

SELECT COMMITTEE ON ARTIFICIAL INTELLIGENCE**COLLATED ORAL EVIDENCE VOLUME****Contents**

| | |
|--|-----|
| Alan Turing Institute and Professor Chris Hankin – Oral evidence (QQ 143–152) | 4 |
| Dr David Barber, Digital Catapult and NVIDIA – Oral evidence (QQ 38–45) | 20 |
| Jeremy Barnett, Professor Chris Reed and Professor Karen Yeung – Oral evidence (QQ 29–37) | 38 |
| Miles Berry, Graham Brown-Martin and Professor Rosemary Luckin – Oral evidence (QQ 181–189)..... | 57 |
| Professor Nick Bostrom, Professor Dame Wendy Hall and Professor Michael Wooldridge – Oral evidence (QQ1–8)..... | 76 |
| Graham Brown-Martin, Professor Rosemary Luckin and Miles Berry – Oral evidence (QQ 181–189)..... | 93 |
| Dr Mercedes Bunz, UK Information Commissioner Elizabeth Denham and Dr Sandra Wachter – Oral evidence (QQ55–64) | 94 |
| Eileen Burbidge MBE, MMC Ventures and Project Juno – Oral evidence (QQ 46–54)..... | 112 |
| Rory Cellan-Jones, Sarah O'Connor and Andrew Orłowski – Oral evidence (QQ 9–17)..... | 129 |
| CIFAR – Oral evidence (QQ 172 – 180)..... | 145 |
| Citizens Advice and Competition and Markets Authority – Oral evidence (QQ 85–94)..... | 153 |
| Competition and Markets Authority and Citizens Advice – Oral evidence (QQ 85–94)..... | 172 |
| Digital Catapult, NVIDIA and Dr David Barber – Oral evidence (QQ 38–45) ... | 173 |
| Dyson, Fujitsu and Ocado – Oral evidence (QQ 105–115) | 174 |
| Professor David Edgerton, Professor Peter McOwan and Professor Sir David Spiegelhalter – Oral evidence (QQ213–223) | 190 |
| Fujitsu, Ocado and Dyson – Oral evidence (QQ 105–115) | 213 |
| Future Advocacy, Professor Dame Henrietta Moore and Professor Richard Susskind OBE – Oral evidence (QQ 95–104)..... | 214 |
| German Research Centre for Artificial Intelligence (DFKI) – Oral evidence (QQ 163 – 171)..... | 234 |
| Professor Dame Wendy Hall, Professor Michael Wooldridge and Professor Nick Bostrom – Oral evidence (QQ1–8)..... | 241 |

| | |
|--|-----|
| Professor Chris Hankin and the Alan Turing Institute – Oral evidence (QQ 143–152)..... | 242 |
| Dr Hugh Harvey, National Data Guardian for Health and Care Dame Fiona Caldicott and NHS Digital – Oral evidence (QQ 128 – 142) | 243 |
| HM Government – The Rt Hon Matt Hancock MP, Minister of State for Digital, Department for Digital, Culture, Media and Sport (DCMS) and the Rt Hon the Lord Henley, Parliamentary Under Secretary of State, Department for Business, Energy and Industrial Strategy (BEIS) – Oral Evidence (QQ 190 – 200) | 261 |
| HM Government –The Rt Hon the Lord Henley, Parliamentary Under Secretary of State, Department for Business, Energy and Industrial Strategy (BEIS) and The Rt Hon Matt Hancock MP, Minister of State for Digital, Department for Digital, Culture, Media and Sport (DCMS) – Oral Evidence (QQ 190 – 200) | 283 |
| Dr Julian Huppert, PHG Foundation and Understanding Patient Data, Wellcome Trust – Oral evidence (QQ 116–127) | 284 |
| IBM, Sage and SAP – Oral evidence (QQ 76–84) | 307 |
| IEEE-Standards Association and Professor Alan Winfield – Oral evidence (QQ 18 – 28)..... | 328 |
| Professor Rosemary Luckin, Miles Berry and Graham Brown-Martin – Oral evidence (QQ 181–189)..... | 346 |
| Major Kitty McKendrick, Professor Noel Sharkey, Mike Stone and Thales Group – Oral evidence (QQ 153–162) | 347 |
| Professor Peter McOwan, Professor Sir David Spiegelhalter and Professor David Edgerton – Oral evidence (QQ213–223)..... | 367 |
| MMC Ventures, Project Juno and Eileen Burbidge MBE – Oral evidence (QQ 46–54)..... | 368 |
| Professor Dame Henrietta Moore, Professor Richard Susskind OBE and Future Advocacy – Oral evidence (QQ 95–104) | 369 |
| National Data Guardian for Health and Care Dame Fiona Caldicott, NHS Digital and Dr Hugh Harvey – Oral evidence (QQ 128 – 142)..... | 370 |
| NHS Digital, Dr Hugh Harvey and National Data Guardian for Health and Care Dame Fiona Caldicott – Oral evidence (QQ 128 – 142)..... | 371 |
| NVIDIA, Dr David Barber and Digital Catapult – Oral evidence (QQ 38–45) ... | 372 |
| Ocado, Dyson and Fujitsu – Oral evidence (QQ 105–115) | 373 |
| Sarah O'Connor, Andrew Orłowski and Rory Cellan-Jones – Oral evidence (QQ 9–17)..... | 374 |
| Open Data Institute, Open Rights Group and Privacy International – Oral evidence (QQ 65–75) | 375 |
| Open Rights Group, Privacy International and Open Data Institute – Oral evidence (QQ 65–75) | 390 |
| Andrew Orłowski, Rory Cellan-Jones and Sarah O'Connor – Oral evidence (QQ 9–17)..... | 391 |

| | |
|---|-----|
| Dr Jérôme Pesenti – Oral evidence (QQ 201–212) | 392 |
| PHG Foundation, Understanding Patient Data, Wellcome Trust and Dr Julian Huppert – Oral evidence (QQ 116–127) | 407 |
| Privacy International, Open Data Institute and Open Rights Group – Oral evidence (QQ 65–75) | 408 |
| Project Juno, Eileen Burbidge MBE and MMC Ventures – Oral evidence (QQ 46–54)..... | 409 |
| Professor Chris Reed, Professor Karen Yeung and Jeremy Barnett – Oral evidence (QQ 29–37) | 410 |
| Sage, SAP and IBM – Oral evidence (QQ 76–84)..... | 411 |
| SAP, IBM and Sage – Oral evidence (QQ 76–84)..... | 412 |
| Professor Noel Sharkey, Mike Stone, Thales Group and Major Kitty McKendrick – Oral evidence (QQ 153–162) | 413 |
| Professor Sir David Spiegelhalter, Professor David Edgerton and Professor Peter McOwan – Oral evidence (QQ213–223)..... | 414 |
| Mike Stone, Thales Group, Major Kitty McKendrick and Professor Noel Sharkey – Oral evidence (QQ 153–162) | 415 |
| Professor Richard Susskind OBE, Future Advocacy and Professor Dame Henrietta Moore – Oral evidence (QQ 95–104)..... | 416 |
| Thales Group, Major Kitty McKendrick, Professor Noel Sharkey and Mike Stone – Oral evidence (QQ 153–162) | 417 |
| UK Information Commissioner Elizabeth Denham, Dr Sandra Wachter and Dr Mercedes Bunz – Oral evidence (QQ55–64) | 418 |
| Understanding Patient Data, Wellcome Trust, Dr Julian Huppert and PHG Foundation – Oral evidence (QQ 116–127)..... | 419 |
| Dr Sandra Wachter, Dr Mercedes Bunz and UK Information Commissioner Elizabeth Denham – Oral evidence (QQ55–64)..... | 420 |
| Professor Alan Winfield and IEEE-Standards Association – Oral evidence (QQ 18 – 28)..... | 421 |
| Professor Michael Wooldridge, Professor Nick Bostrom and Professor Dame Wendy Hall – Oral evidence (QQ1–8)..... | 422 |
| Professor Karen Yeung, Jeremy Barnett and Professor Chris Reed – Oral evidence (QQ 29–37) | 423 |

Alan Turing Institute and Professor Chris Hankin – Oral evidence (QQ 143–152)

Evidence Session No. 15

Heard in Public

Questions 143–152

Tuesday 28 November 2017

[Watch the meeting](#)

Members present: Lord Clement-Jones (The Chairman); Baroness Bakewell; Lord Giddens; Baroness Greender; Lord Hollick; Lord Levene of Portsoken; The Lord Bishop of Oxford; Viscount Ridley; Baroness Rock; Lord St John of Bletso; Lord Swinfen.

Examination of witnesses

Dr Mark Briers and Professor Chris Hankin.

Q143 **The Chairman:** Good afternoon and a very warm welcome to our witnesses: Dr Mark Briers, who is the strategic programme director for defence and security of the Alan Turing Institute, and Professor Chris Hankin, who is the co-director of the Institute for Security Science and Technology, Imperial College London. This is the 15th formal evidence session for the inquiry and it is intended to help the Committee discuss the potential misuse of AI and the implications for cybersecurity.

I am afraid I have a little rubric that I need to read through at the beginning of every evidence session. The session is open to the public and a webcast of the session goes out live, as is, and is subsequently accessible via the parliamentary website. A verbatim transcript will be taken of your evidence and will be put on the parliamentary website. A few days after this evidence session, you will be sent a copy of the transcript to check for accuracy, and we would be grateful if you could advise us of any corrections as quickly as possible. If, after this session, you wish to clarify or amplify any points made during your evidence, or have any additional points to make, you are very welcome to submit supplementary written evidence to us. Would you like to introduce yourselves for the record?

Professor Chris Hankin: I am professor of computer science at Imperial College London and co-director of the Institute for Security Science and Technology.

Dr Mark Briers: Good afternoon. I am the programme director for defence and security at the Alan Turing Institute. My research interests lie at the intersection of artificial intelligence and cybersecurity.

Q144 **The Chairman:** Thank you. I will start with something pretty broad and general, especially in terms of the timing. What does artificial intelligence mean for cybersecurity today, and how is this likely to change over the next 10 years? We are looking in particular at whether, and how, artificial intelligence will impact on conventional cybersecurity today. Does it facilitate new kinds of cyberattacks, how much does that alter the risk profile, and where can AI help?

Professor Chris Hankin: When I think about artificial intelligence in the context of cybersecurity today, I think mainly about machine learning rather than broad artificial intelligence, where my own team at Imperial and many others across the globe have had a great deal of success in using machine learning to analyse network traffic and spot anomalous things happening in that traffic. In fact, there is a UK company, which you may have spoken to or heard of during your hearings, Darktrace, which has made a very successful global business out of machine learning to do that. From a defensive point of view, that would be the main application of AI at the moment in cybersecurity. I heard yesterday, in fact, of people who have been developing chatbots to engage in conversations with phishing attackers to frustrate them, to a certain extent, in their attacks.

The Chairman: So Darktrace is literally finding the source of a cyberattack, is it?

Professor Chris Hankin: It is analysing network traffic. They basically install a monitor in companies or an individual system that learns what “normal” looks like, so to speak, over a very short period of time and can spot if something is abnormal, which could be indicative of a cyberattack. So that is Darktrace technology. It is a very exciting technology, and they have made a great commercial success out of it, but there are still some open research challenges to reducing the false positive signals, for example, which might come out of that sort of system.

The Chairman: Is that still partly under your wing or a fully independent start-up?

Professor Chris Hankin: There is still quite a lot of academic research activity across the world looking at different approaches to machine learning that might be able to give more accurate signals about what is going on.

I also wanted to mention a competition that was held in the States that came to fruition in August 2016 and was about automatic defensive systems. It was about programmes that could understand when they were under attack and take measures to repair themselves and mitigate against the attack. Over, say, a 10 to 15-year horizon, we could be looking at that sort of technology being lifted to the level of systems that can understand that they are being threatened in some way and that can take action to repair the damage that that threat might be causing them.

The Chairman: So there are a lot of very active developments in the AI field.

Professor Chris Hankin: I believe so. Cybersecurity is my area of expertise. I am sure Mark will have a more informed view of the artificial intelligence world.

The Chairman: But these are AI applications, which have implications for cybersecurity.

Professor Chris Hankin: Yes.

The Chairman: Is that why you mentioned chatbots and so on?

Professor Chris Hankin: Yes, that is correct.

Dr Mark Briers: I would echo Professor Hankin's views that artificial intelligence, specifically machine learning, is being applied in a defensive context as we sit here today. If I were to look 10 years hence, I would argue that there will be more sophisticated artificial intelligence in the offensive space. In my experience in industry as a cybersecurity specialist over the past few years, we have seen no evidence of artificial intelligence malware or software encroaching on corporate networks. I believe that will change over the next five to 10 years and that the offensive weapons used in cyberattacks will become more artificially intelligent.

The Chairman: That in itself is a very interesting comment. With a number of viruses such as WannaCry and that kind of thing, do you mean that there was no AI component to them and that they were fully human interventions?

Dr Mark Briers: More or less. It depends on how you characterise and define artificial intelligence. By my definition, I certainly would not say that they were artificial intelligence in any meaningful manner.

The Chairman: You are aware of the developments that Professor Hankin talked about and you think they will make quite an impact themselves.

Dr Mark Briers: Indeed. From a research perspective it is very exciting, but obviously from a risk mitigation perspective it is quite daunting, so yes, I agree with Professor Hankin.

Lord St John of Bletso: If I could ask a question slightly beyond my pay grade, we read about D-Wave computers and their ability to solve large data analytics and about IBM's Watson supercomputer and the future of quantum computing. Is this the next big potential opportunity to crack the cybersecurity threats that we are facing?

Professor Chris Hankin: D-Wave is a kind of quantum computing, and it can still only do the same sorts of things that the computers we know and love today can do, but there are certain tasks that it can do much faster. One of the particular things that quantum computers can do very fast is the factorisation of large numbers. Modern cryptography is built upon factorising very large numbers into their prime factors. This is a very complex thing to do with a modern computer, but quantum computers can do it much faster, so the modern popular approaches to cryptography are potentially undermined by the emergence of quantum computers.

I suspect that we are several years away from that really becoming a threat. I am comforted slightly by the knowledge that there are many very able research groups and people in government looking at what these days tend to be called “quantum-safe approaches” to cryptography. You do not have to build your cryptographic algorithms on the basis of using factorisation; there are other bases that you can use for crypto, some of which will be less amenable to being solved by quantum computers. It is an issue that we have to be aware of, but responsible people, both in the research field and in government, are paying great attention to this potential long-term threat.

Dr Mark Briers: I have nothing to add to that excellent answer.

The Chairman: Lord St John is well ahead of the pack here, as always.

Q145 **Lord Giddens:** Will only state-sponsored hackers have the means to deploy AI in cyberattacks, or is there a risk that AI-enabled cyberattacks will be “democratised” in the near future? I do not know whether this is beyond your brief, but could you comment on the growing literature on threats to democracy from the use of chatbots and algorithms that have definitely intruded very deeply into the political process?

Dr Mark Briers: In the short to near term, on the basis of my previous answer I would expect artificially intelligent cyberoffensive weapons to emerge from the state-sponsored sector. If you look back in history 10 years ago, the types of threat that are prevalent today and that are democratised in some sense would arguably have come from the state-sponsored sector. Using history as an indicator of what is likely to happen in the future, I suspect we will see those types of artificially intelligent cyberweapons being available to a wider audience in 10 years’ time, so yes, I see AI cyberweapons being democratised.

With respect to chatbots and interference in the political process, it is clearly a serious problem and we need to do more research to understand the threat landscape, how malicious actors are manipulating the information space and how we can counter that manipulation such that we can present the messages that we need to present as a democratic society.

Professor Chris Hankin: I broadly agree with what Dr Briers has said. At the moment, with cybersecurity it is becoming much more difficult to differentiate between state actors and organised crime. The sorts of techniques that those two groups are using to mount their attacks are becoming more and more similar. The weapons that they use are also becoming available through things like the dark web, so it is becoming much more difficult, now and going forward, to differentiate between the different styles of attackers that we are having to defend against. That is my answer to the first part of the question.

On the second part of the question, about chatbots and the threat to democracy, that is a serious issue. It may be beyond the brief that I have been thinking about in preparing for this meeting.

Lord Giddens: We might need some partly technical solutions to it, because they cannot be just political solutions.

