HOUSE OF LORDS

Communications and Digital Committee

Corrected oral evidence: The future of news: impartiality, trust and technology

Tuesday 16 April 2024

2.30 pm

Watch the meeting

Members present: Baroness Stowell of Beeston (The Chair); Lord Dunlop; Lord Hall of Birkenhead; Baroness Harding of Winscombe; Baroness Healy of Primrose Hill; Lord Kamall; Lord Knight of Weymouth; The Lord Bishop of Leeds; Lord McNally; Baroness Primarolo; Baroness Wheatcroft; Lord Young of Norwood Green.

Evidence Session No. 10 Heard in Public Questions 100 - 108

Witness

I: Ciaran Martin, Professor of Practice in the Management of Public Organisations, Blavatnik School of Government, University of Oxford, and former CEO, HMG's National Cyber Security Centre.

Examination of witness

Ciaran Martin.

Q100 **The Chair:** This is the Communications and Digital Committee. We are continuing our inquiry into the future of news. The theme of our hearing this afternoon is disinformation and misinformation. We have two sessions. The first we will start in a moment, and then we have a panel of three joining us later. Could I ask Professor Martin to introduce himself? Why do you not just very briefly set out your credentials in this area?

Ciaran Martin: I am a professor of practice at the Blavatnik School of Government at the University of Oxford since 2020. I also work with a number of cybersecurity companies, in terms of declarations of interest. That is all publicly available. I am very happy to give it to the clerk directly. Prior to that, I was a civil servant for 23 years, the last seven of which were involved in setting up and running the National Cyber Security Centre within GCHQ, on whose board I served for nearly seven years. Part of that role was protecting the UK's democratic infrastructure from cyber interference, which touched on, although was not directly responsible for, issues of disinformation campaigns and so forth.

The Chair: We are very grateful to you for being here. Thank you very much for giving up your time. We will cover four themes. We will talk about how disinformation has evolved over the last decade. Then we will want to talk about how new technology is impacting disinformation. Then we will come on to the policy responses from government and the industry. Then we will get to some questions at the end on what you actually think about the general discussion and discourse about misinformation and disinformation. I will hand over straight away to Lord Hall to get us going.

Q101 **Lord Hall of Birkenhead:** Professor Martin, it is great to get your perspective on the issues challenging us in terms of disinformation. I shall start off with a really broad question, about the nature and level of the challenge posed by disinformation to the UK, balancing it between foreign actors and domestic actors as well. If you could, will you give us your perspective from the roles you have been carrying out about how that threat challenge has evolved over the last decade?

Ciaran Martin: The first thing I would say is that, because of the nature of my professional duties in government, I would have approached this practically and operationally from the point of view not so much of disinformation as a standalone problem, but particularly of potential foreign interference in democracy, democratic process and general media discourse, not least because of events such as the interference by the Russian state in the 2016 election in the United States, not via disinformation but by a different route. To that end, disinformation was a strand of an approach by adversarial nations, particularly Russia.

You might think of it as different prongs of the same pitchfork. One might be what we call hack and leak, which is what the Russians did to

the Democratic campaign in 2016. That is not disinformation, because essentially what they did was to steal information that was true and then promulgate it through various outlets. That was accompanied, particularly over time, by the rise of fake videos on Facebook and of creating information deliberately that was not true as part of that campaign. In terms of the evolution of that, we moved from a period where a lot of the effort was focused on hacking information and leaking it to perhaps the more nebulous forms of disinformation campaigns sponsored by states, not least because they are easier to do.

Lord Hall of Birkenhead: It was true information that was leaked.

Ciaran Martin: That was the initial focus. You still get quite a lot of that. In fact, we are dealing with some of those now in terms of attacks on parliamentarians and so forth, but it has been evolving. In about 2018, for example, we were talking about troll farms coming out of St Petersburg, where essentially, rather than have intelligence officers sitting in a building hacking quite well-protected networks and doing difficult things, you just had a bunch of people, perhaps lower skilled, perhaps being paid less, perhaps with a less formal connection to the state, just making up propaganda and lies, sticking it out there and seeing what would happen. Both of those were quite prevalent towards the end of the last decade.

In terms of how that is evolving, we will come to this under new technology, but it is easier to make better fake information now, as we move into AI deepfakes. In particular, something that has emerged in the last year is that it is really easy to make audio deepfakes. It is harder to make credible video deepfakes although, as ever, the costs of doing business for the malicious actors get lower all the time, and we can expect this trend to continue. That all sounds very concerning, and it is. However, it is important to say at the outset that I would not want to try to offer you a preliminary analysis of the problem without actually talking about the impact.

There is a tendency sometimes, particularly in quite a gloomy profession such as mine, to confuse intent and activity with impact. There is a lot of this about, and we need to be vigilant. But one of the things in terms of detecting this, unlike some forms of espionage and so on, is that, if you have a big intervention that seems to be decisive in political discourse, it is not a secret. It becomes known that there was a fake. You can find examples in some other countries, and we might come on to that—but it is perhaps an unpopular opinion thus far, but the UK has not done that badly on this front. Because this is not my direct responsibility, I do not take any credit for it professionally.

