what the impacts or the risks might be. It is only by directly engaging with children and young people to understand their actual experiences and

their actual concerns that we can really understand the approaches that need to happen.

That is perhaps not directly answering your question about vulnerability to propaganda, but it is just to recognise that we really need to start with an understanding of the actual experiences of the different groups that we seek to engage with.

**Professor Sander van der Linden:** There are some changes now. It is true that we have always had propaganda. People were polarised before social media. Social media is not the cause of all negative things, but it is certainly an accelerant. In my book, I calculated how long it took for a false message to travel in the Roman Empire and then compared it to WhatsApp, for example. People used to get their information from a few select sources and now we are bombarded with information. When the brain is stressed out by multiple channels, it resorts to simplifying rules of thumb or heuristics. The way we are interacting with information is fundamentally different. It is Snapchat, WhatsApp, cable news, regular TV, newspapers, media, friends and family. We have so many more sources of information that we are constantly dealing with, so it is hard for the brain to be able to still discern in those contexts. That is an additional challenge that technology brings.

It is also about information inequality. Some people do not have access to good-quality information. It is very unequal. You see it online. If we look at echo chamber formation—this was mentioned already—some communities are embroiled in low-quality information, whereas other communities find themselves in mostly factual information environments. When we speak of average exposure, it is tricky because the distribution is very unequal. We know that minorities are targeted more intentionally with disinformation during elections but also around health, for example. It is complicated because some minorities do have good reasons to be sceptical of official authorities because they have been experimented on in the past. That might not be the case now, but how we rebuild trust with those audiences is important.

Ofcom's data shows something interesting. A third of the UK population does not think that journalists are doing their job in terms of abiding by good journalistic practices. Trust in true news, in mainstream news, is also quite important. To your question about elderly and older generations, they are hard to reach. They are not in the education system. They are less likely to come out to public events. The one area that they do tune into is the mainstream media, so what can the mainstream media and journalists do to help solve this problem? That is another question to look at.

**Q13 The Bishop of Leeds:** You have mentioned several times the role of regulators, of Government, and so on. I am haunted by the fact that in our two previous inquiries we had very vocal arguments against involving Government in anything because they were seen as part of the problem. Both of you have referred to the need for the tech platforms to take responsibility—that they are part of the problem but could be part of the solution. What is the role of Government in particular and then, following

on, regulators, platforms and the media industry in boosting public resilience to the threats that you have talked about and improving media literacy?

**Professor Sander van der Linden:** That is a big question. As a psychologist I will just be offering my opinion on that. If we look at, for example, Elon Musk's feed about how the Prime Minister is supposedly imprisoning people for posting things on social media, there is a fundamental misunderstanding between different countries. Hate speech is illegal in the United Kingdom, where it is not in the US, so that is a big difference. People often misunderstand that there are different laws.

I am not a prosecutor or a lawyer, but I think that what it comes down to is that we need to have systems of accountability. I think that very few people would jump at the idea of having a government-led ministry of truth. Most people do not think that the Government should be policing everyone's speech. However, we need to have a system of accountability that comes from regulators to hold media companies to account. That is, of course, why we have Ofcom and regulators such as Ofcom, but they also need to be able to look at problems of mis- and disinformation and AI risks. If people are creating manipulated videos and duping parts of the population during an election, that seems something that should not be acceptable. There should be intervention and regulators should be empowered to potentially do something about that.

Ofcom will not police people's speech, but if you take the riots as an example, if a story goes viral and if it is determined that the risk of national violence will go up drastically, maybe there should be some intervention and encouraging of relevant authorities and social media companies to try to do something about that problem rather than standing by and saying, "We will just let this happen" without any accountability.

That is the issue for me because, without systems of accountability, you get vigilantism. There is always accountability. If you say something false and somebody else does not like it, they might show you how much they do not like it in various ways. We need to have rules and regulations to co-ordinate speech in some sense, not to police speech but to have accountability for speech, especially when that speech is false and harmful. It is the Government's role to make sure that those systems are in place rather than to tell people what they need to believe. That is everyone's concern. We know what happens when Governments tell citizens what they need to believe, so that is not a system that anyone is suggesting, but a system of accountability would be my take.

Q14 **The Bishop of Leeds:** Governments can easily say, and they do, "That is the role of the regulator". In fact, it is up to the firms then to decide whether they play within the gap, if I can put it like that. Do you think that there is space for more cross-government work on this for joining up how different government departments work to protect the public?

**Professor Sander van der Linden:** Yes. I think that there should be a lot more cross-government co-ordination. In my independent chats with

different parts of Government, people in defence often do not know what is going on in the police council. The police council is not always talking to counterterrorism and then the communications department is not clued in to what other departments are doing. There is no national co-ordination on this problem, which leaves vulnerabilities that foreign actors are exploiting in terms of disrupting and intervening in UK discourse to try to sow chaos and confusion, because we do not have a co-ordinated authority to streamline a lot of the issues that we are talking about.

Foreign manipulation is a different remit, for example, from, let us say, domestic media literacy. Everyone is looking at each other as to who is responsible for this. The regulator? The regulator is not empowered to look at mis- and disinformation. With the European Commission and the Digital Services Act, of course, it is up to individual member states how they implement things, but they do have a Digital Services Act. They can fine social media companies for spreading misinformation. They can change incentives. The CEOs of social media companies might not listen to the regulator or accept fines, but at least they are empowered to try to do something.

**Q15** **The Bishop of Leeds:** Dr Aitken, in relation to that, we know that ethics always follows technology, and the technology is developing and expanding far in advance of our ability to understand or, in one sense, regulate or control it, so there is scepticism, if not cynicism, about some of the conversations we have. What impact, if any, do you think that the implementation of the Online Safety Act will have on the risks and threats that we have been identifying in this conversation?

**Dr Mhairi Aitken:** The Online Safety Act is a significant and important step forward for sure, and it is making important advances in addressing some of these risks. I am particularly thinking about risks to children and young people and vulnerable users online. As has already been mentioned, there is this important gap around not having the powers to address the risks of misinformation beyond deliberate spread of false content with the intention of harm. That is a significant gap, and what we have been discussing here today is around the broader risks and harms that can be caused through the proliferation of misinformation online. I think that is a gap that needs to be addressed, and how to address that needs to be considered.

To your point on cross-government action or the role of Government, it is important that we see this as an issue that cuts across all policy areas, all sectors, all parts of Government. Media literacy and the risks of misinformation are significant to all policy areas, so there is real value in having a joined-up cross-government approach to develop approaches to address these challenges, recognising that media literacy is important to underpin a healthy, functioning democracy. This really is an area that is central to the remits of all policy areas, and having more cross-sector and cross-government joined-up approaches will be valuable to move this forward.

**The Bishop of Leeds:** It is actually not in the interests of some to have

that cross-government co-ordination working effectively.

**Dr Mhairi Aitken:** If it sets the ground rules and the boundaries of what is acceptable, then that is in the interest of us all. It also creates the environment in which we can innovate and develop new approaches safely and responsibly. There are different views on that, but I think that it is in the interests of all society.

Q16 **The Chair:** Thank you. We are running out of time now. Professor van der Linden, is there anything you want to add on the last points?

**Professor Sander van der Linden:** No, I think that was well said.

**The Chair:** Excellent. That is a very good start to our inquiry today. You have answered our questions well on the risks and challenges and we are starting to look at how well equipped we are to deal with them. There is a great deal left for us to do, but thank you very much for your time today.