Professor Chris Hankin: No. Someone somewhere needs to think very seriously, as I think Mark was hinting at, about the counter-narratives so that the messages that we want to get across appear above the noise that the chatbots and so on might be creating.

The Chairman: Can I come back to you on your response? You agreed with Lord Giddens and said that, yes, there was the threat of a black hat issue? Is that the result of leakage of knowhow? Is it that the skills required are getting more common? Is it the fact that the investment required is not very great? Is it a combination of all those things? What creates the fertile ground for the hackers using AI?

Professor Chris Hankin: Certainly, for those who have been educated in coding practices and rudimentary computer science, many of the modern programming languages that people use to code up artificial intelligence applications have libraries that have, for example, machine learning functions within them, so one does not necessarily need to have a very deep understanding of machine learning to be able to use it in some way in a system that you are constructing.

Viscount Ridley: Dr Briers, you said, going back 10 years, that you could see how state-sponsored software ended up in private hands and being used for malice. I remember 10 or 15 years ago some pretty blood-curdling presentations about how viruses were going to make computing impossible in the near future, and that the bad guys were going to win. That did not happen, did it, and we stayed one step ahead of that. We all have virus problems still, as WannaCry exemplifies, but it has not been as catastrophic as some said. Can we learn any lessons from that, or am I being Panglossian?

Dr Mark Briers: No, there are lessons to learn from that. I suspect that communications via appropriate government bodies should be congratulated in some senses for moving the cybersecurity posture of some of our major industries and critical national infrastructure to avoid some of the problems that we quite rightly could have faced during that time. It is great to see the Government continue to invest in cybersecurity with the National Cyber Security Centre and other organisations like that to put cybersecurity at the forefront of the UK and make the UK a secure place.

Viscount Ridley: Was it the Government or was it the Apple Corporation that helped me avoid that fate?

Dr Mark Briers: Based on my expertise and opinion, I suspect it is a bit of both. It will be political pressure to encourage organisations, such as Apple, to patch their systems and so on.

Q146 **Lord Swinfen:** Do AI researchers need to be more aware of how their research might be misused and consider how this might be mitigated before publishing? Are there situations where researchers should not publish where there is a high risk of misuse? Should the Government consider mechanisms, either voluntary or mandatory, to restrict access in exceptional cases in a similar way to the defence advisory notice system for the media? Do you think there should be a code of ethics?

Dr Mark Briers: In short, yes, I do. I believe there is an ethical responsibility on all AI researchers to ensure that their research output does not lend itself to obvious misuse and to provide mitigation, where appropriate. If you look at the analogous situation of the 3D printer, the manufacturers or the designers of it perhaps did not envisage somebody producing a 3D-printed weapon downstream. If we use that as an analogy, should the manufacturers and designers of 3D printers have considered this, and should we have prevented 3D printing from ever making it into the marketplace, even though there are fantastic medical advances that could surface through 3D printing? We need to provide those principles and ethical guidelines, but they need to be principles and guidelines as opposed to definitive rules that one has to follow, so a judgment needs to be made against each of these different algorithms and the capabilities that they offer up.

Professor Chris Hankin: I broadly agree. A number of the large vendors now offer bug bounties to people to disclose vulnerabilities that they discover in systems so that patches can be issued before the vulnerabilities get into the public domain. That is one aspect where there is a potential financial incentive for people to be responsible about disclosure.

It is important, when we educate people in cybersecurity, that we impart some ethical values to them. Indeed, a lot of the cybersecurity work currently going on here in the UK is funded through various schemes that the National Cyber Security Centre has set up. The research institute, which I lead, has no issue with passing papers through contact points within the centre so that they can be checked before they are published. By and large, researchers are fairly responsible about the way they disclose things, but there is the danger that if new vulnerabilities are discovered and published without any of those checks, it gives the attackers and the hackers something to exploit.

The Chairman: Are we in anywhere near a common code of ethics that could be accepted in research in this field?

Professor Chris Hankin: A number of the professional institutions have ethical codes, and if people are registered with those institutions they ought to be living by those ethical codes, but there is no uniform approach to that. If the Alan Turing Institute is to become the AI centre for the UK, maybe it should address promulgating such a code.

The Chairman: Do you think that getting that code up and running and having general acceptance could be part of your agenda?

Dr Mark Briers: I certainly think it could be part of our agenda, and I am sure that Sir Alan Wilson and other colleagues at the Turing would consider that and hope to produce something of that kind to support researchers.

Lord Swinfen: Should workers at certain levels be security vetted?

Dr Mark Briers: In certain circumstances for certain applications, almost certainly, but I suspect it is not practical to security-clear all the research community or a large portion of it who would develop algorithms of this kind. One has to ensure that there are sufficient guidelines, ethical

principles and support mechanisms, communication mechanisms, et cetera, to ensure and encourage the appropriate and ethical publication of research as opposed to validating and checking everybody. Sadly, I do not believe that is a practical solution.

Q147 The Lord Bishop of Oxford: This question is on adversarial AI, which I gather is researchers attempting to fool AI systems into making incorrect classifications or decisions. How much of an issue are recent developments in that field of adversarial AI for the wider deployment of AI systems? Do you see compulsory stress-testing as part of the future?

Professor Chris Hankin: We have been doing some work on using adversarial AI to see how possible it is to train an attacker to evade the state-of-the-art classifiers that we have been developing on the other side of our research activity. It is certainly true that one can use adversarial nets and get very high success rates in learning. If you can get into the right part of the system, you can learn a lot about what the classifier might be doing and introduce sufficient noise into your attacks, such that it evades detection by the classifier. We have some quite interesting results in that. The message I take from that is that, at the moment certainly, AI is not the only answer we should be thinking about for defending our systems.

I will give you a short story about the Stuxnet malware that was used to delay the Iranians in their uranium enrichment process. The attack was essentially a physical attack, mounted through cyber, on the centrifuges that were being used to enrich the uranium, and in one version at least it caused the rotor blades in the centrifuges to spin at very high speeds. You might have been able to detect that attack by looking at some network traffic and seeing what was happening with the control systems, but certainly if you had been standing anywhere near the centrifuges you would also have had a physical signal that something was going wrong. At the moment, in the state we are in we have to use all sources of information that we can to decide what is going on in a system, because AI is not the only answer. In that setting, maybe the adversarial net approach might have enabled the attackers to get round the AI detector, if they had had such a thing, but the noise from the centrifuges probably would have given them away.

Dr Mark Briers: I agree with Professor Hankin. In the cybersecurity industry, there is a large group of individuals known as “penetration testers”, whose job essentially is to try to ethically hack into an organisation’s network and look for vulnerabilities with the intention of trying to secure those vulnerabilities. I see there being a large marketplace in the short to near term in AI ethical hackers, if that is the correct phrase, and I hope to see the UK leading in this field so that we develop that marketplace and lead it internationally as UK plc.

The Lord Bishop of Oxford: Would the people doing this typically have doctorates, or would they be, as in science fiction films, very bright teenagers who are going rogue?

Dr Mark Briers: I suspect as we sit here today that it would be people with doctorates and beyond. However, with all these kinds of technologies there is a democratisation effect—the open sourcing of the

technology in relatively straightforward-to-use language that makes it easier for the kind of individual you suggest to use.

Q148 **Lord Levene of Portsoken:** How prepared is the UK for the impact of artificial intelligence on cybersecurity?

Professor Chris Hankin: As Mark said in answer to an earlier question, the UK has done a lot of work over the last number of years to produce some very good advice to companies, government and private citizens about how to protect themselves. The sorts of attacks that we may be talking about, which are AI-based, will at the moment be probably no different from the sorts of attacks that they might be expecting to see from human attackers. The advice that has come out about cyber hygiene such as keeping software up to date, having appropriate anti-virus software, not sharing passwords with people, all these sorts of controls that have been devised, are very effective in reducing the impact of cybersecurity at the moment when they are properly applied. Unfortunately, the cyberattacks that have been most prominent in the news over the last year—WannaCry, NotPetya, Equifax—have all been the consequence of people running unpatched software, which we have known for many years is leaving vulnerabilities exposed that attackers can exploit.

The Chairman: So it is nothing to do with AI at this point.

Professor Chris Hankin: Nothing to do with AI at all, I would say.

Dr Mark Briers: As Professor Hankin suggested, the UK took some good initial steps with the creation of the National Cyber Security Centre approximately a year ago, and indeed with the industrial strategy White Paper, which was released yesterday. Specifically, the office for AI mentioned six priority business sectors, the first of which is cybersecurity, so we are taking some good initial steps in positioning the UK as a global centre of cybersecurity excellence.

However, as we have alluded to during the discussion so far, there is still the potential for significant impact with respect to cybersecurity from artificial intelligence and the need to ensure that the guidelines that we mentioned, the principles under which researchers publish research and the general democratisation of artificial intelligence, are managed carefully such that we do not get into a situation where all the great work that the UK has done so far is overtaken by advancements in technology.

Lord Levene of Portsoken: Can the bad guys leapfrog where we have got to and move on to the next level?

Dr Mark Briers: Indeed, I believe they can, which is why we need to continue to invest and put cybersecurity at the forefront of what we do in the UK.

Lord Levene of Portsoken: It is never ending.

Dr Mark Briers: Yes, I am afraid so.

Lord St John of Bletso: I understand about the guidelines, and large organisations are certainly well prepared for cyberattacks, but the SMEs are ill-prepared. What is being done to help small and medium-sized

enterprises, which are more vulnerable than they have ever been to cyberattacks?

Professor Chris Hankin: I applaud the Cyber Essentials scheme, which was designed specifically for SMEs of a certain size and contains very good advice on patching, firewalls, good configuration of systems and so on to protect them against the commodity-style attacks that hitherto have been the main kinds of attacks that they have seen, the sorts of attacks that people can readily download from the internet. SMEs will typically experience attacks of that kind rather than attacks from state actors or serious organised crime. Cyber Essentials is an excellent piece of advice that has come out of government to help those companies.

The only caveat to that would be that if you are a micro-SME with one or two people, you will not have a full-time IT manager, so even Cyber Essentials may be a bit beyond your reach. Indeed, in my team at Imperial we have done some research to look at how we might add a slight nuance to the advice of cybersecurity that would fit within the budget of such micro-SMEs.

Dr Mark Briers: I suspect that the majority of micro-SMEs are utilising cloud-based services for the provision of computing infrastructure. Micro-SMEs of that kind could benefit from the security investments of Microsoft, Amazon and the like in cloud security and therefore stay secure by virtue of using a democratised system of that kind.

The Chairman: Do you think the National Centre for Cyber Security, which was only formed relatively recently, is doing enough?

Professor Chris Hankin: It has got off to an excellent start. Like so many things, in the end it will have to be a partnership between industry, academia and government. The Industry 100 scheme that it has introduced, which is to encourage industry to second people into their new building in Victoria, is an excellent initiative. It is too early to say how effective that will be, but the NCSC is already working very actively in partnership with academia through its 14 academic centres of excellence in cybersecurity research.

The Chairman: Is it paying sufficient attention to AI?

Professor Chris Hankin: This is an area that, with the new industrial strategy and so on, we may see more of in the future.

The Chairman: It is interesting that there are two grand challenges close together.

Professor Chris Hankin: Yes. I have been quite involved in some of the certification of a Master's in cybersecurity, and the depth of knowledge in machine learning is not one of the criteria that we have been looking at in that certification process. There is a slight danger that we end up with an artificial intelligence silo and a cybersecurity silo when they need to be talking to one another.

Dr Mark Briers: In one initiative that is trying to impact the point that Professor Hankin just made about two communities, the Turing Institute is bringing together cybersecurity researchers with artificial intelligence

experts to try to cross-fertilise ideas from different disciplines and to ensure that both can communicate in a common language.

The Chairman: So you accept full responsibility for driving that part of the agenda forward?

Dr Mark Briers: We certainly will push as hard as possible and we accept full responsibility.

The Chairman: Excellent, with all the resources necessary.

Dr Mark Briers: Indeed.

Q149 **Lord St John of Bletso:** With the GDPR coming into force in May next year and the Data Protection Bill, on which our Lord Chairman is playing a very active role, will the law be able to adequately prosecute those who misuse AI for criminal purposes?

Professor Chris Hankin: The law will almost inevitably be in a situation of catch-up with any technological development. What tends to happen is that the law makes do and mends. A few years ago, I led an expert group for the Government Office for Science looking at the future of identity, which also touched on some of these data protection issues. At that time, we took expert evidence from lawyers about the state of the law with respect to identity theft, and it turned out then, and I am sure it is still true now, that the majority of prosecutions for identity theft ended up being prosecuted either under computer misuse or were traditional fraud prosecutions. That is possibly the way the law will always have to be, because it is not possible, given the process of introducing new laws, to stay abreast of the latest technological developments.

Dr Mark Briers: I believe there are significant positives in both the GDPR and the Data Protection Bill with respect to cybersecurity. I am particularly comforted by the breach notification section, because that will have a positive impact on organisations' cyber incident response plans and will improve the majority. If that only results in organisations proactively patching their systems because of this law, which will exist in the near future, as we discussed earlier, that will help to prevent things such as WannaCry ever becoming as major a problem as it did previously. The GDPR certainly has some very positive components to it with respect to cybersecurity.

On my reading of the GDPR and the Data Protection Bill, I do not believe they cover the complete potential sets of misuse of artificial intelligence, not least because it is very difficult to envisage artificial intelligence research output in five years' time and therefore what the applications and the potential misuse cases of that technology could be. I would echo Professor Hankin's comments. To an extent, the law will be playing catch-up, but with sufficient research and sufficient interaction between the law-makers and the research community in understanding the strategic direction of AI, we can hope to stay at least present with research, if not slightly ahead.

Lord St John of Bletso: Professor Hankin, you spoke about the advances in technology. Somewhat unrelated to that, we are on the verge of driverless cars, and concerns have been expressed about the

misuse of drones for military purposes. To what degree do you think that cybercrime law should be tightened to cover these cases?

Professor Chris Hankin: That is certainly quite a long way out of my comfort zone. I have looked at some of the discussion that you have had about driverless cars and drones, and certainly the law already has provision for restricting the flight of drones in sensitive areas, but of course there will be many more of them in the future and they could be targets for attack by criminals and so on. Given that we know that that technology is about to be with us, one could consider developing new laws specifically in those areas. I do not think I am necessarily competent to tell you what those laws should be.

The Chairman: Do you both think specifically that legislation such as the Computer Misuse Act would need updating, irrespective of the data issues?

Dr Mark Briers: Like Professor Hankin, I am no expert in this area, so it is outside my comfort zone.

The Chairman: That is all right. We like people taking a fly.

Dr Mark Briers: I suspect, based on my limited understanding, that the Computer Misuse Act would need updating to reflect the changing technological landscape. My limited experience in this law-making space suggests that those kinds of law would inevitably need to be updated to reflect changes.

The Chairman: It is amazing to think that it is a 1990 Act, is it not?

Dr Mark Briers: Yes.

The Chairman: If it does not need updating, it must have been an incredibly prescient piece of legislation.

Dr Mark Briers: Indeed, so that would lead me to conclude that it will require updating at some point in the near future.

The Chairman: Do you agree with that, Professor Hankin?

Professor Chris Hankin: You asked us to go out on a limb a little bit.

The Chairman: We like taking you out of your comfort zone a bit.

Professor Chris Hankin: I would suggest that it certainly needs updating if it were last passed through Parliament in 1990, and I would like to see it being developed in the direction of saying something about cyber in the context of cyber physical systems. After all, that is what drones and autonomous vehicles are, and suddenly we are on the brink of launching a whole raft of these different systems, which are not only cyber-enabled but have a definite impact on our physical environment. The Computer Misuse Act probably says nothing about that area at the moment, because I do not think that in 1990 we had any concept of this link between cyber and physical, but it should have something to say about it.

The Chairman: That is very interesting. That might be the point at which you would insert a new regulation or whatever.

Professor Chris Hankin: Yes.

Q150 **Baroness Grender:** How can personal data be effectively secured against misuse, especially given the potential conflict between secure and open data? Does the increasing availability of AI have implications for securing this data? Also, how does AI affect the security of anonymised datasets, and is there a level of anonymisation that is good and secure enough?