Actually, I am wary of exaggerating the impact of this because, if you tell people often enough and loudly enough that there are all sorts of nefarious interferences with democracy and political discourse, they will start to believe you, with all the consequences that involves. For me, the success measure for the Russians in terms of their interference in America in 2016 is not the outcome of the election, but the fact that eight years on, and two electoral cycles on, the American people are still

arguing about it. That is actually the impact. We ought to be very careful about confusing a great deal of malicious intent and a lot of malicious effort, with as yet—I stress "as yet"—very hard to detect and limited impact on British political discourse and the integrity thereof.

Lord Hall of Birkenhead: You have talked about the Russians, but do the Chinese or the Iranians also deal in disinformation? If so, what is the nature of the tactics that they use and the way that they campaign?

Ciaran Martin: They are all slightly different. Forgive the phrase, but you would have Russia, in terms the amount of effort that it would put into this and the targeting of it at western politics, as probably the market leader. It is a deliberate objective of Russian policy, as revealed by a decade or more of activity, to destabilise western democracies in the European Union, the UK and North America, and some other pretty obvious allies.

This reflects the fact that the origins of Russian and Chinese digital aggression, if you might call it that, are slightly different. Russia's has always been geopolitical, whereas China's has been, while not exclusively, primarily economic. A lot of that has been a hybrid of influencing through human agents. You have seen some scandals around that in this country, around the promotion of information and so forth. It is hard to say this without sounding very dovish on China, which I would not consider myself to be, but it is warier about being caught outrageously interfering in open discourse than Russia, because it keeps telling people to stay out of Chinese affairs and so forth. It is quite obvious that Russia wants to exacerbate polarisation in western societies; it is less obvious that China has a deliberate programme of doing that. It is more about advancing its economic interests, which leads to slightly different activity.

Iran has waxed and waned in this area a little. A lot of its digital activity has tended to be around quite regime-specific things, such as Iranian dissidents in this country and so forth, rather than that sort of systematic interference.

I am afraid, Lord Hall, I did not touch at all on your domestic question. I am less well attuned to it, partly because of where professional interests have taken me but also, as a board member of GCHQ, it is a foreign intelligence-gathering service as well as a cybersecurity one, so I am just less expert in it. My observation, which others may contest, is that the relative amount of UK-based disinformation is not on a par with the United States. Towards the end of my time in government at the end of the last decade, I worked very closely with US counterparts on protecting the totality of electoral systems from a whole range of threats, from actual cyberattacks on infrastructure such as electoral rolls to stop people voting, to broadcasts, which obviously you will know very well, and all the way through to disinformation.

If you look at what the Americans found in 2020, there was the whole controversy that ended up with President Trump firing the senior official in charge of the electoral protection effort. But they set up this thing called Rumor Control, which was an outstanding effort. Say there was an

allegation that polling lines in Georgia were being closed early or the ballot paper ink was not working in Arizona—these are all real cases—what the Americans found was that a concerning proportion of these rumours were being generated domestically. I just do not have that same sense in the UK. Again, I am not trying to be dismissive of the threat, but, if you are looking at managing disinformation around American politics, you are now looking at a very significant set of domestic operators. I do not have the same sense of that level of activity here in the UK.

Lord Hall of Birkenhead: I am not going to run around every country in the globe, but just one other is North Korea, which I seem to remember at some point was fairly active. Has that also been an area of disinformation?

Ciaran Martin: It has not in this space. As you say, it is the last of the classic gang of four that we keep talking about. It is an aggressive cyber actor, but mostly these days it is just trying to steal money from banks to shore up its hard currency position in the face of sanctions. To my knowledge, it does the odd information operation, but I would not put it in the same bracket in terms of systemic risks to the UK.

Lord Hall of Birkenhead: Are there links in disinformation? You were clear about the difference of the Russians versus the Chinese versus others. Are there links between those players and domestic players? Are there solo operators globally that also can be a real force that we should be concerned about?

Ciaran Martin: On the solo operators, it is certainly a potential threat, because unlike a sophisticated cyberattack you can do something relatively on your own. There have been sporadic moments. For example, a few weeks ago there were fake remarks attributed to the former Prime Minister Liz Truss, which she never said. They were widely picked up. It is hard to see where they came from, but most of the circumstantial evidence points to, frankly, a prankster based domestically.

On the first one, can I just ask you to clarify what you mean by links between foreign and domestic?

Lord Hall of Birkenhead: In other words, are there links between foreign states and other individuals in communities working here who are dealing in disinformation?

Ciaran Martin: No, not significantly, to my knowledge. In classic cyberoperations with the criminal underworld, which is headquartered in Russia, you get the odd transaction that is attributable and that is discoverable here, on money laundering and that sort of thing. I do not wish to cast aspersions on whole groups of people, but you will find at least some anecdotal concerns about people with connections to Russia who live in this country talking to some people to get a better feel and that sort of thing. But I would not have detected, either here or in the US, any large-scale formal tasking and organised activity; the US homegrown disinformation problem is entirely self-sustaining. This is one

of the problems in tackling it. All this outside interference digitally can be done without setting foot on the territory of the country that you are targeting. That makes it much lower risk for the people doing it, because they are beyond the reach of our authorities.

Q102 **Lord Dunlop:** You said at the outset that the cost of doing business is falling. Looking at the threat that lies ahead, do you think the balance between malign state actors and lower-level domestic actors is likely to change as we look ahead?

