HOUSE OF COMMONS

Foreign Affairs Committee

Oral evidence: Disinformation diplomacy, HC 703

Tuesday 11 February 2025

Ordered by the House of Commons to be published on 11 February 2025.

Watch the meeting

Members present: Sir John Whittingdale (Chair); Alex Ballinger; Aphra Brandreth; Richard Foord; Uma Kumaran; Blair McDougall; and Abtisam Mohamed.

Questions 1 to 42

Witnesses

I: Professor Vera Tolz-Zilitinkevic, Sir William Mather Professor of Russian Studies at University of Manchester; Dr Jon Roozenbeek, Lecturer in Psychology and Security at King's College London; and Professor Martin Innes, Director of the Security, Crime and Intelligence Innovation Institute, Police Science Institute and Professor in the School of Social Sciences at Cardiff University.

Examination of witnesses

Witnesses: Professor Tolz-Zilitinkevic, Dr Roozenbeek and Professor Innes.

Chair: Good afternoon. This is the first session in the Committee's new inquiry into disinformation and diplomacy. It will act as a sort of primer, with our witnesses educating us on the details of the area we are to examine.

I should begin by conveying the apologies of the Chair, who sadly has had an accident that temporarily means she is unable to join us, but hopefully she will be back in her place quite soon.

Thank you very much for coming along. Could each of you introduce yourselves for the record?

Professor Innes: Thank you for the opportunity to address the Committee. I am Martin Innes. I am a professor at Cardiff University, where I lead a large-scale research programme into information operations.

Dr Roozenbeek: It is a pleasure to be here. My name is Jon Roozenbeek. I am a lecturer in psychology and security at King's College London's department of war studies. I work a lot on mis and disinformation, why people believe it and what you can do about it, with a particular focus on Russia.

Professor Tolz-Zilitinkevic: Thank you for inviting me, Chair. I am Vera Tolz. I am a professor of Russian studies at the University of Manchester, and I co-lead a large project on the history and current practices of disinformation from the first world war to the present, looking at seven different linguacultural environments.

Q1 **Chair:** Perhaps we can start by making sure we are all talking about the same thing. The Committee has discussed disinformation in several other meetings, and different terms have been used. Some people have talked about disinformation, some about foreign information manipulation and interference, and then there is misinformation. Could you briefly outline what you see as the definition in each case, and whether it is being used consistently across Government and other bodies?

Professor Tolz-Zilitinkevic: The term is a minefield—the whole range of terms is a minefield. They are not consistently used, and probably cannot be consistently used, because they are historical terms going back to the 19th century.

For the purpose of our discussion, disinformation is deliberately falsified information, whereas misinformation may not be deliberately produced. I can explain later why I think talking about foreign disinformation probably is not the right perspective, because we are dealing with a global environment where various dubious information providers collaborate

across borders. That is absolutely essential for understanding how the processes work.

If we are talking about information manipulation campaigns sponsored by particular Governments or actors, I think the term disinformation can be applied well, analytically and constructively, because the actors who organise those campaigns try to conceal the origins of the messaging; that, in itself, is a form of deceit.

Q2 **Chair:** Would you say there is a consensus on each of these terms and their distinctions?

Professor Tolz-Zilitinkevic: No.

Q3 **Chair:** Where is the confusion?

Professor Tolz-Zilitinkevic: I get that academics try to complicate the simplest things, but to me, the key issue is whether you reduce disinformation to something that is demonstrably and factually false. To me, this is too narrow a definition.

We will talk about people's reactions to dubious information: usually people buy a particular news report because of the whole nature of the story and framing of it might seem plausible; then, if the framing is very questionable, whether individual facts are correct or not becomes less important. For the purpose of our discussion, the key issue is whether or not we reduce disinformation to factual falsehood.

Q4 **Chair:** Professor Innes, there is a lot of talk now about disinformation. How serious a threat is it? How much is it taking place, both in the UK and globally?

Professor Innes: To follow on from the previous comment, there are three other terms of art at play that are worth mentioning. There is the FIMI concept—foreign information manipulation and interference—co-ordinated inauthentic behaviour, and then information operations. None of those terms is perfect and each captures something different. What you think the threat is depends on which of those you select, because it frames the threat we have.

My preferred term to focus on is information operations, or information influence operations, because that gives you a greater sense of analytic precision. To illustrate that, rather than just talking about FIMI, if we look at the conflict in Ukraine, there are probably 12 separate Russian information influence operations continually at work trying to influence public perceptions and understanding of the conflict in Ukraine. That matters because each of them has their own signature methodology, they operate across different platforms and try to do different things.

Some of those information operations are focused upon the military conflict and trying to influence public understanding of it; others are trying to overload fact checking or spoofing media sites. That is where we need to get to in terms of understanding what the threat is. The threat is multi-faceted, and it is evolving.

Then we move on to who is behind the threat. There are two good sources of information that take a global perspective: the data coming out of Princeton University in America and the reporting coming out of Meta's threat analysis. Both of those suggest that the majority of information influence operations that have been detected have Russian origins, and that the next most common threat actor is Iran and then probably China.

Q5 **Chair:** You have slightly added to my confusion, rather than reduced it. You used three different descriptions, the first of which was FIMI. What were the other two?

Professor Innes: Co-ordinated inauthentic behaviour, which is the preferred term of social media platforms. It is good for the Committee to understand that, because they are important to the work that goes on. The reason they prefer the term co-ordinated inauthentic behaviour is that, by focusing upon the online behaviour, deceptions and co-ordination that take place online, they argue that they do not have to get into policing freedom of speech and freedom of expression. But that curtails a certain amount of what they can do.

The final version was information influence operations, which is a slightly more focused concept. That is about trying to key into campaigns that have specific, defined objectives and use specific assets that we can attribute to particular aspects of the Russian or Chinese state.

Q6 **Blair McDougall:** Professor Tolz, I am interested in what you said, about the clandestine nature of the troubling information coming out of Russia, China or Iran, versus that coming from other actors, which may be more obvious.

This inquiry is very much about actors, rather than just states, so how do you see the development of disinformation in terms of the risk from individuals—we have seen Elon Musk taking a particular interest in the UK and spreading particular information here—versus the threat from states? You are nodding, Dr Roozenbeek.

Dr Roozenbeek: The distinction between state operatives and individual operatives is often very muddled, in the sense that, as Professor Tolz correctly pointed out, these actors very much try to hide their tracks. Establishing that they are linked to a given state, or a given set of interests, let us say, is often very difficult, because they deliberately try to conceal that. When you find an individual who seems to be personally, individually interested in spreading disinformation, or to be engaging in these kinds of influence operations, you do not have a lot of evidence to go on as to whether there are any links with state actors and so on.