Dr Mark Briers: There are numerous ways of answering your excellent question. With respect to the potential conflict between secure and open data, there is a class of techniques known as differential privacy-based techniques, which essentially allow noise to be added to data such that it becomes more and more difficult, as you add more and more noise to the data, to de-anonymise the data in some respects. Therefore, as you add more noise, the data is mathematically semi-guaranteed to become more secure. That class of techniques is emerging from academia, and they are being utilised in some instances to protect individuals' right to privacy. I hope that continues and that the research in that domain continues.

With respect to the increasing availability of AI having implications in this area, as the sophistication of AI increases, inevitably the ability of AI systems to correlate across most data sources, to pull in contextual information, to bring in human expertise and to combine it all in a statistically rigorous manner lends itself to being able to de-anonymise data and misuse data in ways in which we would not want data to be misused. Therefore, I see it as being a major risk and something that research will help to counter as we progress forward.

Professor Chris Hankin: I am not sure I have anything very insightful to add to that excellent answer from Dr Briers. Anonymisation is a significant research challenge at the moment: doing it in such a way that it is not possible to de-anonymise the data afterwards. Of course, the sort of AI that we have mainly been thinking about during our answers has been machine learning, which is very data hungry. If you want good machine-learning algorithms, it is essential that you have good data to train the algorithms against. There is a demand for data coming from the AI side, but that has to be done in a way that protects the privacy of individual citizens and so on. I have not studied the latest draft of the Data Protection Bill in great detail.

The Chairman: You are one of the lucky ones, in that case.

Professor Chris Hankin: I believe there is some provision for people to have control over their data in the context of algorithmic decision-making, which relates to Article 22 of the GDPR, but there are other provisions in our Data Protection Bill that go beyond that and give an element of protection to individual data.

The Chairman: Some of your colleagues in Imperial are rather concerned about one of the clauses in the Bill relating to criminalising re-identified data. Clearly, there are public policy issues on both sides here; you do not want to discourage research, and on the other hand you want to discourage re-identification. Where do you stand on that, without dropping your colleagues in it, Professor Hankin?

Professor Chris Hankin: The state of the art in anonymisation at the moment is such that one could inadvertently de-anonymise data, and I

am not sure where the Bill would stand on that. The point about de-anonymisation is that it will occur if you sometimes link data from different datasets. Inevitably, it will highlight a small group of individuals who might have provided the data that had previously been anonymised. That could happen in a completely inadvertent way. If the people who did that were criminalised, that would be unfortunate. This is probably an issue of the law trying to be slightly ahead of the technological curve in that if we do not have perfect anonymisation, and that is an example where we do not because it is possible to de-anonymise inadvertently the data—

The Chairman: Is it too far ahead?

Professor Chris Hankin: Yes. Unless we can be absolutely sure that data can be perfectly anonymised, it could be de-anonymised inadvertently. It would be a great miscarriage of justice if people were criminalised as a result of that.

Lord Swinfen: Dr Briers, what is the noise that is added? Is it music? Is it mush? If you can add the music or the noise, surely someone else can take it out.

Dr Mark Briers: Indeed, there is a kind of cat-and-mouse game going on within these systems in that if you can understand the statistics of the noise that has been added to the data, in some cases it is possible to remove that noise with additional contextual information, be it data or expert opinion, so there are always ways to counter the inclusion of noise in data to de-anonymise. “Noise” is a statistical expression that I am using, and it is essentially about trying to modify the data such that it is away from its original—

The Chairman: It is not what the rest of us would call noise, basically.

Dr Mark Briers: My apologies for not being clearer.

Lord Swinfen: That is fair enough, because we are in a new realm and new words need to be invented, which you are doing at the moment, so congratulations.

Dr Mark Briers: It is a statistical word that has probably been around for many decades, but it is new in this context, certainly.

Lord Giddens: What are the implications of face recognition technology for data security? It seems to me that it cuts both ways in that it could improve security but it could also be misused very easily. It is already being misused, as I understand it. The consequences could be pretty serious, because picture identification is a key part of security with driving licences, passports and so much else.

Dr Mark Briers: My experience was in the defence industry prior to my current role. We focused on using multiple biometrics as opposed to a single biometric, such as face recognition, to improve the security of whatever system we were trying to secure. My personal view is that we should be fusing together multiple weak indicators of a person’s identity, be it fingerprints, facial recognition, gates, or whatever it may be, to maximise the probability that we correctly identify the individual we are trying to assess.

Lord Giddens: There is not much an ordinary punter can do against misuse, is there? For example, you can see a picture of President Obama saying things he never said, and all those kinds of manipulation are possible. It is not quite the same as facial recognition.

Dr Mark Briers: Indeed. If I am interpreting your question correctly and it is about the manipulation of the information space where people are suggesting that President Obama said something that he did not actually say, by manipulating the data in some respects, as I alluded to earlier, we need to counter that information space manipulation and present the facts in the way they should be presented. We need more research to ensure that those facts are presented in a way that the public can understand and can trust what is presented and how it is being presented to them.

Q151 **Baroness Rock:** We have talked quite a bit already about national resilience, and I want to focus a bit more on long-term resilience. You have mentioned investment for the long term and the importance of working across all sectors and not in silos. Could you take that a step further to how we can maintain the security of AI systems, particularly those of a safety-critical nature, for the long term? Who should be responsible for securing these systems, and how long should that responsibility be expected to last?

Dr Mark Briers: Professor Hankin, when he referred to the National Cyber Security Centre, hit the nail on the head in that it has to be a collaboration between academia, government and industry. That collaboration needs to endure. This is a problem that we will face at least for my lifetime, so we need to incentivise all three parties to work together to counter the threats and to mitigate the long-term problems that we currently face.

The Chairman: It is a step change in terms of risk, is it not? If it is bad now, with AI and autonomous systems, it will be even worse.

Professor Chris Hankin: I would echo to a certain extent what Dr Briers has just said: that the responsibility is across the whole supply chain from the vendors through to the users who are operating the systems. Some of the aspects of cyber hygiene, if I can call it that, which I mentioned earlier, would be the responsibility of the users to use the system in the way that it was intended, but the vendors, system integrators and so on also have a responsibility all the way along that supply chain to make sure that systems are as resilient as they can possibly be.

Viscount Ridley: Where would the malice come from in an AI virus? Human beings have four billion years of malice behind them to inject into the things they do, including computer viruses, and you can imagine setting up an AI bot to do nasty things, to continue doing them and to learn how to get better at doing them, but could it invent that? Could a benign AI go bad?

Dr Mark Briers: If one of the goals of artificial intelligence research is to mimic humans and human behaviour, inevitably we will encounter the

situation you have alluded to, which is that artificial intelligence could learn to be malicious by mimicking society.

Viscount Ridley: By copying us.

Dr Mark Briers: Essentially. You can kind of see it in the biases that are introduced in machine learning on a microscale, but if you extrapolate from that point, yes, I can believe that in the future at some point a generally artificially intelligent system, if that ever appears, could be malicious and naturally malicious as it has evolved artificially, if that is the correct phrase to use.

Viscount Ridley: So “Hal” and the “The Terminator” films are not all wrong?

Professor Chris Hankin: Time will tell.

The Chairman: You have mentioned the National Cyber Security Centre’s Cyber Essentials programme. Does it yet have any kind of AI-related component, or is that all in the distant future?

Professor Chris Hankin: Explicitly, no, it has no AI component, but it talks for example about having appropriate firewalls in your defence, and one could imagine the evolution of firewalls. In my answer to your first question, I said that one of the primary uses of machine learning these days is about intrusion detection and spotting anomalous network traffic, which is part of the role that a firewall is fulfilling, so one could imagine that implicitly already in Cyber Essentials, if the technology were readily available, the advice would be to use those kinds of defences rather than the firewall defences.

The Chairman: So cyber hygiene, for one, can be right. I am getting used to this idea of cyber hygiene and I like it.

Q152 **Lord Hollick:** What, in your view, is the single most important recommendation that we could make?

Professor Chris Hankin: I was delighted to read the AI review talking about the creation of capability at all levels of the educational pyramid, as it is presented in the report, and particularly that the industrial strategy has picked up on that and is funding 200 new PhD student places per year for the next few years and the creation of Master’s courses, which is very important. For the future, it is very important, going back to the issue of silos, that we recognise that cybersecurity is a priority within the artificial intelligence area and that a good number of those, or perhaps additional studentships at all levels, are funded to support this linkage between cybersecurity and AI.

Dr Mark Briers: My recommendation is very similar to Professor Hankin’s. I believe that the investment potential or the investments made by commercial organisations in artificial intelligence will increase exponentially over the next few years as they realise business value. I believe that there needs to be a proportionate increase in the resources available in cybersecurity-related research specifically to counter the AI threat. My recommendation would be on the increase in cybersecurity-related investment to counter those potential misuses and threats that we face.

Lord Hollick: So more trained people and more money.

Dr Mark Briers: Essentially.

The Chairman: It is very straightforward. That was very clear indeed. Thank you both very much for a very interesting session.

Dr David Barber, Digital Catapult and NVIDIA – Oral evidence (QQ 38–45)

Evidence Session No. 5

Heard in Public

Questions 38–45

Tuesday 24 October 2017

[Watch the meeting](#)

Members present: Lord Clement-Jones (The Chairman); Baroness Bakewell; Baroness Grender; Lord Hollick; Lord Holmes of Richmond; Lord Levene of Portsoken; The Lord Bishop of Oxford; Viscount Ridley; Baroness Rock; Lord St John of Bletso; Lord Swinfen.

Examination of witnesses

Dr Marko Balabanovic, Dr David Barber, Dr Timothy Lanfear.

Q38 The Chairman: Good afternoon and a very warm welcome to you. We have Dr Marko Balabanovic, Digital Catapult, Dr David Barber from the Alan Turing Institute and Dr Timothy Lanfear from NVIDIA. I have a little rubric that I need to read out on every occasion, which I am sure will bore everybody to death by the end of our evidence taking. The session is open to the public. A webcast of the session goes out live and is subsequently accessible via the parliamentary website. A verbatim transcript will be taken of your evidence. This will be put on the parliamentary website. A few days after this evidence session you will be sent a copy of the transcript to check for accuracy. We would be very grateful if you could advise us of any corrections as quickly as possible. If, after the session, you want to clarify or amplify any points made during your evidence, or have any additional points to make, you are welcome to submit supplementary written evidence to us. Would you like to introduce yourselves for the record and then we will begin with the questions?

Slightly unusually, we have Divisions taking place at just about the time that we are sitting, so if there is a Division in the Chamber while we are sitting, the Committee will adjourn as soon as the Division Bells have rung and will resume after 10 minutes. It will be pretty obvious when the bells go off; all hell breaks loose and we come back after 10 minutes and resume proceedings. I am sorry if that is going to cause a little bit of disruption but, hopefully, as little as possible.

I am going to open the questions. My first question is what, in your opinion, are the biggest opportunities and risks in the UK over the coming decade in relation to the development and use of AI? That is a fairly

general question. We would like to follow up on that in view of Dame Wendy Hall's review. Do you think the Government's review of AI will help with these? Does it go far enough? Perhaps we could start with you, Dr Balabanovic.

Dr Marko Balabanovic: In terms of the opportunities from our point of view—and I think this is very similar to what is mentioned in the review—there are going to be very few industries or walks of life that are not affected by artificial intelligence, so we would see the opportunities in the UK across every part of our society and every part of our industry. That will include the formation of entirely new markets. For example, the use of machine-learning techniques to improve speech recognition and speech synthesis is now leading to a new market for assistants in the home which operate with voice interactions, such as the Amazon Echo Alexa system. That was not there before and it is artificial intelligence that has made that possible. It is now in more than 11 per cent of US households that have broadband. It is a very rapid rise.

The Chairman: What percentage was that?

Dr Marko Balabanovic: It was 11 per cent but that was some time ago. I suspect it has gone up but I am not sure what it is now. They are not always keen to release up-to-date figures, as you might imagine. For the UK particularly, though, it is worth thinking about where we have differentiated advantages. There are a few industries, such as digital manufacturing, which have the higher-value and more datacentric side of manufacturing, and this is a strength that the UK could build on and use to compete globally and successfully. Artificial intelligence tools for the creative industries is a very interesting area that has not been picked up as much by some of other countries investing globally, and the UK has great potential because of its strength in the creative industries. An obvious one is the area of digital health, where there are great advantages to be had in improving health outcomes and producing a more efficient health system. Shall I go on to talk about the risks side?

The Chairman: Hold it there for the moment. Do you agree with that, Dr Barber?

Dr David Barber: The AI report is a great report. It goes a long way to address many of the key issues. It is perhaps a little light on some of the core research competencies and it could be strengthened in that sense, but I would fully agree with the issues that Marko has addressed in terms of opportunities. To reiterate, efficiencies in transport, personalised healthcare, drug discovery, freeing people from menial tasks such as being checkout assistants, et cetera, are obvious opportunities. Accelerated delivery of products, driverless cars and drones, et cetera, are all coming relatively soon, I would say. We can also see the potential for energy reduction by improving the efficiency of systems. If we think about plant monitoring and chemical processes, et cetera, efficiency could be dramatically improved by the use of AI. There are perhaps more contentious issues around things such as national security, where the ability to track and to monitor conversations automatically is rapidly increasing, and that is an opportunity and potentially also a risk.

The Chairman: An interesting aspect is the role that Dame Wendy sees

for the Turing. I am not sure whether you can answer the institutional issues about the future of the Turing, but I am assuming that is under consideration in terms of how the Turing is digesting Dame Wendy's report.

Dr David Barber: Yes, I believe so. I think the Turing would be delighted to play the role of the UK national artificial intelligence institute. In my personal view, it would require some significant restructuring of the current framework of the institute and, obviously, a huge increase in funding.

The Chairman: That is quite runic in a way. Is that because of the academic ownership of the institute?

Dr David Barber: Partly, yes. Also, the structure of the institute is largely based on the secondment model at the moment, so academics are seconded to go there, which is fine but not necessarily the best way to proceed if you want to make progress in focused areas of science, particularly in AI. In industry there tend to be quite large teams of 40 or 50 people working in very specialised areas and that is very difficult to achieve without having a large number of people who are full-time working in the institute.

The Chairman: You might want a model a bit more like the Crick where you have this match?

Dr David Barber: Yes, for example.

The Chairman: I am with you. Dr Lanfear, you are at the sharp end of investment in NVIDIA, so you really have to put your money where your mouth is, do you not?

Dr Timothy Lanfear: Absolutely. I think the company has definitely done that. A large percentage of NVIDIA's revenue is reinvested back into the company as R&D. I am not sure exactly what the R&D budget is, but the units are billions of dollars per year. All this goes back into creating the technology that is going to underlie this artificial intelligence revolution. All three of us are agreed that there is no doubt that artificial intelligence is going to have an enormous impact on the daily lives of every one of us. There are two themes that have already started to appear that are going to be critical for us every day. In the area of transport, self-driving cars is a clear direction of research and the other one is in healthcare. There are lots of things that could be done in healthcare. We see start-ups which are developing new concepts in healthcare. One of the companies that NVIDIA has worked with in the last year or so is using artificial intelligence to comb through all the medical literature to try to assemble pieces of disparate data that they can put together and that will perhaps give an indication of where we can look for better drugs in future. Another area, again staying in healthcare: I have read reports that there seems to be a shortfall in radiologists. Radiologists are being brought out of retirement to process the number of medical images that are being produced. If they can be assisted with artificial intelligence, perhaps that shortage can be addressed.

The Chairman: Those are your top-of-the-tree wins.

Dr Timothy Lanfear: The two areas where we are going to see the biggest societal impact are transportation and healthcare, in the short term.

The Chairman: In your view, therefore, is Dame Wendy's review helpful in the way you see things going forward?

Dr Timothy Lanfear: I must confess to have not studied the report in detail so I will not go out on a limb. What I am putting forward I would not say is the accepted wisdom but a commonly put-forward view from the leaders of the AI community that this is where things are going.

The Chairman: Were you consulted by Dame Wendy?

Dr Timothy Lanfear: No, I was not.

The Chairman: Not at the time?

Dr Timothy Lanfear: Not at the time, no.

The Chairman: That in itself might be somewhat telling. Thank you very much indeed. Now we go to Baroness Bakewell.

