



Select Committee on Communications

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

Tuesday 26 June 2018

4.35 pm

[Watch the meeting](#)

Members present: Lord Gilbert of Panteg (Chairman); Lord Allen of Kensington; Baroness Benjamin; Lord Bishop of Chelmsford; Baroness Chisholm of Owlpen; Viscount Colville of Culross; Lord Goodlad; Lord Gordon of Strathblane; Baroness Kidron; Baroness Quin.

Evidence Session No. 11

Heard in Public

Questions 93 - 102

Witnesses

I: Dr Ewa Luger, Chancellor's Fellow, Digital Arts and Humanities, University of Edinburgh; Professor John Naughton, Senior Research Fellow, Centre for Research in the Arts, Social Sciences and Humanities, University of Cambridge.

USE OF THE TRANSCRIPT

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

Examination of witnesses

Dr Ewa Luger and Professor John Naughton.

Q93 **The Chairman:** I welcome our witnesses to the second evidence session this afternoon in the House of Lords inquiry into regulation of the internet. I remind them that the session will be recorded and broadcast online, and a transcript will be taken.

Our witnesses are tech experts. We are very grateful to you for taking time to be with us today. Would you start by briefly introducing yourselves?

Dr Ewa Luger: I am a chancellor's fellow in digital arts and humanities at the University of Edinburgh in the Centre for Design Informatics. My disciplinary background is political science and, more recently, human-computer interaction, specifically the ethics of intelligence systems. I am a consultant researcher to Microsoft Research on the subject of artificial intelligence and ethics.

Professor John Naughton: I am a senior research fellow in CRASSH, the Centre for Research in the Arts, Social Sciences and Humanities in Cambridge. I am the technology columnist of the *Observer* and a historian of the internet. My background is as a systems engineer, and over the last two decades I have been studying the impact of the internet on society. I recently finished a research project, which Professor David Runciman and I ran at Cambridge, on the implications of digital technology for democracy.

The Chairman: Thank you. Could you start by giving us a brief account, if that is possible, of how the original infrastructure of the internet developed. At the time of its development, and in its early years, what were its values? How does the internet as we know it today differ from the internet as it was invented and the values envisaged when the internet first came into our lives? Who is best placed to start?

Professor John Naughton: Having written a history of it, perhaps I should start.

The Chairman: I think there is a history lesson for us.

Professor John Naughton: Before we start, can I make one point? It is important to distinguish between the internet, which is the underlying technology, and the services that run on it. I have found over a long period of discussion with politicians and others that that distinction escapes many of them. For example, many people think that the worldwide web is the internet; many people used to think that Google was the internet; and there are now probably a billion people in the world who think that Facebook is the internet.

That is important, because in any discussion about regulation we have to distinguish the infrastructure from the services that run on it. Much of the conversation is about what we do about the manifest social harms that

some companies and services that run on the network are doing, so it is important to distinguish.

Going back to the history, the network we use today is very old technology; its origins go back to the 1960s, perhaps before. The network we now use, based on the TCP/IP family of protocols, was switched on in January 1983. Design work on it went on for 10 years before that, starting in autumn 1973.

The people designing the network were faced with an acute problem: how do you design something that has a reasonable chance of being future-proof? They approached that problem by having two fundamental axioms. One was that there should be no central ownership or control of what they designed; the second was that they should design a network that was not optimised for anything they knew about at the time. That meant in the end that they designed a network that was, in their words, extremely stupid. It did only one thing; it took in data packets from one of its edges and it did its best, with no guarantees, to deliver them to the destination at the other side.

The implication of that was that they left all the ingenuity to the edges of the network. If someone had an idea that could be realised using data packets and they were smart enough to write the software to do that, the internet would do it for them with no questions asked. There was nobody they had to ask for permission to do that. In other words, although they did not use this term at the time, Vint Cerf and Bob Kahn, the two originators of the idea, produced an architecture for what later came to be called permissionless innovation. That is absolutely critical, because it meant that they designed a system that enabled a huge and absolutely staggering explosion of creativity. In essence, those two axioms enabled us to create a global machine for springing surprises. That is the best description I can think of for the internet as an architecture.

As we know, some of those surprises have been very pleasant. The worldwide web is one; it is more or less the invention of a single person – Tim Berners-Lee. We regard Voice over IP, Skype, Wikipedia and a whole range of things as the great benefits of the network for society, but because it enabled permissionless innovation, and not only good people are ingenious, it turned out that it also enables a lot of rather nasty surprises, which is one of the things we are concentrating on.

I bring that up because the network itself, as I have described it, requires some attention from a regulatory point of view, but it is not the same kind of attention that we need to focus on the companies that have captured and dominated it. The values, as I said, were openness and the sponsoring of creativity, and everything else followed from that.