In the case of Musk, we can only really look at what he does in terms of the risks that that might pose to the United Kingdom. Those risks translate practically to questions like, is he amplifying certain content over other content on his own platform, X, algorithmically speaking? There is some evidence of that, but it is not really clear, in the sense that him simply retweeting someone gives them a lot of attention; he does not need to change the algorithm for that at all. In that sense, in terms of evidence of

the deliberate manipulation of the information environment for political purposes—to act against the interests of the United Kingdom—that is too much to say out loud; the evidence is not that clear. But we are also thinking about threats in the sense that we do not necessarily need to have 100% certainty on the evidence in order to consider something a threat.

Q7 **Blair McDougall:** Even if Musk did not own and control the algorithm and manipulate it in that way, how does he compare, just as an incredibly powerful individual, to the threats from states, in terms of the worry about this? We have heard that the security services here are monitoring his output in a way they would normally monitor stuff coming from hostile states.

Dr Roozenbeek: I find it difficult to comment on that in detail. I think that the difference at the moment is that there is, on the part of Russia, China and Iran, a singular focus and a set of very strategic goals that they are pursuing through these means. With Musk, it is possible that it is more opportunistic, but the consequences might none the less be the same, if that makes sense.

Professor Tolz-Zilitinkevic: It is very important to say that when we talk about risks to states from individual actors or these organised influence campaigns, there is a huge distinction to be made. To take operations from Russia specifically, which is the domain of my research, in terms of established Western democracies such as the UK or the US, internally the influence on the population and the threat are not great. However, the influence of Russian campaigns in the rest of the world, including in areas of strategic interest to the UK, for instance, could be great.

The influence of Russian disinformation operations in Africa is significant, may be well documented, and has an impact on people's perceptions. The influence of the Russian state and intelligence services actors loosely affiliated with the Kremlin, according to my research, is much lower in the UK or the US than of actors who are internal to democracy. I think that Elon Musk's influence is potentially greater in the UK than that of Russia's.

I specifically researched the influence of Russian disinformation actors or proxies—say, RT or Sputnik—in the US during the 2020 presidential elections. You could see that the uptake of their communications on social media was incomparably lower than that of domestic American disinformation providers, such as Project Veritas, with which, for instance, RT collaborated. The following of Project Veritas, with the disinformation around elections and "stolen" elections, was incomparably higher.

Q8 **Aphra Brandreth:** I want to pick up on a couple of things that Professor Innes and Dr Roozenbeek were talking about before, to do with social media and some of the information—which could be disinformation or misinformation—that gets put on there. Do you have thoughts about how we enable free speech while addressing misinformation?

Dr Roozenbeek: That is a very good question. I think Governments are rightfully very wary of interfering in speech, especially political speech, and that is a good thing, overall. In a sense, what we are concerned about is not necessarily what someone is saying, but the authenticity of what is being said and how it is being amplified.

Where I think Governments are safe in considering some kind of intervention is not in the speech itself, but how it is produced and spread; to say, "You can say anything you like, but if you use a bot campaign to make it more popular, we can say something about that," with an objectivity that does not touch upon interfering with someone's free speech, to a certain degree.

Similarly, this debate is almost never spoken about in terms of market economics: there is a market for online manipulation, and it is wide open and thriving. You can buy fake likes, comments and accounts and so on, on the open market—there is no ban on it. In the UK, this is very cheap. You need a SIM card to verify a fake account, and in the UK, SIM cards cost pennies; in that sense, the price of manipulation in the UK is very low. We have an ongoing project, which I am very happy to present once it is out for public viewing, and the initial data shows that this is incredibly cheap and easy to do in the United Kingdom.

Speaking about it in those terms also skirts the issue of having to talk about what you are allowed to say in the United Kingdom—what is allowed political speech. We are all, rightly, very uncomfortable with that, and I think it should stay that way.

Q9 **Uma Kumaran:** Professor Tolz, I was reading some of your work and multiple things stood out to me, but two of the starkest were when you wrote that Russia continues "to perceive the media environment as a 'strategic tool'" and that we have witnessed factually falsified content on a scale unprecedented even by the standards of the past decade.

We have been looking at the disinformation campaign as a Committee, and one of the really alarming things we have uncovered is that where the BBC World Service has shut down its analogue broadcasts, RT and Sputnik have taken its place. They are pumping out propaganda, if you like, on these airwaves that used to be the BBC World Service. I have asked a number of our witnesses, "Should Britain in the future be taking into account these things and their impact on our own democracy?" Looking at the content that they are putting out, how concerned do you think we should be? On one hand, you could say, "This is a radio station that serves a region in the Middle East. Why should we worry?" What do you think the impact could be on us here?

Professor Tolz-Zilitinkevic: Western states, including the UK, obviously want to engage with foreign countries, whether it is those in the Arab-speaking world or others, and to have these countries on their side. There is a strategic competition between Russia and the UK in relation to Asia, the Middle East and Africa. In that sense, it is a concern, and RT, which

you mentioned, is doing much better at engaging and probably influencing audiences in those regions than it did in the UK, even before the ban.

My latest research on audience engagement compared RT with the BBC and other western broadcasters. In the Arab-speaking world, RT was ahead of the BBC and CNN in engagement, and it was below only Al Jazeera. RT was doing very well with its influence there, and it has also been doing well in Latin America. It is now increasing its market in Africa because that did not used to be its target, but now it is after the bans in the west.

All this comes with a caveat, in that what we looked at was social media engagement. As RT was set up as a broadcaster later than the BBC and other major broadcasters, it started using social media platforms quicker and more effectively than western broadcasters. That is the reason. RT knew how to use this medium from the beginning and, in a way, broadcasting was more for symbolic significance. Most of its resources have been put into social media communication.

We used analysis of social media to follow the types of engagement, and RT, on paper, is doing well in these parts of the world. The only issue is that Russia is falsifying lots of things at the moment on a large scale, so I asked members of my team to check the percentage of genuine accounts that follow RT.

There are particular tools that cannot tell for sure whether individual accounts are genuine, but they can give a ballpark figure of the number of accounts following an outlet that behave suspiciously. If you look at the BBC following, the number of suspicious accounts that are likely to be bots is just 1.5% to 2%, so it is very low. In relation to RT, in our analysis, 39% of accounts were potentially bots created by Russia itself. This, of course, throws all the figures on following and engagement into question and doubt.

Q10 **Uma Kumaran:** Thank you; that is fascinating. I have a follow-up question for Professor Innes. How common are disinformation campaigns like this, whether it is through bots, social media or more mainstream broadcast channels? We are looking just at Russia, but globally, is it a common occurrence now?