Q39 **Baroness Bakewell:** Gentlemen, it is quite clear that you are at the brainy end of this whole enterprise. You have invested a lot of training and knowledge in what you do and your careers, but where is the training of the future going to focus? Where are those skills going to come from? How far down the food chain are we all going to have to learn about AI? I would quote you Andrew Orłowski who was here the other day. He talked about his children going to an outstanding primary school in north London where they are taught algorithms every week but only taught history once or twice a term. He represents someone who knows the whole of your business, but he also is anxious about the emphasis shifting so heavily from the humanities. I am here to speak up for the humanities and would like to know how you think it balances out. Who would like to respond first? Dr Barber.

Dr David Barber: Gosh, humanities is an interesting one.

The Chairman: Do not believe that there is no right answer to this question.

Dr David Barber: I can address the first part of the question a little more comfortably. I would say that there is a requirement for a breadth of training across the whole spectrum. For example, if you are looking to plant flags in the IP landscape of the future in terms of AI, clearly you are going to need much more training at PhD level in this country. Currently, many of our competitors are investing very heavily in that area and we are way behind that. We are very far behind in terms of investment in training.

In terms of usage and application of AI, the master's level training is probably sufficient. We need to increase rather dramatically the number of MSc students who are graduating in this area in the UK. It would be appropriate for most undergraduates to have some understanding of what AI is, what the applications are, and what the risks and potential benefits are as well. Whether that goes right down to the primary level, I

am not quite sure—or how far it is pushing out other areas such as humanities.

Baroness Bakewell: Presumably, of course, it will affect the lives of the people who are tiny at the moment even more than it will impact on the lives of people who are halfway into the old-fashioned balance of things. There must be a need. Can I ask you, Dr Lanfear, what you think about the youngsters learning this stuff?

Dr Timothy Lanfear: First, I think they should learn from the beginning. The sooner you can start teaching people these things, the better. I have a lot of sympathy for your position on wanting to defend the humanities. Personally, I am a technologist during the daytime but in the evening I am a musician. That is an important part of my life and I see that humanities are being pushed out of our educational process. They should have a much stronger place than they do today.

Baroness Bakewell: Should the humanities be safeguarded in some quite direct way?

Dr Timothy Lanfear: That is probably a little outside the topic we are talking about today.

Baroness Bakewell: I always try to look beyond.

The Chairman: We are used to this, do not worry.

Dr Timothy Lanfear: To come back to the more AI-focused part of your question, education has to be a key part of it. It is going to affect the lives of everyone, so starting in schools is going to be critical. When we come to tertiary education, it should not be in a special department of its own. It is going to affect the lives of all professionals, so all professionals are going to need to know something about this. Education will be a third theme that we are going to bring up repeatedly here today.

Baroness Bakewell: Do we have the teachers to do it?

Dr Marko Balabanovic: Do we have the teachers to teach artificial intelligence? I am not sure, but can I come back to you on the humanities first? It is a mistake to view artificial intelligence as a purely technical endeavour. Anybody who studies it, as I did in a computer science department at PhD level, will find that it crosses over with philosophy, linguistics, psychology and design. It is a very multidisciplinary effort. As soon as you are applying artificial intelligence in the real world you are immediately having to deal with issues of design, entrepreneurship, ethics, the philosophy of how we want these kinds of systems to work and the sociology of them. I would say there is an inbuilt degree of humanities in any case. As to whether the teaching needs to start at a very young age, I am not qualified to say, other than if you look at all the current superstar academics in artificial intelligence, they probably did not start it at primary school, so that is at least one proof point that perhaps you can get by without having to start at age five.

Lord St John of Bletso: It is well known that companies in the AI space struggle to find qualified staff and are having to retrain staff. What scope is there for more collaboration between the academic institutes and the private sector to shape modules?

Dr Marko Balabanovic: It happens already. I agree the best way to think about it is as a track that starts from a young age with increasing study of STEM subjects and going all the way through to PhD, and it is a very long-term investment. The workforce is a much more interesting place to look at in the short term. Online training is now quite a large field, and, effectively, most of the large companies in that space were started by people who originally came from an AI background. Perhaps that is just a coincidence but those were the popular courses at the beginning. Most of those courses are co-created and co-sponsored by industry and academics. They are often free or very cheap. It is a somewhat disruptive model compared to in-person training and going to university, but if anybody wanted to learn about these subjects now, there is a wealth of material available. It is a very interesting time for that.

Q40 **Viscount Ridley:** Can I bring you on to the question of ownership: the ownership of data and of the outcomes of algorithms? You will know that DeepMind's deal with the Royal Free NHS Trust led to a lot of criticism and Sir John Bell in his recent review of the UK life sciences industry said, "What you don't want is somebody rocking up and using NHS data as a learning set for the generation of algorithms and then moving the algorithm to San Francisco and selling it so all the profits come back to another jurisdiction". When UK taxpayers contribute to training and research in AI within the UK, how can we ensure that they see a meaningful return on this investment, and that UK state investments are not simply subsidising AI enterprises in other countries?

Dr David Barber: That is a complicated question. I believe there is a concept called the "golden share" which enables companies to give money back to the contributors of the data. I believe the Department for Transport is heavily involved with that. That is one potential mechanism to see a return on the investment. There are sometimes difficult relationships between companies and, for example, universities where the companies are looking for privileged relationships with the universities and it is difficult to protect some of the intellectual property. There are difficult questions around how to retain the independence of the universities in the light of companies wishing to have privileged relationships with them.

Viscount Ridley: If we are too squeamish about this and we protect the rights of individuals to be rewarded too fervently, the industry will not come here and do anything.

Dr David Barber: Absolutely.

The Chairman: You mentioned the difficulties of protecting intellectual property in datasets. The trouble is that it is often not entirely clear that there is intellectual property in that sense. There may be property, but it is quite tricky to protect that if somebody wishes just to take it abroad. Would the Alan Turing Institute be a useful repository for a data trust or something of that sort?

Dr David Barber: In a sense, yes, I think we could be, but we are also talking on an international scale. For example, if you think about healthcare, the paucity of datasets in radiology is very difficult to

overcome at a national level. We are talking about international collaborations for these datasets. There are very difficult questions around regulation. For example, even if an individual gives consent now for their data to be used, they could revoke that consent later on, and it is a question of how you have systems which can manage those kinds of datasets. It is a very complicated question which I believe people are fervently researching right now, but we are just at the beginning of this discussion

The Chairman: There is a lot from the Royal Society and others who are trying to address this data issue. Because of the importance of data—good data and datasets—that AI is applied to, is that not one of the issues? Has Digital Catapult addressed this issue at all?

Dr Marko Balabanovic: We have helped with the Royal Society and Royal Academy data governance report. In fact, I was there this morning and had a session on that. They have a recommendation for a new what they call stewardship body to take this forward in this country and really look at this landscape that has evolved. It is much more complex than the piece that the Information Commissioner's Office might handle or the piece that the National Data Guardian might handle for health. They have set out a really strong direction that we would support. It will take some time to figure out structurally where that should go.

The other thing to bear in mind is that the market dynamics in the world of AI mean that for the organisations that have the large customer bases and a large amount of data already, there is a network effect that leads them to get bigger and bigger in that sense. It is really important to look at addressing the market barriers for UK growth as opposed to the incumbent what are called data monopolies or AI giants. These are six or seven of the biggest companies in the world—you mentioned DeepMind—which have been so successful they will have customers in the billions and therefore they will have data about people on that sort of scale, and they will attract skilled and talented people—and all this builds upon itself. The risk of things going abroad can be countered by helping UK companies grow bigger and not having them be acquired when they are very young by some of these giant companies.

The Chairman: That is exactly what Viscount Ridley and I were trying to stimulate you to express.

Viscount Ridley: Can I play devil's advocate for a second? The data from a scan of my pancreas, say, is of awfully little use to me because I cannot interpret it or do anything with it, but if it goes off to California and someone turns it into a tiny part of an enormous project that ends up in a cure for pancreatic cancer, which I then benefit from in 10 years' time, what is the problem? Why should I be fussy about something that is not very valuable to me at this point? It is only valuable if it is processed and turned into something else.

Dr Marko Balabanovic: The problem would only come if the cost to the UK health system to access the system that you have described is very high and if it was built using taxpayer-funded data from the UK.

Viscount Ridley: Is it me as a taxpayer rather than me a pancreas owner that you are really interested in, or that I am really interested in?

That is my question.

Dr Marko Balabanovic: You are right. The model where you have full control of your own data would be a very powerful way for things to move forward. It does not tend to be how things are today where, obviously, the scan would be held inside some kind of hospital trust, but, according to the way these regulations are going, you would have every right to access that data, hold it yourself and give it to a company in California if you so wished, and there would be nothing to stop you doing that.

The Chairman: I am going to bring in Lord Swinfen and Baroness Grender, but first I am going to ask Dr Lanfear if he has something to add to what has been said.

Dr Timothy Lanfear: I do not have a great deal to add other than to reinforce the fact that data is really the crown jewels in the artificial intelligence world. NVIDIA is a large enough company to have resources to be able to generate its own data. Self-driving cars is one of our key topics. We have a programme of gathering data that is going to help create the algorithms that will shape the future of self-driving cars.

Lord Swinfen: Is UK national law, and indeed international law, up to date?

Dr David Barber: Could you be slightly more specific; what is your concern?

Lord Swinfen: On the ownership and use of the data. This is a developing industry and law must come into it at some stage. My feeling is that it is way behind at the moment.

The Chairman: If that is too tough, I think we will park that question and interrogate the lawyers on a future occasion—or you could come back to us.

Dr Timothy Lanfear: We are technologists on this panel.

Dr Marko Balabanovic: I would think the problem is not so much whether the law is up to date but whether we know how to interpret the law that is coming. The new Data Protection Act has various requirements—for instance about being able to explain the outcome of decisions made by algorithms, such as AI algorithms. I do not think anybody has a really good understanding of how that is going to work in practice. I do not know if there is a problem with setting down what the law is or whether it is a problem of understanding how to interpret it.

Baroness Grender: Dr Lanfear, you mentioned that the crown jewels is the data. We all own the National Health Service and it is taxpayer funded. Is there a danger that we are somehow selling short or not realising this data as a rich asset?

Dr Timothy Lanfear: First, there is no doubt that it is extremely valuable. If you look at the behaviour of these large companies that we mentioned—the Facebooks, Amazons and Googles of this world—they are very open with their algorithms and very closed with their data. They always expose their algorithms because they would like other scientists to work on those algorithms and improve them. Google can reap the results

of that by applying those algorithms to the data which it keeps entirely secret. There is no doubt that the data is the key point here. The second part was about whether we are selling ourselves short in this country.

Baroness Greender: The NHS is almost unique in that respect. It is owned by us as a nation. It is deep and rich in data. That may be hidden in a filing cabinet somewhere but, by the same token, it is there.

Dr Timothy Lanfear: It is invaluable. The country should find a way to unlock this resource. It is such a valuable resource and if you can stimulate the start-up community in this country and it can make some use of it, some great benefit for the country may come from that.

Baroness Greender: I will address the question to other members of the panel: is there a danger that this is an asset that we are almost letting go on the cheap?

Dr David Barber: Potentially. Clearly it is a great asset, but if we do not use it in some way that is potentially a huge lost opportunity. One of the really tricky aspects is how to use the regulations as they currently are to gain access to that data so we can gainfully use it in some way. If it is so heavily regulated that we are not able to use the data in any usable sense that would also be a great shame. The tricky issue is how to gain access to the data such that we can make use of it in some way. It is a tremendous asset and, absolutely, we should not be selling ourselves short in letting anybody get access to it.

Lord Hollick: Do the academic sector and the academic institutions have the skills to fight their corner in this area?

Dr David Barber: We have the skills; we do not have the resources.

Lord Hollick: So you have the commercial skills to make sure that you get fair value your side of the table.

Dr David Barber: Do you mean in terms of data access and negotiation?

Lord Hollick: Yes, correct. You are negotiating against some fairly successful global companies.

Dr David Barber: I am not sure about that. That may be something that we need help with. The Turing Institute, for example, could potentially play that role for us.

The Chairman: Which brings us neatly on to Lord St John.

Q41 **Lord St John of Bletso:** Should state and privately funded research institutions be encouraged to work more closely together—and, if so, how?

Dr Marko Balabanovic: One thing I would imagine would be more valuable than that is to create the whole spectrum including privately funded research and state-funded research, and the research and innovation that happen in small companies and start-ups and the larger companies, which often need to be the adopters of this technology, do not necessarily have the mechanisms or skills to do that. Collaboration across that whole spectrum and a collective voice in this country would

give us a great strength. Research often needs a problem to solve and that might come from the industry side.

Lord St John of Bletso: I notice, Dr Lanfear, that NVIDIA runs a series of programmes to provide assistance and support to academia around the world. Perhaps you could elaborate on this.

Dr Timothy Lanfear: Yes. We have identified a number of top artificial intelligence research institutes around the world, Oxford University being one of them, and we have a programme in which we make donations of equipment to them. We have people whose job it is to build a tight relationship with those organisations and ensure that they get what they need. It is a two-way street, so we also learn a lot from them. We bring AI experts to our headquarters once a year. We learn from them what they want to do and what sort of technology NVIDIA could develop to help them do that process. It is a two-way interchange between us and the AI academic sector.

Lord St John of Bletso: I was reading in the robotic space that there is a company called Robiquity which has been training students and trying to prepare them for the practical aspects of robotics as well.

Dr Timothy Lanfear: At Imperial College there is now an MSc course on robotic assisted medicine.

The Chairman: The Turing has quite a lot of relationships with business now, does it not?

Dr David Barber: It does, yes. Addressing the universities point, obviously universities want to have relationships with companies, and that is all well and good, but care has to be taken with some of these relationships. Some of the tech giants want very privileged access to certain universities in this country, particularly those at the leading edge of AI. There are some careful questions that need to be asked about those kinds of relationships. Of course they are good for the university, they are probably good for the individual academics, and they are almost certainly good for the companies, but is it good for UK plc? That is an interesting and important question to ask.

The Chairman: That is very interesting. How do you patrol the border in the Digital Catapult?

Dr Marko Balabanovic: It is not our role exactly to patrol it, but we are measured according to whether we can see UK growth and UK productivity, and we would seek to do that in collaboration. If there can be structures where some of the larger technology companies are able to apply their investment into this country, we would see that as a positive thing. If they are employing people in this country, we would see that as a positive thing. If there are opportunities for UK academics to take their ideas out of the universities and into industry, we would definitely be helping that along the road and working with partners such as those around the table to help make that happen. On the whole there is a strong, positive story there.

The Chairman: These relationships are very fluid between the different institutions, are they not, private and public and academic and so on? That is the impression I get in AI in particular.

Dr Marko Balabanovic: At the moment, the interesting thing in AI is that in the main AI conferences the majority of the research publications come from corporate affiliations rather than academic affiliations. It is very much a model where it is sometimes hard to unpick exactly who is in a university and who is in a company, and people have dual appointments or part-time consulting, so I would say that it is very fluid.

Q42 **Lord Hollick:** The UK has a rather long record of coming up with some brilliant inventions but failing to exploit them and translate them into developed products. Graphene and the fact that most of the patents are now owned offshore is a good recent example. What steps do we need to take to ensure that we benefit from the excellent and leading research that is being done here in the UK?

Dr David Barber: I think we are a little risk averse in terms of investing in the UK, particularly in this AI space. There are a lot of successful tech start-ups which came out of the UK which in the end became successful only when they went to Silicon Valley for investment. It is unfortunate that has needed to happen over the last few years. It would be great if there was a stronger emphasis here on nurturing early start-ups particularly coming out of universities. Some form of national incubator or accelerator programme might be very useful within the UK.

Dr Timothy Lanfear: First, I would like to say that the UK is doing very well in this area. A McKinsey Global Institute study has shown that, while the number of papers coming out of the UK may not match a very large country such as China, the influence of those papers is disproportionate to the quantity. We have a programme of good innovation in this country. How can we keep it here? I think my colleagues here have asked for investment, and that seems to be the right thing to do. There is a burgeoning start-up industry in Europe in general. I know NVIDIA is involved with something like 1,000 or more start-ups around Europe. I do not have offhand the exact number for the UK, but start-ups are generating the ideas here and bringing new ideas that can be brought to fruition and real products can be made out of them.

The Chairman: You are able to give advice to spin-outs from academic institutions and help them commercialise?

Dr Timothy Lanfear: Indeed, yes. We have a programme to do that. We have these large events called the GPU Technology Conference. One of the features of those GPU Technology Conferences is a start-up track, and within that we will select perhaps a handful of very promising start-ups and give them financial awards so that they can continue their work.