The Chairman: Dr Luger, what is your perspective?

Dr Ewa Luger: I have nothing to add.

Q94 **Lord Gordon of Strathblane:** I think I know the answer to my question

in light of your introductory remarks. Was the development of platforms almost inevitable?

Professor John Naughton: The evolution of platforms was inevitable; the evolution of the platforms that we got was not. The first great platform, the greatest of all, was the worldwide web, which was conceived and implemented as a public good rather than private property.

Lord Gordon of Strathblane: What do you think are the consequences of the way the internet has developed?

Professor John Naughton: If you look at it from a historical point of view, it has followed a pattern for which there is a great deal of historical evidence. Some years ago the American scholar Tim Wu looked at the history of the great communications technologies of the 20th century in the United States: the telephone, broadcast radio, broadcast TV and the movies. He found the same pattern in each case. The technology arrives; it is exciting, chaotic and open, and encourages all kinds of utopian hopes, but it is hard to use at the beginning. Eventually, along comes a charismatic entrepreneur who makes an offer to consumers. Most normal people, by the way, are not early adopters. The only people who are early adopters are masochists like me, because we like the challenge, but normal, sensible people do not; they just want something to work.

At the beginning, none of those technologies just works, but along comes an entrepreneur—at the moment it is always a “he”—who makes a proposition to the consumer: “I’m going to give you something you can use out of the box and you won’t have to think about it any more”. In the case of the telephone, it was Theodore Vail of the Bell network whose offer to the consumer was, “If you have a Bell telephone, I make you two promises. The first is that when you pick up the phone you’ll get a dial tone. The second is that you can talk to anybody else in the continental United States who is on the Bell network”. Those were the two propositions he made. The same is true of the movies and other things.

You get to the point where an entrepreneur arrives, makes a proposition and the industry is captured, sometimes with government or regulatory approval, as in the case of AT&T or the Bell network. That happened to all those four technologies in the 20th century. Tim Wu’s question was, “Is this going to happen to the internet?” We now know the answer. It has been captured by a number of giant corporations.

Lord Gordon of Strathblane: But it is not one individual; it is a number of corporations.

Professor John Naughton: Yes, it is.

Lord Gordon of Strathblane: Are there advantages in the dominance of some of these platforms, in that there is interoperability and other advantages for users? In other words, is dominance a not wholly bad thing?

Professor John Naughton: This is one of the great arguments. I am not an anti-trust specialist, but having listened to the earlier evidence I can see that there are people in the field of competition law who think a lot about it. The point is that it is two-sided. On the one hand, because of the network effects that are very important for technology, it is very convenient for consumers if there is a dominant search engine, such as Google, because, apart from anything else, it gets so much data from its users that it constantly self-improves, and that is a real benefit.

The argument in relation to dominance and anti-trust is very simple. Our original concept of anti-trust was that, if a company became dominant and abused that dominance by, say, gouging consumers with prices that it could not justify, that was a bad thing. The problem we have with these companies is that they do not charge their users for what they provide; the users are not their customers.

Lord Gordon of Strathblane: They are users of the product.

Professor John Naughton: That is right. The real customers of Google and Facebook are the advertisers. There may be a case for consumer harm when it comes to advertisers, but it is very difficult to use the anti-trust thinking that we have in relation to the corporations, simply because their services are free, and because their services, as far as the users are concerned, are rather good. You could argue, and people do, that to punish Google because it has 95% of the market and charges nothing for its services is punishing excellence. There is a case for that, and that is why it gets complicated very quickly.

Dr Ewa Luger: Users tend to expect interoperability. The idea of a seamless user experience is hard-coded in our expectation as consumers and in the rules we adhere to when we design interfaces. Anything that breaches that is problematic. The models that underpin some of that interoperability, and the models by which data is collected and used, are not expected, seen or understood by users. If they were using a Google service and an ad popped up, the majority of users would not necessarily understand what had triggered that advertisement; they would not know that it was data from their personal emails, for example. We do not have good, robust mental models for how these systems operate.

Q95 **Baroness Kidron:** It is fair to say that the evidence we have had so far is slightly binary. One set of people says that the OTT companies are there to share creativity: "Get closer to your family. We help you find things that you might be interested in". The other side says that it is compulsive, quite mindless and deliberately designed to make you do things that you may not otherwise do. I do not think the word "spooky" came into the evidence, but there are those accusations.

My question is about the design. I would like both of you to talk about the interface of where we are now. Professor Naughton, you said you loved the promise of the first order—Wikipedia, the worldwide web, et cetera—but how does the interface affect what users see, and specifically how does it both create their behaviour and then capture it?