Ciaran Martin: That is possible. It is very hard to say. It depends on a number of factors, including the level of polarisation in the country. While it is a fairly flippant example and was fairly short-lived, the thing I mentioned about former Prime Minister Truss appears to have persuaded a few Twitter account holders with large followings to amplify it. It is nothing more complicated than that. They did not make a video—they did not make an audio. Already it is possible to have some kind of an impact, even though the story ultimately did not get much traction.

One of the things in the tech industry at the moment, which those developing AI are talking a lot about, is that there are freely available products out there that make it relatively easy, particularly in audio, less so in video. We might come on to this, given where you are going with the questioning. Organised groups, if they can be bothered and can do the targeting well enough, will be able to produce this stuff. But amplifying it and doing it in the right way can take quite a lot of skill. On the same weekend as the probably domestic Liz Truss deepfake, we saw a technically sophisticated but entirely unsuccessful Russian deepfake about the terrorist attack in Moscow. It was a really high-quality production of a well-known Ukrainian TV programme, with the Ukrainian national security adviser apparently admitting involvement. It is very hard to say whether it had any domestic impact on Russian opinion, but it obviously did not fool anybody outside, not least because the TV programme makers in question said, "This is nonsense".

I am breaking down the barriers to success for the malicious actor into two. The barriers in terms of producing disinformation are becoming close to nil, other than human effort and having the appropriate technology. The skills, techniques and tradecraft needed to land it as something that will actually persuade people are still relatively high sometimes.

Q103 **Baroness Healy of Primrose Hill:** Professor Martin, you warned in advance about the cybersecurity threat when setting up the centre. I wonder now what impact new technology has had on the development, reach and impact of misinformation and disinformation. I know you are saying that the impact here in the UK is not that great at the moment but, looking at America, how do you feel the future could be for the UK? Particularly, do you have concerns over generative AI, or do you think that is still being hyped too much? I am interested in your view.

Ciaran Martin: It is unwise to talk about the aggressive use of technology without talking about the commensurate steps that can be

done to make it safer and less harmful. The technology itself is neutral; it is maths and engineering—it is neither good nor bad. If you take the tech end of the problem, there are innovations; there are start-ups all over the place. I am not involved with any of them, so I am not selling anything, but there is a spinoff of MIT called PhotoGuard. You use it on images of high-profile people to make the pixelation harder so, if you doctor it, it looks ridiculous. People are doing similar stuff for audio.

Additionally, like everyone, I am sure, on the committee, there are many reasons for me to be critical of the big tech platforms, but some interesting things have been happening there within the last month or so. I am in no way involved in this one way or the other, but Meta has announced an intention to introduce a function on its platform that will allow you to label your posts as generated by AI. If you do not, but it detects that it was, it will do it anyway. There are a bunch of things happening, which are natural market and social evolutions, that are trying to take care of the problem.

In terms of the future threat, there is one issue we have not discussed yet in all this. I am very taken with some comments that Sir Alex Younger made in 2020 in his valedictory interview with the *Financial Times*, where he said it is very hard to divide societies that do not want to be divided. With the way this lands, it is very hard to talk about it just as a technology issue.

I know this is about disinformation, so I am sorry that I keep retreating into my core territory of cybersecurity, but the question of what media outlets do with either fakes or stolen information is really important. It is nothing really to do with the tech. There was a lot of self-examination and soul-searching in the United States after 2016 about whether it was right for the *New York Times*, the *Washington Post* and the other big titles in the US just to take this stuff and make a bunch of stories out of it, rather than realise that it was the prospect of foreign interference.

On the other hand, something the UK has done extremely well is the way in which, when we had the audio deepfakes—I have compared them and I would concur with this—relating to two senior Opposition politicians, the Leader of the Opposition and the Mayor of London, the Security Minister and other members of the Government came straight out and said, "This is a fake. Don't cover it. Don't amplify it". That is much more powerful. Again, it is nothing to do with the technology.

I am sorry to keep sounding like I am dodging the technological questions, but, at the risk of sounding complacently reassuring, I do not know of any tech developments powerful enough to convince us all that some lie was true that would circumvent a determined effort by a society that wanted to try to cohere and behave responsibly in the face of a malicious threat. Of course, societies do not always work like that. The bitter polarisation in America meant that it handled 2016 very differently, and there were a bunch of accidental factors in Slovakia in the election last year. It is the closest European experience we have seen to a fake intervention actually shifting the dial. It is impossible to know whether it delivered a victory for the SMER over Progressive Slovakia. I do not

know. You have to ask every voter why they finally made up their mind, but it does seem to have had an impact. That was due to a whole bunch of technical and accidental factors.

We have to look at this from a whole-of-society, social and political media, as well as technical, standpoint. To my mind, looking at all those factors together, if we get it right, there is no reason to assume that we will be overwhelmed by highly sophisticated disinformation.

Baroness Healy of Primrose Hill: Thank you, that is reassuring. Do you have any concerns about public trust in the integrity of the upcoming elections? Do you think that public trust is being chipped away at by all these ideas about what technology can do?

Ciaran Martin: I have three quick points. First, trust in elections and indeed the political process is a hugely complicated thing. It is about everything from the integrity of the printed ballot papers—if somebody were to do a sophisticated cyber intrusion and just change them or stop them being printed in a particular area. I know that we cannot hack the count or the way people vote, because those are manual, but you can do all sorts of things to stop people voting or have an impact on integrity. You can interfere with the broadcast of election results.