Professor Innes: Yes. I think the best data we have tracks between 2011 and 2023 and suggests that over 70 countries have been targeted. There are multiple countries to which engagement in these kinds of activities has been attributed. As I said, the most commonly identified one is Russia, followed by Iran, followed by China. We are in a situation now where we are living with disinformation. We are not going to eradicate it. We are not going to get rid of it. That is the kind of perspective we need to bring to bear.

Q11 **Uma Kumaran:** Is it mostly in the English language? We are obviously focusing on RT putting out things to perhaps a western audience, but is it in other languages as well?

Professor Innes: It is in pretty much every language that you can think of. Fairly recently, we detected some Russian operations seemingly experimenting with Welsh and Gaelic, which are obviously quite small, niche languages on a global scale. But that is increasingly the direction in which they are heading, because one of the things that social media affords is an opportunity to engage in micro-targeting. Whereas previously propaganda was a fairly mass population kind of effort, you are able to use these tools to target different ethnic groups, different regions and different languages, and that sets up a range of possibilities for you in terms of being able to exhibit malign influence.

Dr Roozenbeek: Now that we have generative AI, it has become much easier to write in any language, pretty much. It is still not perfectly accurate, but it is much better than the Google Translate stuff that we had a couple of years ago.

Q12 **Uma Kumaran:** You mentioned the desire not to stifle free speech and the fact that Governments are wary of that, but in the light of this, and especially in the light of what you are saying about AI, do you think Governments can go some way towards stemming the flow of this?

Dr Roozenbeek: Yes. There are two productive pathways forward. The first is that we need to get better at detecting inauthentic content, however it is created, and that is an extremely difficult question. There are teams of researchers dedicated pretty much only to bot detection, and they are doing pretty well, but they are not perfect at it by any stretch of the imagination. There is a lot that can be done in that regard, but it is a very difficult technological problem that requires substantial amounts of research funding to figure out, wherever that funding ends up going.

The second is the point I mentioned before: there is a market that underlies all this. There are people who sell fake accounts and fake likes and so on. You can just buy this in a particular place. We have almost no view on this market, what it is doing, who is operating in it, who is buying, who is selling and so on. *[Interruption.]*

Q13 **Chair:** I am afraid that sound represents a Division in the House of Commons, so I am going to have to suspend the sitting.

Sitting suspended for a Division in the House.

On resuming—

Chair: Apologies for that interruption. We will recommence. As I was saying, I am afraid the House of Commons sometimes causes these interruptions. Uma, I think you were just finishing your questioning.

Q14 **Uma Kumaran:** I was. Thank you for bearing with us. You were explaining, Dr Roozenbeek, the two pillars of fighting disinformation. You said that one was better detection of disinformation spread through the bots, and the other was detecting the people who sell those fake accounts. Could you pick up on the fake accounts?

Dr Roozenbeek: Very briefly, with respect to detection, Professor Innes has already said that Meta is quite good. It has a threat analysis centre, basically, as does Microsoft—I believe it is called the Microsoft Threat Analysis Center. They do excellent work. They keep close track of—

Q15 **Uma Kumaran:** Have they survived the cuts? With what is going on in the US, and with them getting rid of the fact-checkers, have these units survived?

Dr Roozenbeek: In Meta, I am not sure, but in Microsoft, yes. Microsoft tends to fly under the radar in this domain. There are excellent people working there who are very technical. OpenAI also has a team who work on this, because it is quite concerned about its own tools—ChatGPT and so on—being used for the purposes of manipulation. There is a lot of capacity in that regard that is worth talking to, I would say.

The second aspect is the underlying market. I could be wrong, but I am not sure whether there is anyone working on this. The only people I am aware of are me and my colleagues, basically. That is not to sell myself—that is not the point. We started working on this because we did not see anyone else working on it from the perspective of where you can buy these things, how cheap or easy it is and how much online content online we can see is feasibly fake. That part is substantially underestimated in terms of both resource and policymaking.

Q16 **Uma Kumaran:** We would definitely welcome you back with the findings of that work.

Dr Roozenbeek: We will do so. They should be out by April.

Q17 **Abtisam Mohamed:** I want to follow on from Uma's question on the techniques that influencers use to manipulate or influence their audiences to get to a specific end. There is traditional media and social media, and right at the start Professor Innes referenced a range of different methods and platforms that are used—public understanding, fact-checking, and I think you said military as well. To what extent are the different techniques used to influence and manipulate different audiences?

Professor Innes: I think you can break it down into three basic techniques that are always used. The first is information shaping. That is where you manipulate the information to change how people think, feel or behave. That is your classic disinformation kind of thing, where you manipulate the information.

The second strand is audience shaping, which is where you try to influence the perceptions, beliefs or values of the audience to change how they will interpret information. That is a particular strand that comes from China's cognitive warfare strategy, but it is also the Russian concept of reflexive control.

The third key strand is environment shaping, where you try to shape the environment itself. In the previous question, we were talking about bots. One thing you do with bots is flood the zone to drown out the facts. You

just throw vast numbers of bots at an issue in order to create a lot of noise. It is important to understand the scale at which this is happening.

There is a Russian information operation called Doppelganger, which is a very visible, noisy operation that operates thousands and thousands of bots. It buys them online, deploys them and doesn't expect them to survive. They will get squished, but there are always replacements coming through. As an example, last year we detected them trying to amplify the conspiracies and health rumours about the Princess of Wales. That was about simply jumping on a hashtag or a trending issue in order to get their information into the media stream.

There are a range of different techniques that are operating, some of which are highly targeted and try to hit particular decision makers and shape them, as well as trying to shape public understanding of particular issues. It is that full-spectrum capability that is quite important to understanding what is going on at the moment.

Q18 **Abtisam Mohamed:** How are different factors used? For example, social media influencers and non-traditional social media methods such as WhatsApp are increasingly being used to disseminate information. How does that feature within the three aspects that you just talked us through?

Professor Innes: You will see Russian operations across every single channel and platform that you can think of. Last year we published a report where we found them on Minecraft, for example. On May Day, they built a statue to a famous Russian politician on Minecraft and 12,000 people turned up to visit it. We have got them using LLMs to create ideological chatbots, where people can ask the chatbot, "What does this famous Russian political ideologue think of the war in Ukraine?" He thinks it's great—surprisingly. You have got them on every platform.

Last year in September the US Department of Justice issued an affidavit targeting the Social Design Agency. In the appendices to that affidavit are some really remarkable details that give you a much better sense of how these operations are being organised and conducted.

Q19 **Abtisam Mohamed:** Sorry—the appendices to which document?

Professor Innes: It is an affidavit issued by the Department of Justice in September 2024, relating to the sanctions against the Social Design Agency, who are the people who run the Doppelganger campaign. It was able to show that the campaign design starts with very detailed open-source research on the west. They are monitoring and scanning our media systems, and they are pulling economic trend reports and public opinion data.