Dr Marko Balabanovic: At Digital Catapult we would look at whether there are specific barriers to the growth of UK companies that we can address that would help prevent the sort of situation you are talking about. Late-stage AI funding has been mentioned; skills have been mentioned in terms of access to enough smart people; access to data has already been mentioned. Recently we found, somewhat surprisingly, that there are around 600 AI start-ups in this country out of about 1,200 in Europe, so we are very well positioned. However, more than half of those are constrained at the moment by their access to computational power. That is where we are with artificial intelligence at the moment. It is an

expensive endeavour. The kinds of techniques people are using these days with vast amounts of data, which need vast amounts of computing, are very expensive. That does not matter if you are a Google or an Amazon because those companies run their own huge data centres, so it is not a problem for them, but it is a problem for smaller companies in the UK, and it can be a problem for academics as well. We are working quite hard to see if we can put initiatives together to help address that barrier. That is an example of the kind of thing we would do, and there are others like that.

Lord Hollick: Dr Barber, you mentioned the Government possibly putting more funding to work in this area. That is another long-term record where the Government are not terribly good at doing that. Can you point to any Government who are currently focusing on artificial intelligence and building artificial intelligence competence in their country? What methods have they used? I am thinking in particular that the US and Israel both use the defence budget very heavily to support innovation and then make it available freely to the private market. What sorts of examples are there abroad that have worked?

Dr David Barber: Interesting examples have recently come out of a couple of countries. Canada, for example, is investing very heavily in this; both local and national governments. Last year, the Canadian Government announced a 230 million Canadian dollar investment in AI in universities in Montreal alone; just one city. Similarly, Toronto is doing extremely well. The Vector Institute was launched recently. Canada has a strategy to make the next tech AI giant Canadian, and they will probably succeed, and it will stay in Canada.

Similarly, last year in Germany there was the announcement of Cyber Valley; a collection of universities and research institutes close to Tübingen in Stuttgart, where, again, the whole spectrum from academic to incubators has been addressed, and there will be specific programmes there. They are looking to hire 100 PhD students and create nine new research groups and 10 professorships. The scale is huge, but that is just in Baden-Württemberg alone: one particular region of Germany. Each of those investments is huge. They dwarf the scale of what is going on in the UK. To give you some numbers, 1.48 per cent of the EPSRC's budget—the main funding council for AI in the UK—is allocated to AI, so you are talking roughly £13 million per year. This is a drop in the ocean compared to most of our competitors in this space.

Lord Hollick: Does the Alan Turing Institute see as part of its role to become an advocate for this?

Dr David Barber: Absolutely.

Lord Hollick: Have you published anything along these lines?

Dr David Barber: In terms of advocating this to broaden the space, not that I am aware of, but the remit for the Turing Institute has been more about the data science space. There is absolutely a hope and expectation to move more into the AI space, but that level of commitment would require financial support and it cannot easily be done with the current resources at the institute.

The Chairman: Can I bring in Baroness Bakewell?

Baroness Bakewell: We have spoken of these amazing PhDs and the extremely large number of start-ups and the expansion of this world. We know that it has always been difficult to get young women to study engineering. Is there a similar problem in the AI world, that it is biased against the recruitment of women?

Dr David Barber: I am not sure about the bias against the recruitment of women. Certainly, it is difficult to hire women in this space—this is my understanding—simply because there are not enough women graduates in this area. I do not know if there is a particular bias against them per se.

Baroness Bakewell: I use the word because it crops up in the documents and reports that are given. We want, do we not, women to be a part of this burgeoning world; what steps can be taken?

Dr Timothy Lanfear: I do not have an easy answer to this. If there was an easy answer, I am sure we would have done it already. Personally, I have hired two women in my team and I am hoping to hire a third one in the near future—so I do my bit where I can.

Baroness Bakewell: Out of how many?

Dr Timothy Lanfear: I think I have about 10 people reporting to me at the moment.

Lord Levene of Portsoken: How do you rate China in this race, or climb, if you like?

Dr David Barber: They are doing very well. They seem to have a very specific strategy in AI. They are hoping to invest very heavily in that. They really know what they are doing. They have excellent researchers and very good research institutes.

Lord Levene of Portsoken: Do you think they are going to overtake us or have they already done so?

Dr David Barber: I do not think they have done so yet. If you look at many of the leading conferences and journals in this space, certainly they are very competitive with us, but we are still doing very well.

Dr Marko Balabanovic: If you look at the overall size of the ecosystem of companies, most rankings would have the US and China some way ahead of everyone else. Obviously, there is an enormous advantage in terms of data, as we talked about before. There is a large population and the data protection requirements are different in China, so there are certain attributes of that system that mean it can grow, but the government investment is orders of magnitude higher than elsewhere.

Lord St John of Bletso: They have been making a few acquisitions. Did the Chinese buy ARM and Skyscanner quite recently?

Dr Marko Balabanovic: I did not realise they had bought Skyscanner.

Q43 **The Lord Bishop of Oxford:** You will not be surprised to know that ethics runs as a thread through all our questions. How do you ensure that research and development in AI is conducted in an open and ethical

manner? What would you commend to the group and the wider industry so that the benefits of the research are shared as widely as possible? Could we start with NVIDIA? A number of us are curious as to why the section on ethical implications in NVIDIA's written submission was not answered for reasons of space. It would be very good to hear your views.

The Chairman: You could expand that to include ethics in the translation of research into practical products. We are interested in the gamut of ethical behaviour.

Dr Timothy Lanfear: It is not an easy topic to address. NVIDIA is a technology company. Our goal is to provide the technology that underpins what the AI community will build on. We have no restrictions on who can use it so we will freely offer—I say offer but in fact sell of course—our technology to whoever would like to use it. We are an international company so, although it is headquartered in the United States and has a focus in the United States, we operate globally, and we have no restrictions about where our technology can be used. NVIDIA, by its principles, is an ethical company. We have a code of ethics. It is underlined by the chairman of the company. We have ethical principles. I confess I am struggling a little to address this question because as a technologist it is not my core thinking.

The Lord Bishop of Oxford: It is interesting and reassuring to know that a code is there. Is it, for example, publically available? How do you take steps to educate your own workforce in its values and ensure compliance?

Dr Timothy Lanfear: It is part of the induction programme. When you join the company there are a number of courses you have to do and the code of conduct is one of the mandatory courses that all our employees have to take when they join the company. As a manager I get reminders if they have not done that task yet.

The Lord Bishop of Oxford: Are they publicly available?

Dr Timothy Lanfear: I honestly do not know the answer to that question.

Dr David Barber: There is a set of ethical principles that some companies—Google, Facebook, et cetera—are signed up to. I forget the name of it now but there is a voluntary code they are signed up to. There is a role for the Turing to play here in helping educate people and holding discussions and debates about what this means. It is a very difficult question and not unique to technology. This is about broader society and bringing in people who have philosophical ideas on the impacts on society, et cetera. It is the beginning of this debate. There are very few clear answers to that. There are also interesting initiatives. For example, OpenAI in the States was founded on the idea that AI should be open and accessible to all, somewhat in contrast to some of the tech giants. It may be that a similar kind of institute might form around the Turing.

Dr Marko Balabanovic: We have seen that the AI review has recommended an AI council, which we think is a really good step, although it would need potentially some further structure around it so that it could really make a difference. There is the work from the Royal

Society and the Royal Academy recommending the stewardship body. Between these different suggestions there is clearly a need for a further body that goes beyond the ones that are there today, because the ethical issues are new and many of them need thinking through and research to figure out what it is we want to do. It is not a matter of applying a set of known principles at this point. There is a lot to do. In our own work we are currently setting up an ethics committee for the projects we do.

The Lord Bishop of Oxford: That is good to know.

The Chairman: Is it your feeling that government needs to take the initiative on this or could the collective organisations involved in AI, both private and public, take the initiative on this?

Dr Marko Balabanovic: Part of it will be a regulatory aspect and how that is handled is part of the discussion that needs to happen. I do not think that it can be an entirely voluntary process.

The Chairman: You perceive some form of regulation?

Dr Marko Balabanovic: It may vary by sector and in different industries. We already have lots of structures. In healthcare, for instance, there are various structures already that might need to adapt.

Q44 **Lord Swinfen:** To what extent do you expect established techniques, such as deep learning, to continue delivering meaningful progress in AI?

Dr Timothy Lanfear: We are only at the beginning. The first real appearance of deep learning in the field was around 2012, five years ago. It is a very young topic. Deep learning may indeed be just the beginning. There are lots of new ideas being put forward now. It is a very active field of research. For example, you may have seen the recent announcement from Google. One or two years ago they had a more traditional deep-learning approach to playing the game of Go. The new generation is called AlphaGo Zero, which uses a new methodology of artificial intelligence to generate a skilful Go player with no information whatever. I believe there is plenty more to do in this field. The amount of innovation is endless.

Dr David Barber: It is also worth bearing in mind that, essentially, where we are right now is what we call perceptual AI. I will give you an example. If somebody speaks, the machine can transcribe into words what you are saying but the machine does not understand what you are saying. It does not understand who you are or the relationship of objects in this environment. We have done very well in the past few years. Deep learning, for example, has made huge progress in perceptual AI, but the bigger fruit out there is the reasoning AI; really understanding what these objects are, being able to query this machine and get sensible answers back. That is the biggest and most exciting challenge that all the tech giants are currently desperately seeking to solve. For whoever solves that, the world is their oyster. That is why we are seeing what we are right now. They are parking the bus, if you like, on the machine learners here in the UK. The UK is a rich place for people to do this kind of research and they would love to get their hands on that. I think it is extremely exciting.

Lord Swinfen: Do you see any limitations to this approach, and, if so, how are you going to overcome them?

Dr David Barber: It is always tricky to predict those kinds of things. Many years—50 years—ago people were forecasting this kind of state of the art. I am not really sure what the limitations are. I think there is going to be rapid progress in the next five to 10 years. We will certainly see the evolution of our daily tools and we will get much more sophisticated personal assistants. We will see the inclusion of these perceptual AI systems much more pervasively in our daily lives. We are quite a way away from that wonderful machine which you can just speak to and it can answer any question that you have, or some robot which can assist you in your daily life or you can interact with in an empathetic way. That is a very long way away. These are huge research challenges which people are very excited to address right now.

Viscount Ridley: I want to take off from Dr Barber's point about the difference between perceptual and reasoning AI and use it to explore the issues of transparency and interpretability. We have heard lots of different views on this. We have heard people saying it is absolutely vital that all AIs must be interpretable and other people saying it is an impossible and unrealistic way to go. In NVIDIA's written evidence, Dr Lanfear, you made the remarkable statement that the "assertion that neural networks lack transparency is false", and the fact that the code is written in a programming language, you go on to say about ordinary programs, "just provides an illusion of transparency. It is not possible for humans to read and the only validation can be done through tests. From this perspective, neural networks offer a significant advantage over the handwritten codes". In other words, things are getting more transparent as we move to AI; is that what you are saying?

Dr Timothy Lanfear: Maybe not more transparent. What I want to say is that things have not been transparent in the past. Many aspects of our lives are predicated on things that have been computed; the design of bridges or aircraft has been done by complex computer programs.

Viscount Ridley: Which are black boxes.

Dr Timothy Lanfear: Which are, in some sense, a black box, but that is overstating the case. The people who write these pieces of software take great care in designing them so that they can split them down into smaller pieces that can be tested to be sure they are doing the correct function. They can then assemble those pieces step by step to make more and more complex things. We are using systems every day that are of a level of complexity that we cannot absorb. Artificial intelligence is no different from that. It is also at a level of complexity that cannot be grasped as a whole. Nevertheless, what you can do is to break this down into pieces, find ways of testing it to check that it is doing the things you expect it to do and, if it is not, take some action.

Viscount Ridley: There is a famous essay entitled "I, Pencil" in which somebody points out that nobody on the planet knows how to make a pencil. The information is not held inside a single head. You have to know how to cut down a tree and mine graphite and all these different things. Is that the same point really? We might not understand something

outside of the computing realm, so why should we get hung up on transparency?

Dr Timothy Lanfear: I was not aware of that analogy about the pencil, but it is a very interesting thought. The role of the technologist is to control what we are building; to control it, to test it, to be sure it really does do what it is designed to do and not something different.

The Chairman: I wanted to follow up with you, which I should have done earlier, on the regulation aspect, which is slightly akin to what we have just been talking about. Do you see a role for regulation as a private sector company or do you think there would be a danger of stifling innovation, whether in the context of enforcing transparency or ethical behaviour, or whatever we might think appropriate?

Dr Timothy Lanfear: NVIDIA wants to operate within the laws that are around it and we want to behave in an ethical way. It is part of the core values of the company to act ethically. If there are legal constraints, we will certainly follow them. I do not see any reason why we should try to work around them or avoid them. Those constraints are there for a reason and we absolutely respect the reason that that regulation is in place.

The Chairman: I am going to bring in Lord Hollick. Dr Balabanovic, you have not said anything on this particular question, so if you want to address whatever Lord Hollick has to say, by all means do so.

Lord Hollick: I wanted to follow up on regulation. Can you envisage a situation where the Government would have a keen interest in who bought a leading AI company? DeepMind was bought by Google, and it caused a bit of a stir, but let us say a Russian company had bought it. Russia does not have the best reputation in terms of cyberspace, so should the Government play a role in intervening on ethical and security grounds?

The Chairman: Dr Balabanovic, you are in the hot seat since you have not said anything on this question so far.

Dr Marko Balabanovic: From the economic side, part of the equation when we look at these things is to do with what happens next. Are they maintaining a presence in this country? Are they building up their staff in this country? Are they using this as a place to develop products and so on or is it a small outpost for sales and so on? That is a really important consideration. If there is a question about national security, that falls into an entirely different category of thinking through. I am not sure that is specific to AI skills. I would refer back to say that we cannot pretend the world is not the way it is. Currently, the big acquirers of companies tend to be these big US and Chinese companies and, for us to combat that, we need to build up the industry in this country so that people do not need to be acquired by a big company for their investors to see their returns or for them to grow to the next stage. They need to have enough capital and wherewithal and fewer market barriers to grow in this country.

Q45 **Baroness Greender:** Could we have one recommendation from each of you that you think this Committee should make at the end of our deliberations?

Dr Timothy Lanfear: I was going to follow up, but did not have a chance, on one of the comments that was made about the amount of computing resource that is being invested in this country. First, there is some good news that Oxford University has recently procured a system for artificial intelligence, but let us compare that with what some other large companies are doing. The system in Oxford is one tenth of the size of the system that a single department within Facebook has for doing artificial intelligence research. The whole country's investment is one tenth of what a department in one company is doing. The technology community here is being starved of resource and could do a lot more with good computing investment.

The Chairman: Thank you. Dr Barber.

Dr David Barber: It is hard to make a single recommendation; I have so many things to say. It is vital that somebody in government really understands what AI is about and what machine learning is about. I feel that there has been very little deep understanding of what this topic is about. The community has been poorly served, particularly in terms of research funding and government support generally, over the last 10 years. In some senses, it is despite the lack of support that we have been incredibly successful, and that is testament to the talent in this country. Therefore, in some senses, it is extremely important that somebody starts fighting the AI corner and really has a vision of what we can do in this country. There is too much defeatism and fear around the idea of AI. We need somebody to champion this idea and speak up on the positives. It is extremely important and is going to be transformational for us. Yes, there are going to be many ethical issues that we are going to have to address, but we can do that. It is a positive thing and I would like to see the UK remain a world leader in this space, but we need somebody to start fighting, get the funding and start matching what other countries are doing, otherwise it is not going to happen.

The Chairman: Are you writing the Turing's job description, I ask myself?

Dr David Barber: You could say that.

Dr Marko Balabanovic: I would echo a lot of that. We have talked before about the levels of investment even from countries such as Canada. We are in a really strong position and we may lose that soon if we do not invest appropriately. Having a collective voice in this country for AI research and industry together, and having a much larger level of investment but having it delivered by some sort of body that is accountable to deliver these kinds of results, would be the recommendation—not just an industry council but something that can do the work and make sure that the results happen.

The Chairman: Thank you very much. You have provoked our thoughts and questions. I hope we have done likewise in a number of respects. If you have additional things you want to write to us about, by all means do. In the meantime, thank you very much; we really appreciate your presence here today.