I know this sounds like a really trivial example, but I mention it for a reason. Last year, we hosted Eurovision, a big set piece event important for the country's reputation. The planning for that, in terms of looking at the different threats, was really interesting. It was very complicated, because there were so many points at which you could interrupt the integrity of the event. In 2022, when Italy hosted it and Ukraine won it, the Russians appear to have tried to hack the voting in Italy—just one out of three dozen countries voting. The show went on, everybody could see it, but if you cast doubt on one part of it, you can cast doubt on all of it.

If you take something as complicated as an election, with 400 micro events, or thousands if you take it by local authority area, and lots of different newspapers reporting exit polls, there is a lot of stuff you can do to disrupt that. Since Mrs May back in 2017 in the snap election, we introduced a programme to try to protect as much of that as we could. The template for that is Protect2020 in the United States, where they mapped out the entirety of the way a sophisticated election process worked and protected it as best they could. So there are risks.

On the second point I would make, I feel quite passionately about this, even if it may be surprising to you that I am so cautious both about overhyping the threat and about aggressive interventions. To go back to Alex Younger's interview, he said, "Don't do the Russians' job for them by bigging up the threat". If we talk a lot about this and cast aspersions, we will harm ourselves. There is one specific example from my own professional experience in government, which I felt quite strongly about. I am not passing any comment on Cambridge Analytica or any of those allegations around the time of the Brexit referendum. But the one issue I know very well, which you will all remember, was the collapse of the register to vote system 48 hours before the closure of registration to

vote for the EU referendum in June 2016. Sadly, this does not reflect well on the state apparatus, of which I was part, but when we had a proper look at that we were, unusually for such matters, certain as to the cause. I choose the word "certain" carefully.

The cause was a defective system that could not cope with greater than expected demand to register to vote ahead of that referendum. There was a real surge. It is now fixed, and I am not trying to cause alarm, but ahead of 2016 it was a weak IT system that could not cope with an unexpected surge right on the deadline, so it collapsed. For those of you who know what a distributed denial of service attack is—a DDoS—it was a DDoS by the British public accidentally. That was it. We were certain of it.

Again, I am not commenting on some of the other allegations, but there is an extant report from the other Chamber, by the Public Administration and Constitutional Affairs Committee, saying that, because of the prevailing atmosphere at the time, foreign interference in a specific aspect cannot be ruled out. I think that that was a mistake, if I am honest. We absolutely know what happened: you could trace it forensically and digitally. We would have been quite happy to prosecute it. In my professional judgment, it would have passed a criminal standard of proof in a court of law. Still, there is this hangover: "A digital system went down around an election. It must have been the Russians". If you let that take root, in the face of overwhelming technical evidence, you are inviting people to lose trust in the electoral system, in my judgment.

Q104 **The Chair:** I know that Lord Knight wants to ask about technology before we move on to the next question, but can I follow up on what you have just said? What do you suggest should have been done to make sure that it was then understood that this was a technical defect and not to allow a suggestion down the line that it was due to some kind of foreign interference?

Ciaran Martin: It is hard to say. Perhaps—hands up in my own organisation—we should have been much more assertive and public in terms of the reassurance than we were at the time. We gave evidence but, clearly, it did not convince. I suppose then it is a judgment call for Ministers, supported by their Civil Service organisational heads, to decide how much they want to take that on.

To go back to the US experience, the US has all sorts of problems in this space, so I am not using it as a model, but it has done some very good work on protection. In the 2020 campaign, it had this function called Rumor Control, a rapid response unit set up specifically for elections. When people started making allegations about electoral malpractice that got pickup, it had a pretty large team that would publish, very rapidly, the evidence to the contrary. There is something about being able to contest that when it gets pickup. Of course, there is a judgment. You do not amplify every allegation—but for something as serious as that we ought to have been a bit clearer.

If I may say so—and I am walking into controversial territory without being led there on this—that issue is relevant to the Government's recent announcement on Chinese activity vis-à-vis the electoral system. Most security experts would think that the Electoral Commission, the register operation, was a data-gathering exercise. That is a different problem. It is a very serious problem, even if I suspect this was a rather botched operation by China because, as we all know, it is relatively easy to acquire this data. It did not need to go to that trouble. China is gathering as many population-level datasets as it can for the age of highly advanced computing, including quantum computing, whereby it can build up pictures of the populations of all the states. But it is not taking the electoral register so that it can interfere with elections, in my judgment.

Indeed, on the other aspect of that announcement about intrusions on to Members of Parliament, I am afraid nation state espionage on Members of Parliament, while we do not approve of it, is something that we understand and have always understood goes on all the time. Unless it is accompanied by leaking of information to subvert the political process, it is not disapproved of by international law or custom and practice.

Chair: I will resist the temptation to go down that avenue.

Q105 **Lord Knight of Weymouth:** I just wanted to delve a little more, particularly with the Meta example that you mentioned. With the rapid development of generative AI and its increased power, one would suppose that the ability to create credible content will become easier and cheaper. It is the content amplification by algorithm or broadcasters that then becomes more concerning, just because there will be a lot of stuff out there—and will they be able to detect whether it is fake?

In the area of education, where I do a lot of work, there is a lot of discussion about cheating. The tools that are there to spot fake or AI-generated content are not as powerful as the tools that are used to create the content in the first place, because you are always playing catch-up. Even though Meta is so wealthy and has so much resource, can we really trust the big tech companies to be able to keep up with the innovation in generative AI at the detection end of the process?