With that material they are trying to identify the pain points, as they refer to them, in the west—the issues where they think they can get traction. They will then design a campaign to do this. They have country plans, where they say, "These are the issues that we think we can hit. These are the channels that we think we can do it through." They will then design their social media assets, and we can see that they have performance

targets for how many textual bits of disinformation or narrative they put out, how many images, how many memes and how many videos. They prepare comments in relation to that and then they monitor it.

In the US documentation that was released, we found out that they are systematically monitoring over 2,000 social media influencers to see if they repeat any of the narratives or content that they put out. They are systematically monitoring western think-tanks to see if we are doing it. They are systematically monitoring western media to see if the ideas, concepts and narratives that they are planting are getting picked up and repeated. We have got quite a good nose-to-tail kind of thing on how they are operating and how they measure their performance.

Q20 **Chair:** It is not obvious quite why the Russians would want to spread disinformation about the Princess of Wales. Are you suggesting that was a sort of test to see how well they could succeed in spreading it?

Professor Innes: There are two dimensions to it. First, at the time—I think it was about March or April last year—it was probably the biggest story on social media on the planet. If you can get into that news cycle and media cycle, you suddenly get a lot more eyeballs on the content. They were amplifying this and then they were dropping in things about Ukraine. There was a content layer to it. Secondly, one of the key geostrategic purposes behind this is to undermine trust in institutions. The royal family is a significant UK institution. I think it had a dual purpose.

Professor Tolz-Zilitinkevic: Can I add clarification from my own research? It is important that we do not confuse intent and impact. If we look at impact in relation to the western world, Russian campaigns are very modest at best.

My team looked at the Doppelganger campaign—the same that we have identified. The following of the outlets is modest at best and non-existent at worst. These are not replacements for RT and Sputnik—even though their audiences in the west were modest, they were still visible—because these outlets have invisible audiences in the west. They are a threat because Russia uses them to try to confuse, create moral panic, swamp and disorientate.

We should not fall into that trap by exaggerating their impact on western societies because the goal is to undermine not only institutions like the royal family, but to undermine trust in any media or communication and, most importantly, in significant democratic institutions such as elections. So, we have to be very careful in separating intent and impact.

In fact, in Russia as elsewhere, the actual work is outsourced to third parties who do it for pay. For those actors the goal is not to change somebody's opinion but to make money, therefore, let us say the troll factories do not work as intended. That is not my research, but the research of people at New York University, Harvard University and a group of social scientists at Manchester University—I am not a social scientist.

The impact of troll factories in the context of elections is close to zero. Why? Because these trolls are given tasks such as, "You have to get this many followers" and "You have to produce this many messages and images". That is all true, but they do it at a minimum effort because they do not care about the outcome. The first troll factory was created in South Korea by a businessman who wanted to win elections in, I think, 2009. He was arrested for his activity; he did not get anywhere. All the information about the activity of this troll factory was given to researchers, so it is a fantastic source of data.

The people who analysed the South Korean experiment also looked at the Russian troll factory near Saint Petersburg. They work identically, starting at 9 o'clock, then the communication stops between 1 and 2 o'clock when they are out for lunch, and they finish at 5 o'clock in the evening. A lot of followers are their family members. They keep putting up the same images and the same messages. There is some attempt to target audiences, but it is fairly superficial because their renumeration is connected with the number of posts they put up and the number of followers they have.

There is some fantastic research conducted by a group of people from Harvard University and New York University on the Russian influence operation or interference in the US elections in 2016. That was the largest ever attempt by any state to interfere in the elections—they really worked at scale. There was a whole troll factory near Saint Petersburg producing multiple social media posts, and some of them had more significant traction.

They even sent people to swing states to conduct research. They prepped their trolls on what to do, but it is not in a troll's interest to try to, let's say, influence people in favour of Trump in a swing state. They do not know where they will be followed, so they produced a lot of pro-Clinton posts in, let's say, New York, and pro-Trump posts in Texas. The most followed account was a pro-Trump account in a district of Texas that is the most heavily pro-Trump.

The conclusion of my colleagues from Harvard and NYU was that operationally the work of that troll factory in 2016 in the US election was an operational failure, because it did not influence the result of the election, but it was a huge strategic success because of the amount of publicity created around the interference and it sowed doubts in the electoral process. We have to be very careful not to exaggerate and to see the threat where it is.

Again, in relation to the 2016 election in the US, there was a threat from Russia, but it was from hacking. Part of the operation was the troll factory, but there were also Russian hackers. They hacked the Democratic National Committee accounts and they leaked, through Assange and Wikileaks, the information hacked from Clinton's computer and other DNC computers. Because of the nature of the mainstream media, what the mainstream media in the US and elsewhere discussed was not that there were foreign powers hacking computers of one of the two main US parties, but the

content of the messages, and that went against Clinton. In that sense, the Russian operation might have had an impact on the 2016 election, but we have to be clear where the threat lies.

Q21 **Abtisam Mohamed:** You referred to trolls, but bots are slightly different. You referenced AI, Dr Roozenbeek, and the use of AI in generating information that may or may not be misleading. Can you expand on that? In your view, to what extent is AI overtaking?

Dr Roozenbeek: It is a difficult question. I echo Professor Tolz's sentiment that the Russians do not have unmatched powers of manipulation. They do not have a unique insight into human psychology that we do not have; it is simply a dedicated effort to pursue these types of operations, hoping that they will pay off in some way. They seem to be thinking that for the moment. That means that we should be taking it seriously, but not, as Professor Tolz correctly points out, to refer to it in a panicked manner, as though we are collectively being manipulated by the Russians into doing things we do not want to do. That is the wrong frame—that is not what is happening.

AI is a tool. It is pretty easy now—a bit easier than it was before—to use generative AI, specifically large language models, to post content online. Previously that would have to be either procedurally generated, which is pre-generated content, or written by humans. We don't have to do that anymore. It is now becoming somewhat possible for bots to reply to people in a way that looks kind of natural.

Q22 **Abtisam Mohamed:** Is that overtaking troll farms?

Dr Roozenbeek: It is hard to say. I think they are operating in tandem because they both serve unique purposes. What I do not think AI is able to do better than regular old trolls were doing is overcome content-moderation measures that are already in place, such as rate limits that mean you can only post so many times an hour before you are flagged as suspicious, or that flag you if you keep repeating the same phrase over and over, and so on. There are policies in place—the Online Safety Act, the Digital Services Act in the European Union, or the policies that the platforms themselves impose—that AI cannot magically get around in a way that we were not able to before.

In that sense, AI makes the creation of authentic-looking content easier but not its production. There is no really good evidence that AI is better at lying than humans are. We have tested that and there is a bit there but it is mostly not the case that AI-generated lies are seen as more persuasive than lies created by humans, or whatever we were able to do before.