Dr Ewa Luger: A lot of the ways in which we design user interfaces are to minimise mental workload. You do not want a stressful experience when you are interacting with the system. In that instance, the more data gathered about an individual, the more likely you are to reduce their mental workload by giving them exactly what they want, effectively by predicting their behaviours and, therefore, being responsive in the solutions you offer them. That is where we are at the moment in the design of systems.

There is a trend, which I think will be a future trend, towards the minimisation of the manipulation of the user interface. You sit typing at your laptop, but increasingly we are seeing systems where we interact through natural user interfaces or where data about us is being passively collected, such as applications on our phones collecting location data passively. There are new applications at an experimental stage that would collect snippets of your voice to judge things such as your current mental health and link that to your location data to see whether or not you are depressed.

Baroness Kidron: I am now.

Dr Ewa Luger: It is a bit worrisome. We are also seeing a rise in voice interfaces, but I do not know whether they will take off; it seems to peak and then dip. Research that I have conducted in that area shows that users do not understand what is going on when they interact through a voice user interface, simply because their expectations are based on the model of communication between a human and a human and a theory of mind, assuming you know what the other person is going to say and that you have a common frame of reference.

That does not work with a computer, but we have no alternatives. There are no robust mental models or metaphors that we can use to communicate how systems collect data, what they do with it, how it flows and how it is used through interface design. That simply does not exist right now. People are starting to explore it. We see papers coming out from the leading human-computer interaction conferences, but they are certainly not principles applied at industry level, and that is desperately needed.

Professor John Naughton: I take a bleaker view. We have to make a distinction between the companies. There are five, and two of them have a very specific business model, which is very unusual in our history. The name we have for it now is surveillance capitalism; it provides free services in return for the unrestricted right to exploit the personal data and data trails left by users. In some ways, it is better to think of Facebook and Google not as tech companies but as if they were oil companies. They extract data, refine it and then sell it in one way or another. That is the simple way to think about it.

Currently, it seems a very successful business. To make it work, they have to make sure that the supply of data and the data trails provided by the users of their free services continually increase. That is the key bit.

How do they do that? It is very simple. They deploy, among other things, much of what is known by applied psychologists—for example, to increase people’s likelihood to go back to something else, and to incentivise users to stay longer on their platform.

A serious degree of addictivity is built into these services, which is why when you talk to people who use them and ask whether they are concerned about how much time they spend on them, you get a funny sort of shrug. Somebody goes on to Facebook to check a picture from a family member and an hour later they wonder why they are still there. They are still there, because it is beautiful software that is very cleverly designed. In a way, the core of this is the business model of surveillance capitalism. That is the key, and, if we want leverage on that particular kind of corporation, that is where we have to look. That is why competition law is important, because that may be the interesting place.

As to addictivity and the rest of it, there are other things—for example, the fact that there is huge peer pressure to use the services. In some cases, the social pressure to be on those services is pretty compelling for young people, teenagers in particular. Although you can say to somebody, “You don’t have to do any of this. Nobody is forcing you to use Facebook or Google”, they will say, “I don’t have any options socially”.

Q96 Baroness Kidron: I should declare that I wrote a report about persuasive design, which was published last week. I put that on the record.

I want to move to the question of children and vulnerable adults. Going back to your historical description, in a way many of the people who designed it say that it was a democratising technology; there were going to be no gatekeepers and all users would be equal. I want to ask about the “all users would be equal” piece of it, because children spend huge parts of their childhood in an environment that was not designed for, nor did it anticipate, their presence, and in which as a concept childhood largely does not exist. I think you can extrapolate from that to some other users. Could I have your opinion on whether the design of services adequately meets the needs of different user groups, particularly the vulnerable?

Dr Ewa Luger: Absolutely not. There are a number of issues. Most services are designed by a particular demographic that does not represent humanity as a whole; it does not represent me. That is a massive problem.

Another thing, which I have been concerned about in my research, is that we have run roughshod over the notion of consent and what it is to consent to the use of your data. I know people push back on the term “your data”, but the traces of my existence made manifest through these systems, because of the way they are architected, are not the same as a tangible block that belongs to me or someone else and is actually a reflection of my beliefs and values.

We know that the consent model—the tick and click terms and conditions model—is broken and does not work. People spend less than 30 seconds reading those kinds of documents; fewer than 1% pause to read the small print of end-user licence agreements. There has been research to show that we are trained to accept. When we are presented with a tick-box option, in comparison with any other option, we tick it much faster. Everything about the way consent is manifest on the internet is problematic.

Add to that the idea that consent has to be voluntary, competent, informed and comprehending, and you can see immediately that there is an issue, because what you are describing is competence. A child is not yet able to give consent in any contextual way, so even before we get to the fact that the mechanism is faulty we are putting children in a position where they are agreeing to things they simply cannot understand.