Ciaran Martin: I see the problem. Obviously, generative AI can create all sorts of content. As somebody else in the education sector, I think that one of the biggest problems of misuse is in tertiary education.

On the ability to keep up, in all this use and misuse of technology there is an element of cat and mouse. The thing is, whether it is cybersecurity or the ability to tell whether something has been doctored and so forth, because of the way in which the technology works, you tend to be able to build up a defensive capability alongside. It is rare and, in many ways, contrary to mathematics to say, "You can use this aggressively, but you can't do an equal and opposite". It is a question of who has it.

This would lend itself to the more sophisticated actors. Could the people who can do it at no cost do something that would outrun Meta's defensive capabilities? It is not impossible, but it is less likely. Could an organised unit, sponsored by a wealthy nation state with lots of time,

effort and prioritisation, do something that would be beyond Meta? That is plausible.

I know I am perhaps putting a little bit of cold water on some of the threats, and we need to be vigilant—but there are two broad problems. One is relatively cheaply mass-generated nonsense that people find a way of believing, for whatever reason, because we misfire as a society, we do not handle it properly, we are polarised anyway and so forth. The other is the sophisticated campaign.

To go back to a previous generation of technology, the one quite concerning episode in recent British political history, which was covered and was officially avowed by the Government last year, relates to the politics of 2019, which culminated in the election. That was the leaked letter from the Department for International Trade on potential UK-US trade negotiations leading to an opening up of the National Health Service to greater private involvement, including from the US. There was a very good report by researchers in New York from an organisation called Graphika. This was not election-related initially, because the operation started in January 2019. Of course, the election was in December and nobody in January 2019 was anticipating an election at the end of that year. It was a hack from around the Secretary of State's office of a legitimate letter. What then happened is interesting as to whether it was genuine or fake. The really sophisticated and patient way in which they amplified it is charted by Graphika. It ends up getting into credible enough outlets that somebody in the then Leader of the Opposition's office decides that it is worthy of campaigning. You end up with, in the middle of the election campaign in November 2019, the Leader of the Opposition reading it out on an election campaign podium. That is a concern.

We see that type of operation, where you have a sophisticated deepfake. You do not do something stupid, such as making a fake version of a legitimate Ukrainian TV programme, as in the exchange that I had with Lord Dunlop, because that is so easy to debunk. You just get the TV programme to say, "We didn't do this". You do something else. That is the area we do need to watch out for.

Q106 **Lord McNally:** Just listening to your evidence, it is the centenary of the Zinoviev letter.

Ciaran Martin: It is indeed, yes.

Lord McNally: Low-tech can also have its impact.

Ciaran Martin: I am very glad you raised that, because it shows that this is not new. A lot of it is about how we respond to it. I am sorry to interrupt.

Lord McNally: No, it is reassuring to hear that you agree. You have given a reassuring case that this is not out of control. We are on the case. I wonder whether, in dealing with these malicious actors, there are any weaknesses that you would still identify in our approach. I suppose I am asking in terms of government and government agencies rather than

the public. I would also ask whether the tech firms and the media in general are putting their shoulder to the wheel in this regard.

Ciaran Martin: There are things, but I have a few quick caveats. First, I would be in the area of incremental improvements. Not in this space, but there are other areas of the government approach to tech security where I have been publicly highly critical. I am not just someone who goes along with what the Government are doing. But in this area, it is mostly incremental.

Secondly, because of my worries about undermining public confidence, I apply quite a high bar to very aggressive state intervention in this area, particularly in legislation, because of the risks of unintended consequences and so on.

Thirdly, while there is plenty to complain about with the big tech platforms, and they could do more, particularly with regard to X/Twitter and content moderation, it is hard to sustain a charge that they do not care about this. I would not describe myself as a huge fan of Meta and all its practices, but it has put quite a lot of effort into this across the world.

With that in mind, to give some suggestions, I have already mentioned to the Chair the rapid rebuttal capability. I will probably talk a little later about the Defending Democracy Taskforce, with regard to broadening some of the discussions on to a more cross-party basis and keeping the information flows going. We did that for the cyber protection of the 2017 and 2019 elections. I see no reason not to, not least given that the Government's policy is not to capitalise on these incidents. They treat it as an attack on the integrity of the political process, which is right.

As a career bureaucrat, I sometimes hate getting into informal discussions about culture and so forth but, without entering the horrendous territory of media regulation, there is something around a discussion, either within the traditional media industry or between the media industry and government, about how it would work in the event of a large-scale hack and leak of genuine information or a disinformation campaign, where the Government have reason to believe that it is the work of a foreign service but do not wish to interfere with a free press. What are those mechanisms like? We have not tested or sustained them enough, so I would like to see a bit more of that. It is your classic exercising.

I do not know, and I am not sure I particularly want to know the answer to this, how many of the big national newspaper titles or broadcasting organisations have war-gamed a scenario where something just turns up that is an amazing trove of politically sensitive information and it is very hard to ascertain its provenance. What do you do with that? That is a really tricky situation for a free country and a free press. I know there are at least one or two former senior media executives here, so they will have their own views. That sort of discussion and planning is worth having. I am not trying to steer it in a particular direction, and I would be loath to think about state intervention legislatively, but it is a discussion worth having.