In that sense, I do not want to overestimate the importance of AI in this debate. It is there and used as a tool because it makes things easier but, at least as the weight of the evidence currently shows, it is not a unique way to short-circuit the human mind and make it easier for you to acquire false beliefs.

Q23 **Abtisam Mohamed:** Are we able to use AI to counter the threat, or try to

identify bots out there?

Dr Roozenbeek: That is an open question. In theory, yes. In practice, you are running a race against someone who is running at roughly the same speed. It is a kind of Red Queen problem, like in "Alice in Wonderland", where everyone is running just to stay in the same place.

As the manipulation tools become more advanced, so do our detection tools but, relative to each other, they are staying more or less in the same place. Work has been done on using AI chatbots to reduce conspiracy beliefs and so on. That is helpful because it can be scaled easily but these are not unique tools that offer protection against manipulation in a way that is unmatched by what we had before. It is roughly in the same realm, so it is not bad but not a cure-all.

Q24 **Abtisam Mohamed:** Do we counter the threat at all? Do we flood the system in the same way that our adversaries may do?

Dr Roozenbeek: I will speak only to the realm that I am familiar with and leave the rest of the answer to Professor Innes and Professor Tolz. The UK Government are now investing quite heavily in media literacy, as part of the curriculum. There is a debate currently ongoing as to whether that should be implemented in every curriculum in schools. That is what people like to see as a psychological defence against manipulation.

I think there is value in that but it is only limited to schools. To do it as, for instance, the Finns do it—they are quite gung-ho about teaching media literacy for various historical reasons—requires a lot of additional effort. But bear in mind that schools also need to teach maths, English, French and so on. They have a lot on their plates already. To uniquely put it on the plate of the education system is not fair.

There are a lot of initiatives in the UK on fact-checking, which has been quite helpful, and what they call "pre-bunking", or pre-emptively debunking misinformation. That equates roughly to media literacy, in a lot of ways and it is ongoing but that is the psychological component, raising individual-level defences. It says nothing about, if you will, systemic protection. What do we do about algorithms, and so on? What guidelines do social media platforms need to meet to be able to operate in the UK? Those are all discussions that are quite separate from that, and which are much harder and more scaled.

Chair: We have a couple more questions about AI.

Q25 **Blair McDougall:** On AI, we are talking a lot about prevention but not necessarily accountability. If Russia drops a bomb on a hospital, we hold Putin accountable but we also look at the supply chain that allowed him to have the missile in the first place. Should we have a similar approach to AI, and hold AI companies to account for the use of their systems, if they are offering hostile regimes something that can be used without limit?

Dr Roozenbeek: I will give a very brief answer—because I have been talking a lot. Yes, in principle. Companies like OpenAI are already taking this quite seriously, and they are, in my experience, at least, relatively

easy to talk to about this. But the problem with the nature of large language models is that anyone can train one, or fine-tune one, remove the guardrails and use it for their own purposes. So, in that sense, if you have an agreement with the big companies—Anthropic, OpenAI and Meta, basically, and Google, to some extent—then that does cover a lot, but someone can train their own LLM, as has been demonstrated recently with DeepSeek, as well. In that sense, it is only effective to a limit.

Q26 **Uma Kumaran:** Sorry to make you speak even more, Dr Roozenbeek. You have said that prevention is better than cure, and, following on from Blair's question, you said earlier that it is not bad, but it is not a cure-all, so those things are quite interesting juxtapositions.

You worked on a game—Bad News—during covid to fight the disinformation around that. What can we learn from that? Are there lessons to be learned, and can we put things like that in place here to teach not just the next generation but even some of us around this table?

Dr Roozenbeek: Well, that is very kind of you. I think these are the kinds of things that should be seen within the limits of learning, if you will. These tools are useful for making media literacy more accessible, because they use humour, let's say, and maybe original storytelling, to circumvent that resistance that a lot of people have to—well, to put it crudely, a boring kind of, "Here's something you need to know about media literacy." Many people would not necessarily be interested in that, right?

For a number of people, they are like, "Oh yeah, fine, that sounds interesting; I'll play it." For those people, there is a learning effect that occurs in terms of the ability to identify the various techniques that are used in online manipulation. And that is healthy, I think. But the limits to that are that people forget things. Playing it once does not mean that you will remember forever, but that is with all learning, right? If you don't practise how to do differential equations, you forget how to.

So, there is that component, and then there is the fact that, inevitably, it is a voluntary thing, so it is a persuasive exercise; how do we convince people to do this? For some of the games, that works quite well because maybe there is a media buzz or something like that, and, again, that has a positive result, I think, but, again, not a cure-all, because it does not translate to cognitive immunity against manipulation. That is far too high a barrier for any interventions like this to be able to achieve.

It is just that, at scale, there is such a thing as learning, right? Today, more people can read than could 100 years ago, so literacy rates have gone up, so why could we not achieve higher media literacy rates as well? In theory we could, but that requires that being taken seriously in schools, for example, and so on. We also need to find a way to get people excited to engage in this kind of learning. That is roughly the task, which is not easy, especially for stuffy academics such as myself.

Professor Tolz-Zilitinkevic: Again, I have not done my own research here, but there are, of course, people who are against pre-bunking—scholars who question pre-bunking. There are experiments showing that

warning people about disinformation can actually provoke the opposite reaction—they become more interested in it and more willing to engage. Similarly with fact checking, there is a danger that people hear the same kinds of "facts", which are false facts, but it does not fully register with them that they are hearing a critique. Something is repeated, repeated, repeated, and even though it is from fact checking—it is disputing the "fact"—people start accepting it as true simply because they hear it very often. That is one issue, here.

Also, I am doing historical research. Research on disinformation started in the 1920s, in the aftermath of the first world war. There were media literacy programmes in the United States in the 1930s, but eventually they were seen as not effective. I think the consensus on what disinformation or propaganda campaigns can achieve in terms of influencing people is keeping loyal audiences on side. It is very difficult for a specific campaign to change people's opinions. People are invested in particular positions.

It is very clear that Russia, for instance, has different strategies domestically and for foreign audiences in terms of manipulation of information. For the foreign audiences, it is just confuse, distract and create a panic, whereas domestically at the start of the full-scale war in Ukraine, it was interesting that the main Kremlin-sponsored influence campaign was on the social media platform used by older citizens of Russia, which is the group that is on the whole most pro-Putin and pro-Kremlin, and most embraces official state propaganda.

It was not on social media platforms that are used more by younger people who could be more critical, because they wanted to keep this audience on side and they knew that they could do that, whereas they were not sure that they could change the opinions of those who were sceptical. We should remember that as well. It is all very complex.