It is not just children. Most people do not understand the implications or harms that can arise from sharing data with particular types of platforms. Most experts would struggle to give you a fair prediction of what will actually happen in the long term.

Professor John Naughton: As regards the history, I would describe myself as a recovering utopian. I am an engineer and I thought we had really cracked it. We had created a wonderful network that would do things that were impossible before, and so on. What I omitted to notice was that all technologies are socially constructed.

This technology was invented essentially by a very select group of males, broadly speaking, working primarily in elite research institutes in the United States, although some of them were in University College London. They knew one another. The network they designed was one that conceived of its users, first, as equals, and, secondly, as people who could be trusted. For example, at the time of the design of the SMTP internet protocol, which determines how email servers work, nobody thought to build in authentication, so a mail server did not check that the mail coming in was in fact coming from the person who purported to be sending it. That is why we have spam. We had a hole designed into the network because of its social construction. Nobody in the 1980s worried about stuff like that. You knew who the person was or you knew where they were coming from.

To go back to Dr Luger's point, we see that now in the services that run on the network. The demographic of the people who work in these companies is fantastically skewed. It is amazingly male and amazingly white or Asian; it is definitely not black. Sometimes you get absurd outcomes—for example, somebody developing a healthcare app that omits to notice that women have menstrual cycles. Demographics really matter in these areas. This is a very skewed demographic in a strange part of the United States.

Q97 **Baroness Chisholm of Owlpen:** This may be like asking, "How long is a piece of string?", but can I ask both of you to look ahead? What do you

feel will probably be the biggest changes, and what are your concerns about the internet and all the enabled technologies in the next five to 10 years?

Professor John Naughton: The answer depends partly on whether or not democratic societies decide that they need to do something about these companies. I have a provocative proposition, which is that the only regimes that will be able to control these outfits are authoritarian. The Chinese Government have no problem at all with this stuff. The problem for democracies, apart from the rule of law and all kinds of other stuff, is that it is more complicated. The future depends on whether or not democratic Governments and legislatures summon up the political will to address some of the harms that come from this kind of dominance. If they do not, my feeling is that the future looks pretty bleak, simply because the business model of surveillance capitalism requires that surveillance becomes more and more intrusive.

You may have noticed, and it is no accident, that all the big companies are desperate to have a listening device in your house. The reason is because up to now what happens inside your home is, broadly speaking, rather opaque to them, and they want to make sure that it is not because they need a constant supply of data. Apart from human beings deciding they have had enough, I cannot see an obvious end to that at the moment. If we do not have regulation, we will have real trouble further down the line.

The Chairman: Do you think that in the tech companies the discussion about listening devices you have just described takes place in those terms?

Dr Ewa Luger: No.

Professor John Naughton: Dr Luger is more knowledgeable than I am.

Dr Ewa Luger: I have never heard it discussed in those terms in my experience. The values that exist in tech companies are not necessarily the same values that exist outside them. There tends to be a much narrower product focus. There might be talk about efficiency or creating better and more competitive products. As far as I know, nobody ever says, "We're going to put something in your home that monitors your behaviour".

The Chairman: Do you agree, Professor?

Professor John Naughton: I do, and it is very significant. For example, after the recent Cambridge Analytica scandal, the CEO of Facebook was eventually hauled before the United States Congress. You will have observed that he has declined to be hauled before this body, but never mind. The strange thing is that, if you analyse the transcript of those two sessions, the business model of the company is never mentioned—never. It is like the old saying that nobody would ever eat a sausage if they saw how it was made. It is interesting that many parents who work in the

tech industry are very careful about how much they allow their children to use devices.

Baroness Chisholm of Owlpen: Education is paramount.

Dr Ewa Luger: Education of computer scientists is paramount.

Baroness Chisholm of Owlpen: No, education for us. There is no point trying to educate them. I was thinking more of the general public.

Dr Ewa Luger: People are talking about issues such as data literacy and algorithmic literacy, so those things are important. I did not mean to make a flippant point. I believe that we should educate computer scientists in ethics. That does not exist currently.

Professor John Naughton: If you were to take a longer view, the great cultural critic Neil Postman wrote a book in the 1980s entitled *The Disappearance of Childhood*, in which he argued that our concept of childhood is socially determined, largely by the dominant communication medium of the age. He said that was why you never see children in Brueghel paintings; you see only small adults. That was because in the Middle Ages a child became an adult when they achieved competence in the dominant communication technology of the era, which of course was speech. That was why the Catholic Church, from the Middle Ages onwards, set the age of reason at seven.