Jeremy Barnett, Professor Chris Reed and Professor Karen Yeung – Oral evidence (QQ 29–37)

Jeremy Barnett, Professor Chris Reed and Professor Karen Yeung – Oral evidence (QQ 29–37)

Evidence Session No. 4

Heard in Public

Questions 29–37

Tuesday 17 October 2017

[Watch the meeting](#)

Members present: Lord Clement-Jones (The Chairman); Baroness Bakewell; Lord Giddens; Baroness Greender; Lord Hollick; Lord Levene of Portsoken; Viscount Ridley; Baroness Rock; Lord St John of Bletso; Lord Swinfen.

Examination of witnesses

Professor Chris Reed, Jeremy Barnett, Professor Karen Yeung.

Q29 **The Chairman:** A very warm welcome to Professor Chris Reed, who is professor of electronic commerce law at Queen Mary University of London—in this context, I need to declare an interest as chair of council at Queen Mary University of London, but it is not a fix, I assure you—Jeremy Barnett, a barrister from St Pauls Chambers and Leeds & Gough Square Chambers, and Professor Karen Yeung, professor of law and director of the Centre for Technology, Ethics, Law and Society at the Dickson Poon School of Law, King’s College, London. It is very nice to see you all, and thank you for coming to give oral evidence.

I have a little rubric that I need to read out which my colleagues will be infinitely bored of by the time we get to the end of this inquiry. The session is open to the public. A webcast of the session goes out live and is subsequently accessible via the parliamentary website. A verbatim transcript will be taken of your evidence and will be put on the parliamentary website. A few days after this evidence session you will be sent a copy of the transcript to check for accuracy, and we would be grateful if you could advise us of any corrections as quickly as possible. If, after this session, you wish to clarify or amplify any points made during your evidence or have any additional points to make, you are very welcome to submit supplementary written evidence to us. Perhaps you could each start off by introducing yourselves, for the record, and then we will begin with the questions. Professor Reed.

Professor Chris Reed: Do you want the two-minute biography or just the name?

The Chairman: Perhaps a compromise between the two.

Professor Chris Reed: I am professor of electronic commerce law. I have been working in this field for 30-plus years, perhaps 40—depending

on how you count it—and artificial intelligence has been an interest of mine since the early 1980s.

Professor Karen Yeung: I am a professor of law. There are two core strands of my research. The first is regulatory governance. The second is the regulation of new technologies and the regulation through new technologies, and in the last few years I have had a particular interest in computational technologies that refer to big data, predictive analytics, blockchain and robotics.

Jeremy Barnett: I am a practising regulatory barrister. I also sit as a part-time member of the judiciary, but I do not speak as a member of the judiciary today. I also have an academic interest in law and emerging technologies. For four years I led a research group at Leeds University in virtual courts. That led me into online dispute resolution. I am now working on algorithmic dispute resolution at UCL with Professor Treleaven. He has got me interested in the law and governance of algorithms, which is why I am here. At UCL we have the Construction Blockchain Consortium, of which I am one of the directors. A lot of the issues you have already discussed today we are starting to deal with in practice.

Q30 **The Chairman:** Thank you very much indeed. You are all supremely qualified to answer our questions today. First, let me ask a very broad general question. In your opinions, what are the biggest opportunities and risks for the law in the UK over the coming decade in relation to the development and use of artificial intelligence?

Professor Chris Reed: I can give a very brief answer. Essentially, there are no opportunities for the law, because the law has to follow the developments of society rather than try to lead it. Every time I have seen the law try to lead society in a particular direction it tends to have failed.

On whether there are opportunities, the answer, essentially, is no; it is a job where we have to catch up with the way society and technology are developing and try to help society regulate in the best way possible.

On the risks, the biggest risk is the risk of rushing into regulation too early. We have a long, long history of prospective regulation for technology. I can think offhand of only one example that really worked. All the rest were failures. We cannot reliably predict the future.

Professor Karen Yeung: I will start in responding to this question by thinking about what I understand the biggest general opportunities for AI to be, and from my understanding of that to identify what the legal risks and opportunities might be. For this purpose, I am going to leave aside robotic applications, because I think robots generate particular questions of their own. I am focusing only on AI-driven decision-making systems that do not have a physical dimension in time and space. It seems to me that the greatest opportunities for society lie in the capacity to subject these large volumes of digital data that we now generate to machine-learning algorithms for the purposes essentially of process optimisation. That could be business process, governmental process, research processes—doing amazing things for astronomy, medical research and so on and so forth. This is because the decision-making systems that use

these machine-learning algorithms allow for more efficient decision-making that is data-driven. It is a kind of hyped-up evidence-based decision-making system that is much more responsive, that can generate decisions in real time so we can reduce the massive time delay that is often associated with evidence-based decision-making, and, crucially, that has the capacity for prediction.

This is where so much of the excitement about AI comes from. It is through analysing historic data patterns in order to make inferences about what is likely to happen in the future that we can anticipate what is likely to happen next in people's behaviours, in trends across a population as well as in relation to a particular individual or individual phenomenon. In that way, we can have more targeted, agile and future-oriented decision-making across the board in almost any industry you can name in any sector and in the delivery of services. We can all do that in an automated or partially automated way. It is an extremely powerful technology, and you can see very quickly why it is generating so much excitement because of its capacity to be seriously disruptive across the board.

What does that mean for law? It presents a number of opportunities. First, clients will take them up, which means more work for lawyers, and they will see that as a good thing. It is a new stream of revenue, and there are some very interesting legal questions, which I am sure they will love. Also, regarding the legal process itself, there are real opportunities, which are already being taken up, for optimising legal service processes. We see automation in document handling, document management, discovery, standard contractual drafting and so forth. Legal process optimisation is also possible.

On the risks, there are a number of challenges. They can broadly be divided into three. I will say a little about each. The first is that there are real challenges for what I refer to as constitutional and democratic values associated with these decision-making systems. Some of these can be understood as process values, so concerns about: the transparency of these decision-making systems; the explainability of their outputs and their recommendations; and whether or not the affected person, if there is an affected person, has the capacity to challenge and contest the decision and has all the due process safeguards that we are familiar with in public decision-making, at least; and concern that they are impartial and bias-free, and that things such as the presumption of innocence are not undermined by these predictive capacities.

On the question of substantive values, I think there are real questions about ensuring the accuracy of those decisions and that mistakes are not made. Substantive values such as privacy are clearly implicated in data protection, as is the potential to reinforce historic biases that are built into the historic data. It might be unintentional bias, but it is bias nevertheless.

Of course, there are larger constitutional values that we see in relation to the use of these systems, which have population-wide effects. We see them in relation to debates about fake news, to content regulation and to

the potential problems that that generates. This has implications for democracy, autonomy, dignity and solidarity.

The Chairman: Let me stop you there, otherwise we are not going to be able to make any progress on the questions.

Jeremy Barnett: I come at this from a different position. I think there is a tremendous opportunity here for the UK to lead the world in governance and regulation of algorithms. We have been particularly interested in the use of financial algorithms for automated trading, hence the concerns that I and others have expressed when high-frequency trading starts and goes wrong. I can perhaps develop that with some of the other questions. I am proud to say that we professionals believe that the English legal system is the finest in the world and is recognised as such. We think there is a great opportunity here to develop that skill. You may be aware of the recent announcement about a new IT court, down the road from the Royal Courts of Justice. The hope is that London will become the centre for this type of innovation, but the governance and regulation go hand in hand with the innovation. I see a tremendous opportunity if we get this right.

I can be brief on the risks. I think there is already a public loss of confidence in the transparency of decisions, which has been echoed this afternoon. There are concerns about who is going to set the agenda. Is it going to be the big five American companies, which have their own interests, or is it Parliament and the courts? My vote, as you will hear, is for Parliament and the courts.

The Chairman: Thank you. That is a very interesting set of answers.

Professor Reed, there is enthusiasm for legislation/regulation on your left, but you seem to be saying that regulation at this stage would be premature. Your papers establish certain principles that are very similar to those that Professor Yeung put forward. Are you saying that those should not be enshrined in law at this point?

Professor Chris Reed: I am saying that we should be very careful about enshrining them in law. I do not want to tread on a later question, but in many areas we already have existing legal systems that are largely able to cope with answering the questions we put to them and allocating liability and responsibility. Some comparatively minor tweaks can be made to them to make them work more effectively and efficiently, given the difference in the nature of artificial intelligence and the lack of information about how it is working and the difficulties to understand the workings of the technology.

The problem when you say prospectively, “We have this thing called an algorithm and it is dangerous”, and we regulate it, is that you end up regulating activities that we already are happy with regulation on, such as driving and using cars on the road, and the financial services sector, and you take it right down to regulating the contents of my smart fridge. If I have a smart fridge that is busy designing my diet and I do not like so much tofu and quinoa I can simply turn the fridge off. I do not believe that needs to be regulated. At the financial services end, we are not regulating the algorithm, we are regulating the financial services industry and, therefore, indirectly, its use of algorithms. The danger is that we

jump in and say, “We have this thing called an algorithm and we are worried about it”, and we end up regulating the whole of life.

The Chairman: We will explore all that further. First, Viscount Ridley and Lady Grender have supplementaries.

Viscount Ridley: Lord Chairman, you have basically asked my question, although I want to follow up on it very quickly.

Professor Reed, you said there was one example of prospective regulation that was a good thing. You did not mention what that was. Were you talking about the Warnock committee or something like that?

Professor Chris Reed: No, I was thinking of the electronic commerce directive. Mostly that directive says, “These things are possible. Here’s some law to make it easier to do the things that people are already doing and want to do”, rather than, “This is an activity that we somehow have to take control of and constrain to act in a particular way”.

The Chairman: You could say that the jurisdiction aspects in another directive are more important. That created more certainty than the e-commerce directive. That is a debate we could have, I suspect.

Viscount Ridley: You do not see bolting horses? There was a lot of talk of bolting horses in the last session.

The Chairman: That is a very good point, yes.

Professor Chris Reed: In the previous session, they were saying that the horse has bolted and we have to rein back before it is too late.

The Chairman: You do not see them trotting away at this point?

Professor Chris Reed: No. I can see areas where we need to get worried. I look at my Facebook news feed and I do not necessarily see what I want to see. I know that my Google searches are not the same as everybody else’s, and I do not know how they are filtering things out. That is starting to concern me, because I would like to know what is going on and have more control over it.

Baroness Grender: I have the same question.

The Chairman: We are all madly agreeing with our questions, in that case.

Q31 **Lord Swinfen:** You have probably answered a lot of what I am going to ask already. Does the ethical development and use of artificial intelligence require regulation? What should be the purpose of that regulation?

The Chairman: You might want to go on to add the supplementary questions, as they give a bit more granularity.

Lord Swinfen: I can do that at the same time. Would an industry-led, more voluntary approach be preferable? What are the advantages of self-regulation against government regulation?

Jeremy Barnett: I think the purpose should be safety, in the same way as we expect people to put safe products on the market, echoing what Professor Winfield said earlier. The Health and Safety Executive has

brought about safety on construction sites and has changed the way in which the world works. People used to be dismissive, but now people can work in safety because of the good work of a powerful regulator. There is also the Food Standards Agency and the Financial Conduct Authority.

If we really want the safe use of algorithms, there should be a regulator with powers. We are particularly concerned about self-learning algorithms where we cannot find the people to bring to book. If I can make my point simply, with safe products, as with the design of computer programs, there are designers, the people who commission the algorithms, the manufacturers, the people who implement them; in different parts of the food chain there are teams of people who play their part. What happens if the owner of the algorithm disappears, goes bankrupt or sells the algorithm to somebody else and the algorithm begins to learn and do its own thing, if I can put it that way? That is my concern.

The Chairman: Or, indeed, create other algorithms.

Jeremy Barnett: Yes. We get into a very interesting discussion—I do not know whether we have the time and the bandwidth today—about these new DAOs. The algorithms' own objectives will be to make money and will not be owned as businesses. For debate we wrote a rather mischievous paper saying that perhaps an algorithm, as a backstop, like a company, could be brought before a court. If a company commits an offence such as corporate manslaughter, not just the directors, the designers or the owners but the company itself can be subject to sanction of the courts. That is a debate that we ought to have.

Professor Karen Yeung: I would like to clarify a little what we mean by regulation. I have been doing regulation for 20-odd years now, and people mean different things when they use that word. Judging from the way in which this conversation is going, I think it is being used to mean something new, something additional to what we already have. Is that how I understand it?

The Chairman: Yes.

Professor Karen Yeung: It may take the form of legal rules that are enforceable or it may take the form of something else. If that is the understanding, it is fair to say that we already have some legal regulation of these systems when they are applied in a particular context. It is not true to say that there is nothing there. There is already some legal regulation of the systems that we have in place. The question is whether that is adequate, given the capacity, potential and risks associated with these decision-making systems. My concern is that they are already being used in ways that make really consequential decisions about individuals and have an effect on society and the community more generally. Again, the fake news debate is a great example of how there are real tangible consequences from using these systems, yet we have no real effective way of governing those.

In some ways the horse has bolted, and that is exacerbated, I think, by a Silicon Valley mentality. Facebook's motto is "Move fast and break things". They do not have a general approach of, "Let's get out there and talk to the regulators and find out what the problems are and come up with a consensual solution". As it is such a competitive field among the

big giants, they say, “Put it out in beta before it’s ready and let’s try and get the dominant role in the industry”. There is a problem. They are powerful, they are having profound effects and yet there is a vacuum in the governing of the values that are built in and the social consequences of them. The Government have a responsibility to do something.

The Chairman: We love a bit of disagreement among our witnesses.

Professor Chris Reed: This is very interesting, because I would not have any one-size-fits-all answer to this. The problem area is the area of news, freedom of information, dissemination of information, the Facebooks, the Googles and the social networks of this world. That is a difficult and intractable problem, as we have seen with the question of press regulation. It has been hard enough to deal with the press regulation issue. Now imagine that we start taking that out into the regulation of private discourse on social networks. That is intractable not so much because of the technology but because of the nature of the activity itself. If we look at the other kinds of things that we are examining, it is very obvious that there are areas where a society would want to regulate. The use of motor vehicles on the road is an obvious example. We regulate that anyway. If we are going to have self-driving cars, we want that regulated. If we are talking about medical diagnosis, we have a symbiotic relationship between the professions’ own self-regulation and the additional regulation that society, through the state, might want to put on top of that.

In some areas we are quite happy to have little or no regulation at all, because that is appropriate for that activity. I would look at the particular activities, and if I were being asked how I would decide whether regulation was needed or not, the obvious way to look at it would be through a risk-based approach: that is, both the risks to society at large and the risks to individuals. A high-risk approach is likely to be society wishing to regulate it in a fairly comprehensive way. A comparatively low-risk approach, or one that only risks individuals, or individuals suffering small financial losses, is another. For example, we regulate the financial markets heavily and we regulate eBay lightly. If you think about it, that is entirely appropriate, because my main risk on eBay is that I lose a few pounds. My risk on the financial markets is that I lose my life savings. That is more important.

The Chairman: You are not resiling from your view that we are not quite ready for additional regulation in most areas.

Professor Chris Reed: Seriously, we should not rush. There is a later question on which it is more appropriate for me to expand on precisely how I think it should be done and how the decision should be made. There will be some places where we need regulation up front to make the activity possible, but they are comparatively few in number, I believe.

Lord Swinfen: Mr Barnett, as I understand what you have said, and correct me if I am wrong, algorithms can change as they are used. You were asking whether an algorithm has a legal personality, and you indicated, if I heard you correctly—I am deaf, so I may not have done—that we should be able to take action against the algorithm.

Jeremy Barnett: Yes. For example, somebody invented an algorithm that could pick the best winners for companies in venture capital. They put the algorithm on the board, because they thought the algorithm could choose better winners than the humans. However, that is not allowed, because the Companies Act is being amended to say that to be a director of a company you have to be a human. You cannot have an algorithm as the director.

The Chairman: Horrific.

Jeremy Barnett: Yes. If I get my way, the algorithm could sit on a board of directors. That is just one example.

The Chairman: We cannot have a corporate algorithm that we can sue.

Jeremy Barnett: The problem about using the corporate route is that the company may have disappeared and the algorithm may then do its own thing.

Lord Swinfen: How do you punish the algorithm, if you need to punish it? The algorithm does not have any money of its own to be fined, as far as I am aware. Neither does it have a personality that you can actually put in jail.