One to watch, to Lord Knight's question, is how this "made with AI" label works. Is it good enough? If it works fine and takes care of 90% of the problem, we are okay, but what if it does not or is patchily adopted? X/Twitter is really interesting at the minute, because it is talking about doing this. On the one hand, it has decimated content moderation, so it has got even more unpleasant than it used to be in terms of abuse and so on. On the other hand, for all its many faults, the community notes system has done quite a good job in some respects of knocking down egregious nonsense, including the alleged comments by former Prime Minister Truss, which were knocked down pretty quickly by this mechanism. There is a lot.

Looking at the impact, I know the Online Safety Act did not in the end deal with this issue, and I am not sure that was a mistake—but if these pledges to label, watermark or whatever are adopted and work, that is fine; if not, let us come back and have a look at that issue. Those would be the things I would look at.

Lord McNally: That is very helpful. In particular, as I raised in private discussion, it is not unusual with external threats for Opposition parties to be briefed on privy counsellor terms. I do not see why some of these things cannot be discussed on privy counsellor terms, particularly in an election year.

Are there any lessons from abroad? You mentioned this American rapid reaction or correction force. Are there any other countries that seem to have good things in place that we could look at with profit?

Ciaran Martin: There is good and bad practice. We did mention the Slovak experience, which is interesting for a couple of reasons, in terms of what went wrong. One is that the legal restrictions on reporting of elections, which of course do not exist here but are common in continental Europe, worked against clearing that mess up. Ironically, when the Russians hacked President Macron's campaign in 2017, the opposite happened. The Russians do not always get their tradecraft right. They hacked the En Marche campaign using the same sorts of techniques as with Hillary Clinton's campaign a year earlier. Unfortunately for them, fortunately for western democracies, they dumped it all on the Saturday before voting, without realising there is a 24-hour legal blackout on mainstream media reporting the election. They dumped a trove of embarrassing information about candidate Macron, which nobody could print or cover.

In Slovakia, the opposite happened. Given that most people still engage in some way with mainstream media, the attempts to get the message out that this was a fake audio were unlawful. The other problem in Slovakia, which speaks to potential regulatory issues but just as much to dialogue issues with big tech, is that one reason why it got amplified for so long was a loophole in Meta's policies. It was an accident. Had it been a video, it would have been taken down far earlier, but we were not really used to deepfake audios back then. It had all these policies, saying, "If it's a video, report it and we'll take it down". But it went into

paralysis, because it did not have a policy on audio; it does now, but then it did not. There are those types of unintended consequences.

There are some examples in continental Europe of authorities perhaps doing a bit more of the cross-party work that I talked about. Actually, I am not sure that I have any standout example to point to, I am afraid.

Lord McNally: In your evidence so far, you have not mentioned Ofcom. Does it have any role in the area that we are looking at?

Ciaran Martin: You have to put that to Ofcom. In law, as I understand it, it does not especially. The Online Safety Act, which expanded its role significantly, created an offence. I cannot remember exactly how it is described in the criminal code, but if you wilfully make up a high-impact lie, it is a criminal offence in this country. That is for the police to enforce. There is no commensurate law as per, for example, child safety, whereby Ofcom has to enforce duties on the big tech platforms to have policies safeguarding children.

Since I am here to answer questions, forgive me for posing one, but it is one reason why I am quite cautious about this in totality. Every time I talk about this, I ask, "To which state agency do you want to give effectively what sounds a bit like the regulation of truth?" I remember a very interesting conversation around 2017, because we had the snap election that we were not expecting. It was just after the pretty disastrous experience of the US, so I am not complacent. It was an awful thing that happened in the US in terms of confidence in the political system. We mapped out what we could do to protect the British political system from interference as far as we could. Where we got stuck in terms of the National Cyber Security Centre's role was disinformation. We said, "We're not doing this for this election". "Why not?", we were asked, and we said, "We're a subunit of GCHQ. Are we supposed to validate the information put out by political parties?" There has always been this question: if you are going to have an independent arbiter of information, essentially, where are you going to put it? That needs a lot of careful thought.

One of the reasons why this stuff can be effective is that, if you are the aggressor—and this is the classic Russian playbook—you think, "My objective is to destabilise the United Kingdom, so I'm going to do a whole bunch of things. I'm going to hack its political parties and spread it all over websites. I'm going to make up a bunch of lies and I'm going to do some other things as well". If you just take those first two, then look at it from the British state's defensive point of view, the first is operationally hard but conceptually quite simple. You put in place a bunch of things that try to deter and defend against people hacking your networks. It is always a crime—you can do law enforcement stuff, if they are here. You can strengthen the cybersecurity of political parties and MPs. You can try to do all these sorts of things. They are hard to do, but they are easy to conceive of.

It is not a crime to make up a lie and put it on the internet, most of the time. There is now an offence but, if somebody is doing it in a sophisticated way from another country, what can the Government do?

It tends to be carried by a private sector platform headquartered in the United States, or something you have never heard of, or it is something that is new and you are chasing it. That is a much more difficult problem—that is where you get into the wider social resilience and how we handle this as societies. I apologise for giving such a longwinded answer, but I would be very hesitant about dragging Ofcom into this. Obviously, it can speak for itself.

Lord McNally: As a final thought, one always thinks in military hardware in terms of deterrence. We have lived for 70 years by deterring. Why can we not deter with technology? Do we just sit pat while they do all these things to us, or do we get to a situation where we say, "We can do you some serious damage if you don't stop this"?