Q27 **Alex Ballinger:** Thank you. I am glad that you mentioned young people because we were talking earlier about elections, and we were speaking to people from Romania about the cancelled elections a few months ago. They certainly believe that there was large-scale Russian interference through TikTok. Professor Tolz, could you say something about the interference there? Should we be worried about that sort of interference in the upcoming election in Romania, and maybe in Germany later this year?

Professor Tolz-Zilitinkevic: I am sure that there was a campaign in Romania on TikTok, organised by Russia. It is documented, and there is no doubt about that. Whether the campaign led to the election outcome nobody knows. It is actually very difficult to prove.

The debate that is now emerging in Romania is that the incumbent politicians are using the fact that there was this Russian campaign, to challenge results that they do not like. Again, that is generally accepted in Romania and elsewhere. This is quite a dangerous tactic, because elections are the core of any democratic state. Questioning the integrity of the electoral process potentially is lethal for democracy. One has to be very careful.

Within our project, originally we did not have Romania as a case study, but we now have more funding to look at the Russian campaign there. I will be happy to share the results with you once we have them. I have no doubt that there will be a Russian campaign related to German elections.

Again, I think it will be very difficult to prove that there is any sort of significant impact. But historically, we know that Russian actors have been collaborating with Alternative für Deutschland, and I presume that they will do that again. But there are other reasons apart from Russian manipulation why people, including younger people and particularly in former East Germany, tend to vote for Alternative für Deutschland. To me, this is the first question that needs to be addressed.

Disinformation actors, including Russian disinformation actors, try to amplify existing problems and cleavages. As Martin has already mentioned, they usually don't invent their own conspiracy theories or stories; they use something that is already there, has salience and is disruptive, and they try to amplify it.

Q28 **Richard Foord:** Has that changed since the cold war? Is the approach that Russia takes now—using a modicum of truth or some basis to build the disinformation upon—something that the Soviet Union used to practise during the cold war, or is it a new departure?

Professor Tolz-Zilitinkevic: I don't think it is a new departure, and it is not just the Soviet Union. As I say, the first really modern disinformation campaigns were practised by all incumbent Governments during the first world war, and then by the Soviet Union and Nazi Germany.

Even in Nazi Germany, Goebbels acknowledged that disinformation—manipulated information—has to be factually accurate, and the Soviet campaigns were, in a way, very similar. They used a lot of stuff that was factually correct, and they added to it things that were false and were useful to the goals of the Kremlin. Most importantly, they used stories, and particularly conspiracy theories, that were already salient. I think the techniques are actually very similar to what they were 100 or 50 years ago; now we simply have new media technologies that allow these campaigns to be organised more quickly, more cheaply and on a greater scale so they reach more audiences.

Q29 **Alex Ballinger:** Professor Innes, we talked earlier about the UK's response. Can you share some reflections on how effective our sanctions have been against RT and other agencies? How well do you think the Online Safety Act might prepare us for this information space?

Professor Innes: The impact of sanctions in this area is not really understood. Some of the individual organisations that we have looked at, which have been sanctioned by ourselves, the US and the EU, seem to have taken that as a badge of honour, and it has resulted in them receiving awards for their work in Russia—

Alex Ballinger: Gosh—okay.

Professor Innes: They have secured more contracts, and secured closed contracts, so they are not bidding on the open market. The unintended consequences of some of this are quite significant.

We could look internationally, because there are things that we could learn from elsewhere. The Viginum unit in France is excellent, in terms of the quality of its analysis. Reference was made to the Psychological Defence Agency in Sweden, which is a very good, very interesting agency.

We can look back at what has happened over the last year or so to start to understand that, although there are limits to what we can do here, there are things that can have a practical effect. We can create friction. One of the really interesting elements was the disruptions that the US ran prior to the elections, whereby Facebook and others deprioritised political content in their algorithm. We could see in the chatter online that the Russian operators were responding to that and talking about how they were suddenly having to work much, much harder to get their content in front of people.

There are some really tricky issues here. What does a concept of deterrence for information operations look like? What are the things that we can do? How do we do them in ways where we don't induce the unintended consequences that we are trying to avoid?

Q30 **Alex Ballinger:** Finally, when we were in France, we heard about their work with Telegram. Obviously, the CEO was arrested after Telegram did not co-operate with the French authorities. We met the Brazilian ambassador and spoke a little bit about how they suspended Twitter, when it was not complying with disinformation requests from the Brazilian Government. Do you think we need to be tougher on social media companies? Are we allowing them to have too much sway?

Professor Innes: One of the recurring themes that has come up in the evidence today is, "What is the balance between doing something about this, protecting free speech, and establishing the thresholds for intervention?" I don't think that has ever been done in a particularly clear and well-defined way.

We need to appreciate that, when we are doing that, there are number of different stakeholders and parties in the room—there is Government, the public and the platforms in the middle. Personally, I certainly think there is a case for saying that we need to revisit the extent to which platforms are responsible and accountable for the things that are published on their services, out of which they are making considerable amounts of money.

Q31 **Aphra Brandreth:** Professor Innes, we have talked extensively about the main hostile state threat actors that are spreading disinformation. Could you speak about which are the most important non-state threat actors, and what challenges there are in attributing disinformation from those actors to states?

Professor Innes: One of the key movements that we have been tracking in how information operations are organised and conducted is a shift to

what we refer to as "living off the land". Earlier versions of the Russian and Chinese information operations around the 2016 election were very much saying, "We have employees. We curate accounts. We build accounts." That is not what is happening now. Instead, state information operations are scanning social media and looking for narratives that they can exploit and amplify. It is not a question of there necessarily being a divide between particular individuals, who are non-state, and particular state-based operations. It is the interaction and the blend between them.

Dr Roozenbeek talked about the monetisation of this. There are undoubtedly influencer-type actors online who are monetising this kind of activity. They are realising that, by promoting particular positions, generating conflict and attracting attention online, they can make significant amounts of money out of doing this. There is pretty good evidence of the fact that is something that the state-based operators are looking at, and they are seeking to get behind and amplify those. It is a really tricky problem to work out what it is that you do about that.

Q32 **Aphra Brandreth:** Related to that, you have written about conspiracy theories in the aftermath of terror attacks, and we saw the spread of misinformation after the Southport attack, for example. How are states amplifying domestic misinformation, particularly in the context of terror attacks?

Professor Innes: Southport is a very good example of the dynamics that happen. What we can see is that the conspiracy rumours occur organically, but there are a group of Russian actors who are monitoring the media space. They look at these and they go, "Right. That one looks like it has traction. We can get behind it. We can boost and amplify that, and that will help us to achieve our geopolitical aims of sowing chaos, discord and all those kinds of things." That is a repeated pattern that we see happening.

That is one of the challenges for Government in this space. For a considerable time, the Government have established a division between the foreign policy and foreign affairs dimensions and the domestic policy. Information does not work like that; information is not confined by borders. Understanding the flows between these agendas and which parts of Government are responsible and should be held accountable for that is quite challenging and quite new.