Jeremy Barnett: If I might say, we have now got to the nub of the issue. We talked earlier about ethics and encouraging people to be trained to do things properly. My view is that the question has to be: what is the sanction for people who do not do it properly? There has to be a whole debate on the sanction. In regulation, when I go to a tribunal I try to persuade the tribunal that my clients have done nothing wrong, but if they have done something wrong and we move on to the sanction, what are we going to do?

Lord Swinfen: If you go back to the designer of the algorithm—

Jeremy Barnett: Easy.

Lord Swinfen: Yes, that may be easy, but the algorithm changes, as I understand it, of its own accord.

Jeremy Barnett: Yes.

Lord Swinfen: How could the designer be responsible for the changed algorithm?

Jeremy Barnett: Because he or she has designed an algorithm that can change. If he or she has disappeared, the algorithm has bolted. I would say, “Turn it off”—a kill switch. That is what they are talking about at the SEC for algorithmic trading systems, or high-frequency trading systems, that go wrong. You have to have a kill switch. You have to be able to say, “Stop”.

The Chairman: I can see the other witnesses twitching here. I must bring in Professor Reed and Professor Yeung.

Professor Chris Reed: I only wanted to say that the question about how we take money off the algorithm if we hold it responsible was fully addressed by AP Herbert back in the 1960s. In one of the collections of his stories, *Bardot MP*—I cannot remember which story it was—the

computer was so intelligent that it took over, decided the case against itself and then suggested to the judge that it made an attachment of earnings order. I take the view that you never apply law to technology; you always apply law to humans and the way they use technology, so there will always be someone who is using the algorithm on whom responsibility can be placed.

Lord Swinfen: The problem there is if the algorithm changes of its own accord without human intervention.

The Chairman: It is by design, is it not?

Lord Swinfen: I am not a lawyer, but I am descended from lawyers. Then, you have a very strange legal problem.

Professor Karen Yeung: The question whether you should ascribe legal personality to an algorithm depends on why you are asking the question. These questions are part of a broader set of quite difficult questions about the distribution of authority, responsibility and liability, and, in particular, if there is harm, who is going to be held legally responsible and ensure that the person who is innocent in that show is not left holding the baby. One possibility is that you might want to think about giving legal personality to the algorithm, but the core answers to that question must be driven by how you envisage the distribution of loss, liability and responsibility more generally. You have to focus on those questions first and think about whether you want a no-fault compensation system or a negligence-based compensation system, in which case the breaking of the chain of causation because of the lack of reasonable foresight is a problem.

The Chairman: Thank you. You have given us quite a range of options there.

Q32 **Lord Levene of Portsoken:** It is likely that AI systems could, at some point, malfunction—I reckon it is probably more than likely; I think they will, at some point, malfunction—underperform or otherwise make erroneous decisions that cause individuals harm. Do we need new mechanisms for legal liability and redress in these situations to be considered, or are existing legal mechanisms sufficient?

The Chairman: You can see where a lot of these questions are quite cognate; they are quite similar. If you want to expand on the particular nub of the issue, we would be very happy.

Professor Chris Reed: This is one I have done a lot of work on. Of course I think I have a very strong understanding, but then I would, wouldn't I? Do existing legal mechanisms work? The answer is yes, they are always going to work. The law will find a solution. If we have a liability claim, the law will find somebody liable or not liable. It will answer the question. However, there are two problems with that. One is in producing the evidence that is necessary to answer the question. Let us take the self-driving car as an example. Suppose there is an accident involving a self-driving car and it is alleged that somebody was responsible for the accident: the designer of the technology, the designer of the car, the operator of the vehicle who failed to intervene. It is alleged that some person should have behaved better.

To answer that question, we will need to dig into very difficult questions about how the AI is working, which may be answerable only by obtaining information from the designers who are from a different country. It will make litigation horribly expensive, slow, and very difficult. Where this as a problem there are some comparatively easy answers. Indeed, in the Vehicle Technology and Aviation Bill introduced last year, the answer was quite simple. It said in effect, "Okay, the insurers are liable to pay out in the event of an injury. There is no need to prove fault, and the insurers can have pretty much any debate they like with anybody else they think is responsible". As I understood the Bill, it was basically saying the insurer is liable if an accident is caused by a self-driving vehicle.

The Chairman: You are touching a raw nerve here, you know, because Lord Levene was chairman of Lloyds before.

Professor Chris Reed: From the insurance industry's perspective, so long as the premiums are raised to reflect the increased risk, Lloyds should be perfectly happy, should they not?

Lord Levene of Portsoken: I will not answer that. Professor, you were talking before in a slightly different context about the fridge that had a menu for you that you did not like. It has been suggested that the cause of the Grenfell Tower disaster was a fridge malfunctioning. I think it is unlikely that it was a very sophisticated fridge, but let us suppose that it was and that is what caused it. The implications and ramifications of that fire are enormous. How would one deal with that? You could say, "Go to the insurers", but where are they going to go?

Professor Chris Reed: In a way, you have answered your own question, because the fridge was thought not to have had AI and yet is still thought to have caused the whole Grenfell Tower disaster. We have an existing scheme that we are currently happy with. I have not seen any calls on the back of Grenfell Tower for the regulation of domestic appliances to come in. We already have electrical safety regulation for these things. If I have AI controlling my fridge, I am not sure it greatly increases that kind of risk. We have to look at the genuine risks that are created by these things.

One of the interesting things about the self-driving car or medical-diagnosis systems is that these are already areas where we heavily control the technology used, and as a society we are not going to allow the technology to be used until we are convinced it is safer than the current human alternative. We are going to allow self-driving cars when we think they are safer than human drivers. We already have a benefit to society from adopting the technology. We have to deal with the problem that we will still have accidents and we will want to compensate the victims. The question is: how should we compensate the victims? We could use our existing system of negligence. That is going to be made much more expensive and difficult by use of artificial intelligence. There is a very easy answer, which New Zealand already adopts for injuries: have a no-fault compensation system. In that area, there are fairly easy answers. In other areas there are more difficult answers, such as decision-making systems that infringe people's human rights against discrimination. That is very much more difficult to deal with.

Professor Karen Yeung: If I understand the question, it is whether or not existing legal mechanisms for distributing liability are adequate to respond to the challenges associated with AI-driven decision-making. Lord Swinfen has hit on the conceptual problem that we give these decision-making systems considerable decisional autonomy that they learn recursively and, therefore, that even the developer cannot reasonably foresee the outcome of that decision. According to existing conceptions of liability based on negligence, for example, that breaks the chain of causation and the developer is not legally liable. On the other hand, you could say that means there is a responsibility gap left by this problem. If the AI cannot be sued or held liable, the poor person who is injured or harmed is left with no recompense. That is the conceptual challenge.

My understanding is that I do not think that our existing conceptions of the liability and responsibility have yet adapted. Of course, as Chris says, if it comes to court the courts will have to find a solution, but somebody will have been harmed already, and there is a lot of uncertainty in the industry about how these decisions will be made. For reasons of legal certainty, it is in the interests of industry and the general public to clarify and provide assurance that individuals will not suffer harm and not be uncompensated.

The Chairman: Do you agree there is a responsibility gap? We have an opinion gap forming here, which is rather useful.

Jeremy Barnett: With the greatest respect, I do not agree. We are holding the line in relation to current law and regulation. We have done some work with the department for systemic risk at the LSE on the collapse of the markets. There was the crash of Knight Capital that you may have read about, where an AI system managed to lose \$450 million in half an hour, causing panic in the markets. The SEC led an investigation, which is available online, and it concluded that it was because of bad training. The algorithm was put into action before it was tested properly. That can be dealt with at the moment. We are just holding the line.

Our concern, put very simply, is that when these systems start trading with each other in clusters, it will be very difficult to work out what may happen. We will not be able to predict what will happen. At the moment, the only real way to test an algorithm is to fire dummy trades at the algorithm. People who write their algorithms will not give you the secrets. It is their secret source, people in this area say, so they will not allow people to investigate exactly why the algorithm has failed. The only way you can work out if the algorithm may fail is the testing regime at the outset. Work is being done on testing these systems, but our concern is that the gap will exist when they trade with each other.

The Chairman: Thank you. If you are not careful, we could spend hours firing propositions or hypothetical situations at you, but it is very, very interesting.

Q33 **Lord Hollick:** Mr Barnett, you offered the beguiling opportunity of leading the world in governance and regulation. Dr Mike Lynch has made a similar suggestion. How realistic is it that we could hold a leadership

position when, as one of you has already said, the commercial power lies in the United States? The United States is generally regarded to have a pretty much “free for all” approach, in contrast, let us say, to a more prescriptive European Union approach to this. We like to think that we sit somewhere in between, but how feasible is it to have regulation just in one country?

Jeremy Barnett: I will give a very short answer, if I may. In the regulation of cryptocurrencies and what are called ICOs—you may have heard of these methods of raising money for ICOs—there is regulatory arbitrage. Every country is setting its own bar. Do we want the companies here or are we worried about the effect? There is no standard regulation. People are saying that because the business is global there should be global regulation. Can London do it? Our view is yes, because the courts are just down the road from the tech community, which is near the City and the universities—King’s, UCL and the LSE are around the corner from each other—whereas in the States one of the big consortium in blockchain has moved from the States to London, because its clients are Wall Street but the tech community is in Silicon Valley in a different time zone. They have relocated to London. There is a genuine feeling, because we are good at law and the tech and we have all the players here, that the UK, not just London, can lead the world.

Lord Hollick: Have any other countries adopted a regulatory regime or framework that you admire?

Jeremy Barnett: Not so far as algorithms are concerned.

Lord Hollick: More generally.

Jeremy Barnett: They are all fighting over the cryptocurrencies and how to deal with those issues. There are papers about whether the SEC should begin to regulate algorithms, which I can let you have.

Lord Hollick: Talking more generally, not just algorithms, are there any regulatory initiatives that have been taken in other countries that we should be looking at?

Jeremy Barnett: Not that I am aware of.

The Chairman: Do either of you, Professor Yeung or Professor Reed, want to add to that?

Professor Karen Yeung: I would say simply that the European Union has definitely been the world leader on the regulation of automated decision-making, without question. There are some mechanisms in the GDPR. I do not think anybody has taken any steps in relation to meaningful ethical regulation of AI systems. There are voluntary initiatives abroad, such as the one the IEEE has talked about, but grappling with these questions is the Wild West; nobody really knows what data ethics is. I think there is a real challenge there.

Professor Chris Reed: As this is being recorded, I probably should not give my uncensored views on the GDPR, partly because it entrenches the system of computing that was in operation when I learned to program in the mid-1970s, and that is not the model that anybody uses today. There is real difficulty with the GDPR.

The answer to the question whether the UK could lead the way is yes, if it adopted something on which its thinking was global rather than parochial. We already have a number of examples. The law on the most complex form of electronic signatures is based on the model established by the state of Utah going off on its own in the late-1990s, saying, “Wouldn’t this be a good idea?”, and everybody else in the world picked up on it and said, “That’s the idea. At least someone knows how to regulate this”. Regulatory ideas spread very effectively, but they have to be ideas that can be adoptable elsewhere.

If we go back to data protection, without mentioning the GDPR specifically, and why that has not been copied everywhere—it has been copied quite a lot—the answer is because one of the most important parts of the world for information technology is the United States, and the idea of a centralised state regulator is political anathema there. You could implement exactly the principles of data protection in a different way that would be entirely acceptable to the United States. I had this debate with colleagues 20 years ago, and we agreed that it was rather like gun control; it is much the same kind of idea, but different traditions about how to do it.

The Chairman: People prefer to litigate rather than regulate, so to speak.

Professor Chris Reed: Pretty much, yes.

Baroness Bakewell: We speak of Britain leading the way, global outreach and such. Can you tell me imaginatively what would happen over a maverick or pirate country such as Kim Jong-un’s? What is he going to do with this whole new technology—unregulated and unwilling to be regulated?

Professor Chris Reed: I do not think regulation would affect what he is going to do, would it?

Baroness Bakewell: His own industries.

Professor Karen Yeung: All I can think of is applications as a form of cyber warfare. Cyber warfare is totally contemplable. As Chris says, there is not much we are going to do in relation to those particular challenges.

The Chairman: Sadly, we do not have extra-territorial jurisdiction. There we are.

Q34 **Viscount Ridley:** I was going to ask about autonomy and legal personhood. We have covered that ground quite well, so if you do not mind and with the Chairman’s permission I will ask you a completely unscripted question. The word blockchain has come up a couple of times in this session, which it has not in our previous sessions. I am curious about the degree to which the whole blockchain idea—smart contracts, distributed ledgers and that sort of thing—is part of the AI conversation or a separate one. Perhaps cheekily, are you as lawyers worried that blockchain is going to do you out of a job by doing away with the middleman?

Jeremy Barnett: I became interested in smart contracts, because I was told that it was the end of lawyers, but I think it is the beginning of lawyers; the lawyers have something else to argue about. Blockchain is very important. We at the CBC at UCL believe that it is blockchain in conjunction with other technologies—with artificial intelligence, with the internet of things—so that we can track people around the building, which relates to the Grenfell Tower issue that was raised, and in conjunction with building information modelling and the computerised design of buildings. A paper has just been written, and published yesterday, about the ownership of data, which, again, you discussed in the earlier session. The paper has looked at blockchain-enabled artificial intelligence. Blockchain provides trust, which, again, is one of the points that was made earlier. It builds in trust and it assists automation.

The Chairman: In terms of identity, and so on.

Jeremy Barnett: In terms of everything; who did what, where?

The Chairman: It is an audit trail, basically.

Jeremy Barnett: Yes, it is an audit trail of everything. That is why we feel that blockchain is very important. The new paper by Lawrence Lundy talks about decentralised autonomous organisations owning a person's own data, rather than the companies owning and benefiting, the topic you discussed earlier, the individuals would be able to monetise their own data and control their data. They can control their consent. One of the problems with the GDPR is the right to be forgotten. How can that happen on a blockchain in an indelible record?

The Chairman: You do not need to go much further. The next question from Lady Rock will be exactly in that area.

Professor Chris Reed: I have a paper on blockchain that will be available next week. It does not match up with my work on AI and liability, but there is a potential clear link there. What is interesting is that the benefits of blockchain, the audit trail, the identity, start to become much more difficult to achieve once you start using blockchain for what I describe as off-chain assets—things that have a real-world existence, such as land or even shares in companies.

The Chairman: Off-chain assets. We learn new vocabulary every day.

Professor Chris Reed: As one example, people get excited about the bitcoin model and the anonymity and distributive nature of what is going on, and the fact that it is a cryptocurrency, but if you start to think about, for example, keeping company shares on a blockchain—which is great, because you have all the auditability you need—suddenly you go, “Mm, now we're in money laundering territory. We need to be interested in identity”, and we are not very good at identity. I can tell you I am Professor Chris Reed, but how do you know? Here I am in front of you, and you can look up my photograph on Google. Turn that into the blockchain world. It is even more difficult there.

The Chairman: Professor Yeung, let us keep this to AI and blockchain, if we could.

Professor Karen Yeung: Absolutely. I think Jeremy is right that they are different technologies, but they could be integrated. You can have AI-driven blockchains—

The Chairman: Exactly. It is trust building, and so on.

Professor Karen Yeung: Exactly. They are slightly different, because blockchains are used as cryptographic algorithms; it is not machine learning in itself. That is why they are slightly separate. I am going to speculate here. The excitement is about unpermissioned blockchain, like Bitcoin, and I do not think that is going to fly. I think you will see a proliferation of *permissioned* blockchain systems, and it means that control within that private blockchain will be ironclad. Is that good or bad? Discuss. That is a great essay question.

The Chairman: Thank you. We will hold you to that forecast.

Q35 **Baroness Rock:** I declare my interests as listed on the register. Mr Barnett has kindly started the ball rolling on ownership of data. I would like to see if we can expand on that. We all recognise that it is a global issue. In fact, in written evidence, we have received many respondents who were concerned that public institutions are not being properly compensated for allowing private companies access to public data sets. My question is: when AI systems are developed using publicly-owned data, or even personal data, who should own them? Should we explore alternative models for individuals or trusts to retain ownership over that data?

The Chairman: If you were present at our previous session, you will know we had a little bit of debate about that as well.