Ciaran Martin: There are a few answers to that. There are things we can do. We have struggled in tech deterrence, partly because of the geographic imbalance. When you do not have to send somebody over to inflict harm on a country, it changes the calculation, particularly with regards to Russia, for two reasons. First, the Russian state does not extradite its own people. Secondly, it is unembarrassable. China will not extradite its own people to you, but it is embarrassable. The calling out of China on nefarious digital activity has, from time to time, had some impact. So we have struggled a bit with deterrence—and I know that is a hard message.

Having said all that—and I am not going to go into details and, in any case, my details are years out of date—we have offensive cyber capabilities. I always thought that the degradation of the digital infrastructure of disinformation outfits in adversarial territories such as Russia is a good use of those. It is widely assumed—I do not believe the American Government have ever officially confirmed it, but there have been books about it and extensive reporting—that, after what happened in 2016 in the United States, ahead of the 2018 midterms there was a prolific troll farm churning out all sorts of quite impactful disinformation, not just in the US but around elections in the Balkans and so on. It was known as the Internet Research Agency, based out of St Petersburg. Unlike a Russian cyber intelligence agency, which would have quite good tradecraft, sealed buildings and hard infrastructure, the Internet Research Agency was a troll farm in St Petersburg that was internet connected, just making up lies and putting them out there. The US cybercommand essentially rendered it useless—so there are things that can be done there.

There is the playbook that we use in cyber interventions of sanctions, et cetera, so there are things that can be done. Much as with cyber criminals, it is tactical, but tactical interventions matter. The more we can use interventions to take down the disinformation infrastructure of these groups, and the costlier we make it for them to operate, the better it will be.

Q107 **Lord Kamall:** Thank you very much, Professor Martin. I was struck by two of the examples you gave. I want to be careful about my words, so tell me if I am misquoting you. There was the one at the beginning about

the US election, where there was information that was not disinformation, but it was hacked and leaked for political advantage.

Ciaran Martin: That is correct.

Lord Kamall: Then the other one was the voter registration system going down for technical reasons, lack of capacity et cetera, but then the committee in the other House said it could not rule out foreign interference. Given that, how can we distinguish between disinformation and difference of opinion?

Ciaran Martin: Gosh, I am not sure how to answer that, Lord Kamall.

Lord Kamall: Do you want me to phrase it a different way?

Ciaran Martin: Yes, thank you.

Lord Kamall: To help you, I will try a couple of angles and you choose which angle works best. There are some suggestions that misinformation or disinformation is overblown. Would you agree with that? If you focus too much on disinformation, does it risk being politicised?

Ciaran Martin: It comes back to that distinction between intent and activity on one hand and impact on the other. There is an awful lot of this, a huge amount. It ranges from organised nation state efforts all the way through to bored pranksters and everything in between. In countering it, to go back to quoting Sir Alex Younger, we have to be hugely careful not to big up its impact, if that impact is very limited or negligible. It is hard to point to any wholesale distortion of the integrity of UK political life by disinformation as of yet.

That is not to say, as per the exchange with Lord Knight, that there could not be a very sophisticated campaign. That is not to say that we might not end up falling behind and having all sorts of people put out all sorts of nonsense that is believed. But it is not neutral to sound a very loud alarm on this without the evidence base to do so. That is not a neutral act, in my view, in terms of confidence in the political system.

Lord Kamall: Given that, do you think there is a danger in undermining trust in so-called news because of this issue? If you think about the diversity of opinion we have in our newspapers, for example, you can go from the *Mirror* and the *Guardian* at one end to the *Daily Mail* and the *Telegraph* at the other. Sometimes they will report the same story but from a very different angle, and sometimes they accuse each other of disinformation. When there is that accusation, does that undermine trust in news?

Ciaran Martin: The bigger issue here is outside my field of expertise. Scholars I respect are still quite divided on this, and the evidence base is pretty patchy, but a much more powerful place to look than lies, which is essentially what disinformation is, is the echo chamber polarisation in the age of the internet. It is 23 years since Cass Sunstein wrote his paper about how, because of the way the algorithms would work, we would look only at the news that we agreed with. There are plenty of scholars who take the view that that has not actually happened. It is not my area

of expertise, although I would veer towards the Sunstein view that that is happening.

In terms of trust and the integrity of news, the lack of exposure of many people to differing viewpoints is a much more significant challenge, than the carrying of outright lies. Please correct me or ignore this if you think it is wrong or nonsensical, but the carrying of outright lies in mainstream British broadcast and print is very rare and is generally caught.

Lord Kamall: Going back to a question that one of my colleagues may have asked, is there a role for government and what is that role?

Ciaran Martin: There is, but it should be limited and careful. First, there is a role for government, particularly at the sharper, more progressive—call it what you like—intelligence and security end, in gathering information on what potential adversaries are up to in this space, whether that is hack and leak or disinformation and, where possible, intervening on that. The exchange with Lord McNally is a very important one. That is something that only the Government can do.

There is another thing only the Government can do, although I do not suggest they do it now, because I do not think there is this outstanding gap that is there in other areas of our cyber defences. Only the Government can, through Parliament, change the law to require something different to be done. Others will have different views, but at this stage, in this particular space, particularly now that the offence has been created of wilfully creating a lie, I would not advocate that. That is something for government.