Q33 **Richard Foord:** Is the responsibility across Government Departments too diffuse? There is responsibility for separate aspects of this in FCDO, the Home Office, the Cabinet Office and the Department for Science, Innovation and Technology. Is it too diffuse to actually be effective?

Professor Innes: I think there is a need to decide if you want a decentralised model or a more centralised model. You could do either.

One of the things that strikes me in this space is what we might refer to as a degree of mirroring. Military information operations folks are very good at understanding what their adversaries do. Intelligence services are very good at understanding that. Foreign affairs diplomats understand what the

other side do. We need all those bits, but whether and how you join them together is the critical thing.

The bit that has interested me for a while, particularly when we look at what Russia is doing, is the outsourcing and the use of contracted agencies that are more creative and agile in terms of their approach. They are like digital natives. Whatever we put together, we need to appreciate that there are different dimensions to this, and it probably transcends the traditional remit of any one single Government Department.

Dr Roozenbeek: With regard to individuals being involved in this space, it is not completely possible to see this as disconnected from the financial market, crypto scams and so on. Crypto scammers make use of the same strategies and techniques that influence operators make use of. They buy from the same marketplace, effectively, and they use the same levers that they have available to gain popularity, get people to buy stuff and so on. In that sense, political influencers are very often tied to states, for obvious reasons, but financial intelligence is not always.

Q34 **Chair:** Can I ask you to measure the scale of the problem? Are there any elections or particular campaigns where you think a significant impact has been caused by disinformation sponsored by Russia or elsewhere?

Dr Roozenbeek: Could I ask for clarification?

Chair: Yes. We have heard a lot of evidence about elections where there have been attempts, and other campaigns where clearly a lot of activity was going on. As people studying this, have you actually said, "In this particular case, the outcome or the general belief was significantly affected by what Russia or others were doing"?

Professor Tolz-Zilitinkevic: I would say yes. There were elections or votes in which the amount of disinformation was huge. It is difficult—

Q35 **Chair:** Which would you say specifically?

Professor Tolz-Zilitinkevic: I will say. Probably that had an impact on the vote, but my position—others might disagree—is that that usually comes from actors internal to that state, because they are much more followed.

I cite two elections: the election in the United States in 2024 and the Brexit referendum. There were many, many disinformation actors, an enormous amount of manipulated and very poor-quality information, and you had particular outcomes. There was an interesting Oxford Internet Institute study comparing informational space in the United States, where it is regarded as very poor. The institute created maps of the density of disinformation actors internal to the particular states. The country with the greatest density is the United States. There is an explanation as to why, in the United States, so many dubious online information providers appeared so quickly. Then comes China, and then comes Russia.

The quality of media spaces, combined with very low levels of trust in mainstream media in the United States in particular, created a very

unwelcome environment during the elections in terms of the quality of information that people consume. Again, that was an interesting study compared with the situation in France: the media space during the elections, which French citizens consume, is much healthier than in the United States.

I started my comments with this point: we cannot separate foreign disinformation with disinformation providers internal to democracies; the two very often work together. The Russian state and state-affiliated actors who participate in the campaigns, which Martin described in great detail, collaborate, and they could not do what they are doing without working together with disinformation providers internal to democracy. My examples are Brexit and the US election, but there probably were also elections in other parts of the world; I just don't know about them.

Q36 **Chair:** You say, "internal to the state". To some extent, you could say that what we all take part in, which is political debate where you say, "If you vote for the other side, they will do this," could be termed disinformation. Where do you draw the line between what is robust political argument and disinformation?

Professor Tolz-Zilitinkevic: That is an excellent question, and scholars go round and round in circles discussing that. That is the issue we are raising in our project.

Let us take three outlets I looked at, which were set up by American citizens and were very active in the elections in 2020 in America. They are Project Veritas, *Breitbart* and Alex Jones's *Infowars*. You can see the amount of plainly false claims. I think I could judge, on the purely factual level, the amount of very, very twisted claims that are clearly articulated for a purpose—it is absolutely overwhelming.

Maybe Alex Jones believed in what he was saying. That is another matter. But that is a difficulty because disinformation, as I said, works in terms of how people engage with it, not at the level of individual facts. It works at the level of "This story seems plausible", and what seems plausible differs in different linguo-cultural contexts. Let us say, in Russian campaigns the actors know that, and they tailor the story to a particular context. Some kind of story blaming the United States for creating a virus or something in the laboratory, might look dubious to people in western countries, but in Asia, Africa or the Middle East it could sound very plausible. Then it doesn't really matter what the facts are.

Of course, even with conspiracy theories, some turn out to be true and that creates a problem. Something could be perceived as false or true and then the status of the story can change with time.

Let us say that originally some suggestions were proposed around the treatment of covid with particular drugs and that was a plausible story. Then the efficacy of these drugs was questioned by reputable scientists. But some Russian actors continued to promote those drugs which were discredited. Then the status of the report changes.

It is very difficult to draw the line. It is also difficult to draw the line regarding where we should or should not panic, because if we panic too much it plays into the hands—

Q37 **Chair:** The distinction that we heard about when we were in France—I think Dr Roozenbeek referred to this—was that they said, "We are not there to remove content or say that content is wrong or right. We are targeting the dissemination—the artificial amplification using technology—which is being carried out by hostile states."

That seemed to some extent get around the argument that you are impeding freedom of speech—if you are stopping content being artificially spread, then that is a legitimate action for a state to take against disinformation.

Professor Tolz-Zilitinkevic: I would agree.

Dr Roozenbeek: I agree. With respect to your earlier question about the difference between robust political speech and domestic disinformation, in the margins the differences are very difficult, and also it is not the place of the Government or even academics to decide.

Rather, what we are really talking about is not a classification system, but a set of norms that we rely on in service of democracy. For instance, that you don't lie with impunity for political purposes is a norm that we hold, because we think it is important not to do that in pursuit of your own political career, say.

If someone violates that norm, we have a normative judgment about that violation. But that doesn't mean that the levers of Government are equipped to deal with that particularly well. That is often the clash that we run into. It is very awkward to observe. If a political party or actor lies with impunity, not caring about this norm, and there is no concerted effort to counteract that lying—whether it be through formal or informal punishment, or by being shut out of a debate, let's say—then that norm is changing over time.

That is the finding of a particularly interesting study from two weeks ago by Petter Törnberg, a researcher from Sweden. Very awkwardly, he finds that within the democratic world this particular almost callous disregard for what is true and what is not, in terms of what information you spread and what narratives you create and so on, is almost exclusive to the populist right—not to the right and not to populism, but to the populist right as it currently exists, in this present moment, in many democratic countries.