Postman's argument was that, when print came along, the time it took to get to communicative competence was longer, which was why the age became 12, and we had the beginnings of mass education. He went on to argue, mischievously, that with the dominance of television the age of adulthood went down to three, because you never saw a remedial class in television viewing. The big question for us is: what is this dominant communication technology now doing to childhood? I do not have an answer, but it is an interesting question.

The Chairman: We ought to stick to the ethics of business models, if we may.

Q98 **Lord Bishop of Chelmsford:** Quite a lot of the written and oral evidence that we have received from other witnesses has pointed in the direction of what is known as ethical by design. Dr Luger, this is an area to which you have given particular thought. I would be grateful to hear from both of you what you think is meant by that term, and what principles might be adopted to ensure some sort of ethical by design standard.

Dr Ewa Luger: Ethical by design is one of the terms that has been recently coined. There is no hard and fast definition, but if one were to define it, the most common understanding is that it is the consideration of human values and ethical principles from inception to completion of the design of the technology, from the ideation stage to the point it hits the market. I would argue that we need to extend the context within which it is deployed.

There are four ethical principles that one might adhere to in any context: beneficence—always do good—and non-maleficence, never do bad; autonomy; freedom to act; and justice and fairness. Those things are critical and a lot of them are violated to some extent by the kinds of systems we are discussing.

If we were to drill down a bit to create specific principles, one of them would need to be openness by companies as to how personal data is collected, stored and used, which would include activity tracking and behaviour tracking. There should be an emphasis upon intelligibility or legibility, depending on how you wish to frame it—enough information about how a system operates that a user can meaningfully interact with it. Currently, we do not have that in the kind of systems that are emerging. Next, there is opt-in as default and easy revocation.

Some of those principles are now enshrined in the EU's GDPR, but they have not yet made it into the design of interfaces. How does one do that without, for example, breaching user experience? All those ethical principles have to be offset against developing good user experience, or nobody will use the technology.

Another principle would be enhancing voluntarism. Do people have a genuine choice in using a system? Are there other systems they might use if they choose not to share aspects of their data? Does the system still operate in the same way?

Purpose limitation is also a legal principle, and it is incredibly important. Should the data be used beyond the purposes outlined when the technology was sold or adopted by the user? I would say they should not and the law says they should not, but in reality we see it happening all the time. A really nice mechanism for this in data protection regulation is the idea of the motivated intruder test. Could a motivated intruder re-identify the dataset? Something similar in an ethical context would be pretty important. If you change the purposes, does it become an issue? Minimisation of discrimination is another one that speaks directly to the justice aspect.

In failure handling, how do we deal with transparency and reporting? When a system fails, there is something about interface design that suggests that the computer is always right, but in reality we know that some algorithms are not accurate all the time. Some of that needs to be exposed so that users understand what they are interacting with. In the long term, it is through such interactions that you get a sense of how to engage with the system. Part of it is through good design and part of it is through testing and using a system.

Another aspect is the reproducibility and re-performability of algorithms. Can an algorithm create the same function or output more than once? This is much more complex in the emerging classes of algorithms, such as deep neural networks, but as I understand it it is still possible at some level. We should work hard for that principle.

Finally, we come to provenance. I am surprised that we do not do more instrumentation of record keeping—logging inputs and outputs, and the effects on people through the course of the technology.

Professor John Naughton: This is not my field. I agree with what Dr Luger says in that area.

Ethics also apply at the corporate culture level. Until quite recently, in the technology industry at any rate, ethics were treated like statements about motherhood and apple pie—in other words, as vague bromides. As it slowly dawned on the industry that there might be serious trouble coming, it has started to boost its concerns with ethics. Some of them are quite preposterous in the sense that they are simply public relations stunts. In some cases, there is evidence that a few companies are starting to take this very seriously. The classic one is DeepMind, the Google-owned British artificial intelligence company. I think it is taking ethics seriously.

Having honest business models would be a start. I do not mind paying for, say, Facebook. I am quite happy to pay for it, but I do not want to give it my data. I would like to have that option, in which case we might get to a better place quickly.

The final thing is responsibility. It is absolutely the case that, however critical one is of those who lead some of these companies, broadly speaking they are not evil-intentioned people. In fact, part of their problem is that they believe they are good; they feel that they are transparently good and therefore that they could not be doing evil, but if they operate businesses that produce what look to us like socially damaging consequences, they need to accept responsibility for them.

Lord Bishop of Chelmsford: I was interested in your oil-company analogy. It was a fascinating way of thinking about the companies. Did they set out to be oil companies or did they just discover at some point, “Oh, look, we’ve got this data and there’s money to be made here”?