There is then a harder thing for government to do, which we have already talked about. One thing they have done very well, through the Defending Democracy Taskforce, but mostly through the way they have gone out there, is knocking down and discrediting particular pieces of disinformation. The good examples of that are Minister Tugendhat's two interventions on Sir Keir Starmer and Mayor Sadiq Khan. You could go further with that and systematise it, as the Americans have done. As I have already suggested, there is no reason why you could not have a broader and more inclusive cross-party dialogue on that. I am aware of the very fraught history of government-media relations in this century after Leveson and things like that, so I am not suggesting aggressive regulation, but there is something around the dialogue with media outlets, sharing information on the threat and talking about policies.

If you will forgive me, there is a really good example in a different field from Australia, in two very different areas. One is teen suicides, going back a quarter of a century, and the other is cyberattacks from last year. Some 25 years ago there was a spate of teen suicides in Australia. The experts all said, "They're all copycats, and it is because it has been amplified in the media". The Government and media outlets had a discussion about responsible reporting.

Some 25 years on, Russian criminals hack 37% of the Australian population's medical records through a medical insurer called Medibank. There is a huge panic in Australia. This is October and November 2022.

The Government use that playbook. They get all the editors and broadcasters in, and they say, "We're not going to censor. We're not going to invoke any emergency powers. We're not going to pass any laws, but there is about to be a data dump of 37% of the population's medical records on the dark web. It's obviously legitimate to report that. Report the fear and the failures. Criticise the Government all you want. Just don't report individual medical details. Just don't do it". Nobody did. Twitter and Facebook took them down. Australia managed not to pay—the company did not pay the $10 million in ransom, which would have encouraged further attacks. There has been no discernible harm to Australian citizens as a result of this, because the media basically suppressed the information that was out there. It was hard to get; it was on the so-called dark web. There were no legal interventions at all, no threats of new laws and no censorship. It was a dialogue.

As a career bureaucrat, that makes me uneasy, because what do you do? Do we have that system here? I am not trying to say, "Let's create a D-notice committee for this type of thing", but there is a dialogue there to be had before it happens. What happens if there is a big disinformation campaign that is sophisticated? What happens if someone does a hack and leak? How do media outlets balance their duty to report legitimate information, whether that is the content of it, if it is real, or the fact of the disinformation campaign, with their desire, presumably, not to be unwitting agents of an adversarial hostile state?

Chair: We are running over time. Was there anything else?

Lord Kamall: I just have one more thing.

Chair: Can I ask you to be brief in response?

Ciaran Martin: Yes, of course.

Lord Kamall: If it is a longer answer, please write to us. That is okay with us. I would be interested in your view. I asked you what government should do, and I know you have touched on this, but what should government not do? Maybe you would prefer to write to us about that rather than tell us now.

Ciaran Martin: I will.

Q108 **Baroness Harding of Winscombe:** Thank you for your very thoughtful and thought-provoking evidence, which has all been very calm and in keeping with your several decades of Civil Service background. I should declare that Professor Martin and I have worked together, so it is very consistent with my prior experience.

This is a topic that often does not get discussed in these terms; it quickly becomes very politicised. I am very struck by what you just said: that it is not neutral to sound the alarm. I just wanted to get your sense of how we, as politicians, can find a way of talking about the issue and discussing it sensibly without immediately going down the blind alley of politicisation, and this fear that you are distorting the political dialogue just by raising the issue.

The Chair: To give another take on that, which has already been floated, in raising the issue of disinformation, it can seem to some people like there is an opinion that is being shut down. It is that, really. Is there anything from your professional perspective that you would offer to avoid that? In fact, is there a role for the security services in this, too? I do not mean a formal role, but has there been any self-reflection within the security services about the way this whole matter has developed over the last few years?

Ciaran Martin: I will do my best, thank you. In terms of politicians and the political discourse, we are not starting from the beginning. We have had several years of this era already. There is one core strength of the UK system that is worth keeping, which is not there in some but not all other countries, as we have seen—a sense among competing political organisations and parties that this is an attack on the system rather than on a particular wing of it. Hold on to that, is the first thing I would say.

Secondly, politicians, the security services and the government machine as a whole should talk about this. Given your question, Chair, I am not trying to shut down discussion of disinformation.

The Chair: No, what I mean is people using it as a weapon in debate.

Ciaran Martin: No, absolutely. In terms of Baroness Harding's question, we should talk about it and raise awareness of it. This is an area where, ultimately, the judgment of millions of individual citizens is what is going to determine success or failure. We need a public discourse that says, "There is no evidence of a crisis in the integrity of elections, but there are plenty of people trying to purvey nonsense, whether for criminal reasons, for propaganda reasons, for political reasons or whatever it is". We are in an age where we have to train people to make judgments and apply scepticism to information that they see. This is partly about civic education in schools; it is partly about society. That is a decent way to talk about it.

I get this boundary, Baroness, that you are talking about between something that we say is disinformation or misinformation and an unpopular opinion. That is a very tricky one that I probably want to reflect on a little further, if I am honest. Perhaps I will leave it there.

The Chair: That is okay. If you want to reflect on that and the question Lord Kamall asked, which you promised to follow up in writing, and come back to us with both of those, that would be tremendously helpful. We have already taken up more of your time than we should have—but it has been hugely helpful to have you here, Professor Martin. Thank you very much for your evidence.