That is a very awkward conclusion, because it almost sounds like I am saying the populist right is wrong in this sense. I am not saying that; I am saying that, observably, that is what is happening in a lot of different countries right now. That is an effect that exists at the moment, assuming that the results from that study are correct, but there is no guarantee that that will be the same in 10 years or that it was the same 20 years ago, and so on. These are temporary political moments we find ourselves in

that are very difficult to grapple with. But now we are having a discussion about political norms, not disinformation.

Q38 **Aphra Brandreth:** I want to jump in on this theme. Professor Tolz-Zilitinkevic, you said that the scale of a disinformation campaign might not be reflected in its impact. Thinking back to the discussion we were having about elections that might have been influenced by disinformation, how do you, or any of the panel, think we can measure the impact of disinformation campaigns? You mentioned a couple of elections that you thought might have been influenced, but how can we know with any degree of certainty the impact that a campaign is having?

Professor Tolz-Zilitinkevic: As far as I understand it, we cannot measure with any degree of certainty. There is a lot of research in relation to the 2016 elections in the United States, with leading scholars concluding that the troll factory does not seem to have had any impact because the trolls were not that active in the swing states, where US elections are won or lost.

But in relation to the hacked DNC accounts and the release of emails, where the hacking was conducted by Russia, the impact is speculation—it might have had an impact. On the whole, my understanding is that most scholars believe that the technology is not yet good enough to have a predictable impact. Maybe it will improve and become more sophisticated and that will be possible.

The research on the impact of the manipulation of information on individuals has been going on for exactly 100 years with basically the same conclusion: that it is very difficult to exercise impact, particularly around elections, because it is short term. If you have exposed somebody to disinformation or propaganda—hyper-partisan content—over a long period of time—a drip, drip, drip influence—then yes, you will have an impact, but it is very difficult to change people's opinion with short-term campaigns. Also, there are not robust technologies to prove that disinformation was the factor.

Q39 **Aphra Brandreth:** Professor Innes, can you come in on that point? I think you might have something to add.

Professor Innes: We have just secured a large research grant with colleagues in the US to do a bilateral study to look at not necessarily whether we can measure impact, but how we think about it in the first place. We need to differentiate between impact as a tactical, operational effect and a strategic thing, or between a chronic and an acute phase. There is a tendency in terms of how we ask the questions to think about an acute effect: did something happen and did it produce an effect on a particular election? Generally, that is not how Russia, or any of the other adversaries we are talking about, think about this.

Ladislav Bittman, the head of the Czechoslovak disinformation department in the 1960s, defected to the West and wrote a memoir. One of the things he said was: "Even though a single operation may not visibly change public perceptions of an issue, a series of related operations will eventually

bring the desired changes." That, I think, is very profound, and it is absolutely how we can see the adversaries that we are talking about, thinking about these kinds of things. Recently, I have been very interested to note that a number of Russian universities are setting up courses to teach this kind of stuff in information warfare—so they are obviously seeing this as a long-term strategic aim that they want to pursue. If we appreciate that it is not just acute or chronic approaches but both of those together, then we start to see a different way of approaching this.

In all of the elections that have worried me, it has not just been that a campaign has been run against a specific election; there has been a long-term influencing effort. You talked about Romania earlier; I am not entirely sure about that because the information is not necessarily in the public domain, but we have seen what has happened in Slovakia. Certainly there was a very, very near miss in Moldova in the last round of elections. So it is not just a question of talking about effect, but thinking about, are we thinking about this in the same way as our adversaries are?

Professor Tolz-Zilitinkevic: Bittman's is actually a very good memoir, worth reading. Bittman writes on the basis of his experience in the 1960s and he wrote his memoir when the communist regimes were still in power in eastern Europe. He is of course worried about influence. What happened? Communist regimes all fell. The Soviet Union collapsed. And Putin's regime, with time—I don't know when—will also collapse. Democratic states are much more resilient than these autocracies. They seem threatening, but they are extremely corrupt. Their messaging and their disinformation campaigns are hugely opportunistic.

What impact did the Soviet campaigns have in the West? I was a political refugee from the Soviet Union. My ex-husband was a dissident and we were expelled, dramatically, from the Soviet Union in 1982. My first job was at the Radio Free Europe research department. It was in the early 1980s. At the time, the Soviet Union mounted what they regard as their most successful disinformation campaign. That was about the origins of AIDS—that it was created in a United States laboratory. It had an update in Africa and the Middle East, in India as well, and actually it was my department that noticed that campaign at the time.

I find it interesting that there wasn't the same amount of concern about Soviet disinformation then. To me, the explanation is that it was because western societies were much more self-confident; there was much less perception of internal problems and worries that we didn't know what the solutions were. I still believe that western democracies are far more robust, and—this is important—we need to gain more self-confidence rather than worrying about these actors. In 1989 the CIA produced a report, incredibly exaggerating the Soviet military and economic might—and the country was on its last legs, and in 1991 the whole thing collapsed.

Yes, Putin's regime is into this type of information manipulation. A lot of actors are earning money. It is not the doppelganger campaign—Moscow spends a lot of money on this operation, but it seems to me the main

point is that it is laundering the money abroad. Through this operation, it is transferring money to Alternative für Deutschland and other populist right-wing parties that it is co-operating with. Again, the following of this outlet in the West is meagre.

Chair: I am afraid we will have to draw to a halt. Richard has a quick last question.

Q40 **Richard Foord:** During the cold war, the Foreign Office used to have a department called the information research department. I am not sure if you are familiar with it, but I understand it served to counter Soviet misinformation but also pursued its own information campaigns. Will you comment on whether there would be any value in having anything similar today in the FCDO?

Professor Tolz-Zilitinkevic: Fund BBC Monitoring. That is a fantastic organisation and it is sort of independent. I don't know the quality of the work of the outfit you mention, but BBC Monitoring does an excellent job.

Q41 **Richard Foord:** Have you heard of the information research department, Dr Roozenbeck?

Dr Roozenbeek: I think a robust media landscape does that job really well. UK media is healthy, to an extent. I am quite a big fan of Ofcom, frankly—I think it is doing a great job—but a lot of UK media outlets are doing excellent work and investigative journalism. They are doing a lot. If the Government take on that task, it is easily delegitimised.

Q42 **Chair:** I totally agree with that, except about the media being well financed and successful, because they are struggling. But on their role, I completely agree.

Dr Roozenbeek: They are producing great work; let's put it that way.

Professor Tolz-Zilitinkevic: Ofcom is a great organisation.

Chair: It has been an absolutely fascinating session. May I thank the three witnesses very much? This was an introductory session for us, and you have exposed the vast range of issues we are going to have to grapple with. Please, if you have any further thoughts, do not hesitate to send them to the Committee. On behalf of my colleagues, I thank you again for coming.