Professor John Naughton: In the case of the more prominent companies, they did not have any idea what they were going to do when they started. For a long time, both Google founders and the Facebook founder expressed opinions about advertising that suggested it was pretty awful and they would have nothing to do with it. In the end, they discovered that the only way to fund what they wanted to do was to become advertising companies, and they did so.

It is not that their intentions were not good; it is just that they were naive. For example, in the case of Facebook they built an amazing machine for enabling advertisers to target people with messages. It is a terrific machine, and if you go in as an advertiser you really see its quality. It did not seem to have occurred to them that it is not just advertisers who want to target people; it is people like Steve Bannon and other actors. There is a naivety about human nature, which is strange.

Q99 Lord Bishop of Chelmsford: That leads me to a supplementary question. It is about ethical by design in particular, but it may cover other things. What role do you think should be played by committees such as ours, government, academia and private organisations in the development of ethical standards? What do you want to say to us, which we might say in the report we are preparing, about ethical standards that could realistically be designed into these systems?

Dr Ewa Luger: The first point to note picks up something that Professor Naughton mentioned, which is the culture of the tech industry and the fact that people do not set out to do harm, but they do not know what the alternative is. Responsible innovation is not embedded in the teaching of computer science, machine learning or AI. There is now a bit more in robotics, but it tends to be mildly flippant, looking at what happens if robots become sentient. I am not saying that those questions are unimportant; they might be at some point, but that is getting ahead.

Changing the culture requires investment at HE level, not simply once the cat is out of the bag and everybody is happily working in a corporation. There is no real incentive for tech designers who are competitive against each other and the wider world to consider ethics unless it is forced upon them. Certainly it is something academia could think about, not simply coming out with a critique of the issues or saying, "There's bias in this algorithm". Those things are important, but we also need to consider that we are producing a raft of people who are incapable of doing anything else, so how do we change that?

It is a matter of embedding ethics in teaching, and developing solutions. There are all kinds of work on that. The United States is dominant, with the work of the fairness, awareness and transparency agenda—the FAT people—and the AI Now Institute and Data & Society. They are already starting to look at these kinds of things. We need plug and play solutions for industry. If we have solutions that we know work to minimise bias, industry will use them, but the expectation that solutions will come internally from those organisations is probably not one that will happily produce any results.

As to government, one of the things that we still do not understand is the long-term harm from data-driven systems. We can speculate, and we can identify when bias occurs, but we do not actually know what automation will mean—for example, minimisation of the workforce and that kind of thing. We do not really know the long-term implications. There are no long-term studies of such things. I suggest we start to conduct long-term studies to look at where algorithms are deployed in areas such as public health, social security and credit scoring, where the negative impacts might be massive.

Q100 Viscount Colville of Culross: I would like to ask about the disadvantages of using algorithms online. Professor Naughton, you have talked about the embedded bias of language that may not be picked up by machine learning. You have also talked about the lack of context. Could you explain some of your concerns about algorithms? Is it possible

to design an unbiased algorithm?

Professor John Naughton: That question is of the same order of magnitude as whether it is possible to design an unbiased human. We are talking about specific kinds of algorithms, which are machine-learning algorithms. Machine-learning algorithms are basically programs that can convert vast amounts of data into patterns that can be observed or predictions that can be made. What we already know without controversy is that the old rule about garbage in, garbage out applies. Most datasets are not clean; they are coloured in one way or another with all kinds of unconscious and other biases. In those circumstances, there has to be much greater awareness of that. Awareness in the machine-learning community is now pretty widespread. It is very impressive compared with what it was like, say, five years ago.

The problem then arises with the wider community of people who are dazzled by this technology, by which I mean government Ministers, among others, corporate executives and so on. They know nothing about the technology, but they are dazzled by it. We always have to be prepared to apply the standard levels of human scepticism that we should apply to anything. There was a period when that was not happening.

The most spectacular case was when Google revealed that its analysis of queries about flu enabled it to predict flu outbreaks two weeks ahead of the Center for Disease Control and Prevention in Atlanta. This was claimed as a fantastic advantage of big data. It then turned out that Google knew nothing about flu; it was just that its machine-learning algorithms had picked up a pattern. Then there was another kind of flu and it did not work, but that did not stop people extolling the fact that machine learning was the next big thing.

The industry is trying to find ways of dealing with that. One of them is that if you want to do machine learning, on the one hand you have the machine learning doing its stuff; on the other hand, you have another kind of AI, which is effectively a sceptical AI, questioning it. You have a kind of antagonistic approach. There may be technical solutions to some of this, but in the end the problem is that there are always some kinds of biases somewhere in datasets. That is why one has to be very sceptical about using them in an unmediated way as guides to policy or decision-making.

For example, the United States for some years has been using machine learning with considerable and, some would say, impressive success in carrying out analysis to identify targets for drone strikes. That is real life and death stuff, but it has held back from the idea that the algorithm, having identified the target, can then institute the strike. That is the only model we have for the foreseeable future, and if we shift from that we will be in deep trouble.

Viscount Colville of Culross: With that scepticism in mind, should we not allow algorithms to make decisions that affect humans?

Professor John Naughton: At the present stage, no algorithm should be allowed to make a decision that affects the life chances of somebody else, without human oversight or a body that can be held responsible for that decision.

Viscount Colville of Culross: What about much smaller decisions?

Professor John Naughton: I would have to go back to Dr Luger. There are all kinds of areas where low-level decisions are rather helpful. For example, I can ask Google Maps to tell me the best way to get from the Palace of Westminster to King's Cross. It will do that for me and I will follow it through. There is a level at which it seems unproblematic, but there is a level further up about whether your kid gets into a school, or you get a loan or you get parole. Some systems now do that. In that case, I cannot see any justification for having machines make those decisions.

Dr Ewa Luger: To some extent, there are distinctions to be made about algorithms of the type we are talking about. It comes down to algorithmic capacity, which we have heard about, and the algorithm being the black box. There are three broad reasons for that: one is that somebody looking at it does not have sufficient technical literacy; another is that it would take a trained expert to pull apart the algorithm and work out what it was doing; and the final one is that some algorithms are being developed where even a trained expert could not tell you how it reaches its judgment. For that latter class of algorithms, absolutely not. We do not know enough about that yet to deploy it in sensitive contexts.

As Professor Naughton noted, it depends on the context. If it is Google Maps getting you from A to B, it is fine, but deciding whether somebody should get some kind of health intervention is an entirely different proposition, so these things matter. There is a recent example. Durham police are using an algorithm designed at Cambridge called the HART algorithm. It was intended to predict whether a criminal is a high or low risk.

An academic from another institution did an analysis and suggested they ought to remove some of the predictors, such as location and sociodemographic data, because they might result in a prejudiced view of crime and who was involved. Some crimes are more likely to occur in certain types of places, so the algorithm was prejudiced. That is a nice example of where people are working to minimise bias.

The issue with that algorithm is that it is better at predicting things, which humans are not very good at. There are some things on which algorithms could give us much better judgments. It could save money and create better decisions in the health domain. Lots of algorithms are being used there to help identify whether, for example, multiple sclerosis has moved from one stage to the next.

There are some contexts where it is fine, but you absolutely need human oversight at this point in time. It comes down to whether the algorithm is

explainable. If it is a class of algorithm that somebody who is technically literate could understand and an expert could pull apart and tell you how it reaches that judgment, that is one thing. If not, I would warn against it.

Q101 **Baroness Benjamin:** How can we effectively ensure that algorithms are accountable or transparent? Do you think a code of conduct for algorithmic design would help companies to act morally and with integrity and trust, and for that to be embedded in their DNA? Do you think that type of code would do anything to improve accountability?

Dr Ewa Luger: It depends on how the code is enforced or reinforced. There are codes of conduct for lots of main bodies. The Association for Computing Machinery, for example, has a code of conduct and, for engineers, the IEEE has a code of conduct for ethics, but I do not see any real difference in ethical practice. It is important to have codes; we absolutely should, but there needs to be some way that is manifest in the products that are designed, rather than people just agreeing to things that are a little like greenwashing or, in this case, ethic-washing. That is a concern.

There are some developments in the area of explainable artificial intelligence, which is called XAI in the States. Some of the things people are pushing for us to be able to do with algorithms are the answers to a number of questions. Why did it do that? Why did that function occur? Why did it do that and not something else? When has it succeeded and when has it failed? When can the user trust it?

That is about the proportion of accuracy. Is it 80% or 100% accurate, because in the context that absolutely matters? How do I correct an error? DARPA believes that, if you have algorithms that can respond to those questions, you are broadly in the right ballpark.

Professor John Naughton: Codes of practice have a long history and we have had varied experience with them. In general, they are a good thing because they represent a set of aspirations that we would like an organisation or group of workers to adhere to. They are a necessary but not a sufficient condition. I write a newspaper column, so at one level I could be classed as a journalist. Many of you in this House know that the British journalistic industry has various codes of conduct. I invite you to speculate on how effective they have been in relation to tabloid newspapers, for example. Even though the code is there, it does not seem to bite.

Baroness Benjamin: Why?

Professor John Naughton: I do not know. In part, it may be because it is not enforceable in law. That may be the problem, or maybe it is because in a competitive environment ignoring a code may give you a competitive advantage. I am afraid that could be true in this industry, too. Codes are a good thing, because they represent our aspirations, but they are not enough.

The Chairman: There has to be enforcement.

Professor John Naughton: It is code plus enforcement, and, if the code works, you hope you will never have to enforce it.

Baroness Benjamin: What should be the consequences if you do not follow the code?

Professor John Naughton: It depends on the context. It would be very interesting to think of sending the editor of a British tabloid to jail, and what Mr Putin would do with that. There has to be some kind of proportionality, and that is not easy; it varies from context to context.

The Chairman: Proportionality is at the heart of all regulation, is it not?

Professor John Naughton: I am afraid it is, and it is hard.

Q102 **Baroness Kidron:** I am mindful of the hour. I have an enormous question. You might like to answer some of it in writing, retrospectively. One thing that keeps coming up is verification, specifically age verification, and anonymity as a concept online, and the pros and cons of that. There has been a lot of evidence about what privacy should look like. By that, I mean default privacy settings. What policy should there be with regard to anonymity, default privacy and age verification? What are the technical challenges? I understand that all three of those things are quite huge, so feel free to say the top line and then, with the Chairman's permission, perhaps you could write to us about the detail.

Dr Ewa Luger: There is really only one part that I can comment on based on my research, and that is privacy. We have a large problem, in that our notions of privacy alter almost daily. What people expect and what they are prepared to accept shift in accordance with the systems they use, and that is problematic.

A number of the surveys that have been conducted show that people are concerned about their privacy, but then we see the behaviour-intention gap; they will merrily share any and all data to get the shiny thing at the end of the road. The systems are designed to hit the reward centres of our brains, so that is fairly unsurprising. The whole concept needs to be broken down to some extent; it comes down to what is good for people versus their expectations, and they are not necessarily the same thing. I suppose we could get into a kind of paternalistic role.

Baroness Kidron: Can I push you a bit on that? Their expectations in the current context may be out of kilter with their expectations in any other context, in that they are designed in.

Dr Ewa Luger: Absolutely. I often hear the phrase "Privacy is dead". I apologise in advance for giving a terrible example. We say that and mean it, broadly and socially, but everybody shuts the toilet door and gets dressed and undressed in private. We have expectations of privacy; it is just that when we talk about it our default idea is, "If I've got nothing to fear, I've nothing to hide", but we all know that if there is a context shift,

that becomes untrue. The political climate changes, or suddenly some aspect of your character becomes linked to terrorism, and it becomes problematic for you as an individual, because the monetisation model relies on our being profiled and those profiles being sold. Privacy is important because it protects us, but online it is definitely a different model from our social model.

Professor John Naughton: I agree. Privacy is one of those strange contested concepts; it is ambiguous and it is very hard to get a grip on it. One simple way of trying to clarify it a bit relates to the deal you strike with an internet company by clicking "Accept" on its end-user licence agreement. In those circumstances, if you are a user of web mail, particularly Google's Gmail, you make a contract that says, "Google, you can read my mail", so in that sense you are handing over your privacy. That is fine; you make that decision. You and Google are treating privacy as if it were a private good that you can transact in return for services. That is fine for a while. Then you send a message to somebody who does not approve of that and has definitely not signed up to any of it. Replying to your Gmail account, that person emails you. Suddenly, you have compromised their privacy. The point is that privacy is in some ways both a very private thing and a social thing; it is almost an environmental good in many cases. The technology has blurred all of that and made it very, very difficult.

On verification, were you thinking of age verification?

Baroness Kidron: I was in this specific case, because we have laws coming in that require it. It is a slightly difficult area, because the law went ahead of the technology.

Professor John Naughton: It is difficult. The key thing about it, which is worrying, is that in general one should not enact laws that one is not going to enforce, because you undermine the law, as well as everything else. That is really tricky. There are some very worrying developments—for example, the increasing use of tablets and smartphones as a way of amusing very small children. That is really worrying. There is some evidence about what babies need, and it is not a screen; it is a human face. I think that in some authoritarian societies it will become a crime to let your toddler use a smartphone. In our societies, we will not do that. I am afraid you have opened up a huge can of worms and I have no obvious answers.

The Chairman: In that case, we will draw the session to a close. It has been fascinating for us. We noted your warning at the beginning about the breadth of the subject and loose use of the word "internet", and we understand the point you are making. This is a broad inquiry, which will take a long time. I think you have noted our focus on the business models, not just on immediate harms and abuses. Your evidence in that regard has been very helpful and interesting.

I thank both of you for taking the time to be with us today. Please feel free to write to us if you want to elaborate on anything. If some

interesting reading crosses your desks over the next months that might be of interest to the Committee, please send it to our clerk. Thank you for being with us today.