Jeremy Barnett: It is a big subject. Who owns the data? The individual. For public data, systems have to be developed in order to allow the public to advance learning in, say, the health service without compromising the individual's rights. The concern I alluded to earlier is the idea of these decentralised, autonomous organisations controlling data, being programmed to send the data to certain places in advance and deciding whether or not that organisation can benefit. That is the concern. We cannot say where the money goes until we have systems in place that allow the control of the benefit. The building block is to design systems that allow people to control the data, and then the decisions can be made in each case as to where the benefit should lie.

Professor Karen Yeung: I wish I knew the answer. Your question about ownership of data in relation to both personal data and non-personal data is one of the biggest challenges of our generation. I do not think we are anywhere near solving that. The problem is the very unique nature of data. It is a public good. If we aggregate data we can generate incredibly powerful insight. On the other hand, especially with medical data, sometimes it is deeply intimate and personal, and people do not want to share that, but if we did share it we could generate amazing insight. People are also very mistrusting of the fact that that data may be used to prejudice their interests in ways that they do not understand or recognise. These are real problems, and I wish I had the solutions. I am sorry that I do not.

The Chairman: It is a very hot topic, though.

Professor Chris Reed: For publicly-owned data I have an answer, which is obvious to me as blindingly simple and completely unhelpful, which is that of course it should be shared ownership. You can say, “How should it be shared?” That comes down to a negotiation and a discussion about the benefits that each side receives, and how the private sector part is to be remunerated sufficiently and yet enough benefit comes back to the public sector. There has to be an individual discussion in each case. I do not see how you can set any fundamental principles up front that the public must own it, in which case your private sector partners will probably walk away from the deal. If you want the benefits you will have to share the ownership, and if you are sharing the ownership you have to negotiate. That is the only way of doing it. As I said, that is unhelpful, but I think it is the only possibility.

For personal data, it is really much more difficult. I have come out of an hour’s discussion with a PhD student about personal data being used as part of a big data set for improving healthcare, how that links to the right to proper healthcare under the EU charter and how we try to balance that with the right to data protection. It is a very difficult question, but it seems to me that a society is entitled to take the view that, for example, the wider public good in discovering a cure for ailment A is more important than my individual desire to stop my data being put into a data set for this kind of analysis. Society has to take that view collectively. What I object to is it happening quietly and behind my back and there being no discussion about the proper way of doing it and protecting me.

The Chairman: Thank you. We will move on to our penultimate question.

Q36 **Lord St John of Bletso:** In your introductory remarks you all spoke about the many benefits that AI will provide for the legal profession through natural language processing. Clearly, AI will make a lot of services more efficient. How do you anticipate this developing over the next decade?

Professor Chris Reed: It is a difficult question. I gather that Professor Richard Susskind has given written evidence and may come to speak to you. He probably knows more than anybody else about this. I am more of a sceptic than Richard is, and I have teased him over the years that his predictions that the benefits to wider society through the law becoming understandable to everybody and being able to be implemented by everybody are always “just around the corner”. I think there is a risk for the legal profession in the area of small disputes that to be honest are too low value for legal help to be available.

My example is the eBay dispute resolution system, which very simply gives you a 14-day resolution for free for any dispute that you have with the seller. There is no legal system in the world that can do that, let alone at that cost, but to do that it runs a very simplified version of law that does not give you many rights at all. You could imagine an AI stepping in, as a kind of form of voluntary regulation in an area, to regulate small disputes where, to be honest, the law is too complex, expensive and difficult for anybody to use. People need an answer, and

they are happy to settle for this answer. It will take a lot of the middle-ranking work; the stuff that lawyers do on document discovery and due diligence. A lot of the basic, paper-scanning and shuffling will start to go to AIs. The profession will do something else with the lawyers.

The Chairman: I do not think you are very far away, if you are talking about the next five years, from Richard Susskind.

Professor Karen Yeung: I do not feel that I have the expertise to talk about what I think will happen in the legal services industry, although what Chris says resonates with me. I imagine that all those really boring tasks that I really hated as an articulated clerk will no longer have to be done by a human. That is a cause for celebration, personally. I am a little more concerned about the implications of democratising the automated legal advice systems. Perhaps I am old-fashioned, but I believe that lawyers are custodians of the rule of law and the values associated with respect for the law. It is not instrumental and about doing the absolute minimum to comply with the law. The law, for me at least and I think generally, is a moral enterprise. I am worried that we may lose touch with the normative underpinnings of the law if we rely excessively on machines to generate decisions based on historical analysis of cases and legislation.

The Chairman: Thank you. I like that expression “moral enterprise”. It makes me feel better.

Professor Karen Yeung: I know.

Jeremy Barnett: Our area of interest is using the technologies to resolve big multiparty global disputes, not trivial online dispute resolution/consumer disputes. That is where we think there is real advantage: to speed up and allow access to stop cases going into the long grass—quick resolution—where there is a need for speed. That is a phrase that was used.

Lord St John of Bletso: Has the Law Society come out with a position paper on this at all?

Jeremy Barnett: There is a great move for consumer online dispute resolution, which is generally being supported. I was involved in the academic community for years but I am more interested in using the technology for the bigger, more complex and difficult disputes. The risk, while I have the floor, is known as predictive justice, which you may have come across, where AI is being used for a number of tasks, such as, when the judge has been identified, predicting the decision and combining statistical techniques to decide whether you are going to win or lose and how much you are going to win or lose so that the insurers say, “In fact, there is no point in litigating”.

The Chairman: You clearly have a very sophisticated set of chambers, if I may say so.

Professor Karen Yeung: They are already doing that. Those services are already available in the US.

Jeremy Barnett: I was delighted to learn that France in the 1960s passed a law that a machine cannot make a decision about a human. I,

for one, would love to see that law here in the UK. I do not like the idea of computers being judges and deciding sentence, and the like.

The Chairman: Thank you. I am going to bring in Lord Hollick, who has used an awful lot of lawyers in his time.

Lord Hollick: Have you deployed algorithmic dispute resolution? If so, what is its success rate?

Jeremy Barnett: No, we have not yet. It is too complex. We are looking for use cases. At the moment we have designed the systems in theory and we are in the university—

Lord Hollick: When do you expect to deploy it?

Jeremy Barnett: I am not going to be pinned down. It is a big issue, and it is not easy to say, but I would hope within a year.

Lord Hollick: Would it be binding? Could you enter into it only if it was binding?

Jeremy Barnett: No. The first step is to come up with a win-win, mediated proposal that the parties can accept. If they like that, a lot of work will have been done. If they want to fight, sue each other and deploy the services of my colleagues, at great expense, they can do so. We thought we would try and break in by a mediated win-win.

The Chairman: Lady Grender has the final question.

Q37 **Baroness Grender:** What is the one recommendation each of you think this Committee should make?

Professor Karen Yeung: Can I cheat and have two? Is that all right?

Baroness Grender: As it is nearly Christmas, go on.

The Chairman: One and a half.

Professor Karen Yeung: Thank you. I have to declare an interest: I was on the working party for the British Academy-Royal Society working group on data governance. I think theirs are good recommendations, so I would like to see an oversight body of the kind they suggest.

The Chairman: A stewardship council.

Professor Karen Yeung: Exactly. Secondly, I think it would be really useful to ask the Law Commission to take on the task of examining these questions about accountability, responsibility and liability, and the distribution of risk.

Professor Chris Reed: We did not talk much about transparency, which is a pity, because to answer many of the liability questions, some element of transparency in the sense of being able to explain how a decision was come to is necessary. I would like to see the Committee making some recommendations on how the designers and creators of algorithms could be incentivised to make transparency possible. I would point out that there are two quite different kinds of transparency.

The Chairman: That is very interesting; not regulated but incentivised.

Professor Chris Reed: You can incentivise people by regulating them and having a stick with a nail in the end. That is one way of incentivising. There are lots of other ways of doing it. I can think of ways of working with the liability system such that people would say, “If you cannot, in your algorithm, explain, after the event, why the car crashed, I am not going to buy it because my insurance will be too expensive”. That is one example. Another is to recognise the difference between asking for ex ante transparency—“you must explain this technology before you use it”—and being able to explain after something has happened how it happened. The former is very expensive to do, the latter is much cheaper. You must not mix them up, otherwise you end up insisting on something unnecessary.

The Chairman: Can you make them alternative legal liabilities, so to speak?

Professor Chris Reed: You could do, or you might say that ex ante is more important for something like public trust than it is for deciding liability.

The Chairman: You distinguish between the two.

Jeremy Barnett: I would ask that the Committee recommend the formation of one regulator of algorithms across the piece. That regulator can become an expert; it can decide whether there needs to be a heavy touch in, say, financial trading and a light touch in education. It can gain experience and then understand the difficulties and problems, and build trust and confidence in the public for the use of this technology going forward.

The Chairman: Thank you all very much. You have provided the second half of a very stimulating afternoon. I think it is wonderful that the legal profession is still going to have a future. I am not sure that sentiment is necessarily going to be shared among my colleagues. Thank you very much indeed.

Miles Berry, Graham Brown-Martin and Professor Rosemary Luckin – Oral evidence (QQ 181–189)

Miles Berry, Graham Brown-Martin and Professor Rosemary Luckin – Oral evidence (QQ 181–189)

Evidence Session No. 19

Heard in Public

Questions 181–189

Tuesday 12 December 2017

[Watch the meeting](#)

Members present: Lord Clement-Jones (The Chairman); Baroness Bakewell; Lord Giddens; Baroness Greender; Lord Hollick; Lord Holmes of Richmond; Bishop of Oxford; Lord Puttnam; Viscount Ridley; Baroness Rock; Lord St John of Bletso.

Examination of witnesses

Miles Berry, Graham Brown-Martin and Professor Rosemary Luckin.

Q181 **The Chairman:** Perhaps I may I give you all a very warm welcome to our Select Committee. This is the 19th formal evidence session for our inquiry, and it is intended to help the Committee to consider the education of students about artificial intelligence and in wider digital skills. Our witnesses this afternoon are Miles Berry, principal lecturer at the School of Education, the University of Roehampton, Graham Brown-Martin, an author and entrepreneur, and Professor Rosemary Luckin, professor of learner centred design at the London Knowledge Laboratory, University College London.

I am afraid that I have a little rubric to go through first. This evidence session is open to the public. A webcast goes out live as is and is subsequently accessible via the parliamentary website. A verbatim transcript will be taken and put on the website. A few days after this session, you will be sent a copy of the transcript to check for accuracy. We would be grateful if you could advise us of any corrections as quickly as possible. If after this session you wish to clarify or amplify any points that you made during this session or have additional points to make, you are welcome to submit supplementary written evidence to us. Perhaps you would like to introduce yourselves for the record, and then we will begin the questions.

Miles Berry: I am the principal lecturer in computing education at the University of Roehampton, working on our initial teacher training programmes. I am also a member of the board of Computing At School and the British Computer Society Academy of Computing, as well as the US-based Computer Science Teachers Association.

Graham Brown-Martin: I work in anticipatory research and foresight, with a special focus on education and technology. I sit on a number of advisory boards for educational technology companies providing

consulting services for a broad range of international government agencies and commercial enterprises. I am presently engaged as the chief education adviser of pi-top, a fast-growing British education technology maker and curriculum publisher.

Professor Rosemary Luckin: I have a background in artificial intelligence and computer science as applied to education and I am the president-elect of the International Society for Artificial Intelligence in Education.

Q182 **The Chairman:** Thank you very much. I am going to start with a very general question. First, what does AI mean for education in the UK? Do we need to rethink how and what is taught in schools to take account of AI, or is the current approach the correct one? I suppose one should also ask the question about AI assisting with education both in computer science and across the wider curriculum. Lastly, how much do AI and machine learning feature in the computer science curriculum at the moment?

Miles Berry: I am happy to answer this. In the short term, the applications of AI and machine learning will be in the field of supporting teachers who are working with pupils and students. That will have a number of dimensions to it. Some of it will be assessing children's work or supporting teachers in their assessment of children's work going beyond simple multiple choice and short response questions. Computer assistance in the marking of essays is something that we are seeing already, and I think that a number of teachers will very much welcome that sort of support, as indeed will students who would like to have that sort of immediate feedback from a machine about what their essay looks like.

We can see a potential in setting work for children to do, setting tasks, building up a profile of the learner and tailoring a journey to them from one learning object to another. If we reduce education merely to content delivery, machines can certainly help with that. We can see the use of artificial intelligence and machine learning in guiding teachers to where an intervention might be most beneficial; tracking a whole cloud of data to where a particular student is. Teachers will be able to say, "There has been a change in behaviour and a drop in that pupil's work over the most recent week or so". This is where a teacher can go in and focus support and intervention. There is also a role for responding to students' questions. We see already chatbot applications where a student asks a question and the machine responds on behalf of the teacher or lecturer. So there is plenty of potential to support teachers in their existing role.

For the longer term, it is very hard to say what the changes will be, but I for one would want to avoid a dystopian vision of the future where children are sitting in front of TV screens and having learning objects beamed into their brains. We need to keep true to the vision that education is not merely about content delivery.

In broader terms, we can see a tension between what is on the curriculum since the Education Act 2002, which talks about preparing school pupils for the opportunities, responsibilities and experiences of later life. For a child who started in the reception class last September,

what is that later life going to look like? The role of data and artificial intelligence is likely to be far more prominent in that child's life than it is in ours. How do we best prepare that child for those opportunities, responsibilities and experiences? Some understanding of AI therefore has to be part of their curriculum.

The new national curriculum talks about its aim being an introduction to the essential knowledge that children need to be educated citizens, introducing pupils to the best that has been thought and said that helps to engender an appreciation of human creativity and achievement. That is all well and good and I certainly see it as being a part of what schools and education should do, but I am not sure that we can best prepare young people for the opportunities, responsibilities and experiences of their later life by simply passing on the best from a previous generation to the next. There may be more to it than that.

The computing curriculum is part of the national curriculum, and there is opportunity in it to talk about AI and to teach children about artificial intelligence and machine learning. The scaffold of the curriculum provides a framework in which that can happen, but we are not yet seeing much evidence of its implementation in schools. I will be happy to return to this point later in the session.

The Chairman: We may well ask you to do that.

Graham Brown-Martin: I echo and will try not to repeat some of Miles's comments, but I concur with what he said. As a result of technological progress, we are at an important crossroads in education transformation. The question then, as Miles has said, is whether we continue with an education system that was refined towards the end of the 19th century and the 20th century and is based on what we call the standard or traditional model of education, which is the inculcation of facts and procedures. That model worked quite well for the 20th century. The role of the teacher was to inculcate facts and procedures that are provided by experts or an external agency such as the textbook maker. That would then be tested by a measurement industry, now a multi-billion dollar industry, which would ensure that the student has read the texts or had temporarily remembered those facts and procedures. We could continue down that road and think of education as merely a form of direct instruction: drill and kill or rote learning.

Certainly, if we believe that the purpose of education is to pass tests, we could probably simply replace human teachers with AI. I suspect that is not the purpose of education. I would argue that what is in most parent's minds is how we prepare children to thrive in a transforming world. In this century alone, any child in the education system today will face challenges that no other member of the human species has faced before. We are not talking only about the amazing opportunities and possibilities that are presented by the fourth Industrial Revolution, which AI is a part of. There will be jobs that we have not ever thought of.

However, there will also be downward pressure on income security and social agency. In the second Industrial Revolution, we were concerned about the inhumane nature of work—working long hours in factories and so forth. Now, we have to think about the inhumane nature of not

working. We talk about high employment rates in the UK, but we are also seeing zero-hours contracts and so forth. We have to think about the impact of these technologies, and beyond the opportunities and threats of the fourth Industrial Revolution we are facing what I consider to be existential challenges that are global: environmental change, population growth, antibiotic resistance and so forth.

The latest figures coming out of the UN suggest that by the end of this century there may be up to 1 billion refugees, up from the 30 million or 40 million that we have today. How will we cope with such things? There are some significant changes that I think children in school today and in the future will have to face. How do we equip them to face those? Will simply retaining knowledge temporarily and then regurgitating it at an arbitrary measurement be enough, or can we use artificial intelligence? Can we learn how to live with artificial intelligence so that we use education in a new way to solve new, abstract problems? Is it necessary for us to do in the 21st century the kind of mental arithmetic and things that we have done in later parts of the 20th century?

So the answer to your opening question about AI depends, really, on the purpose of education. If it is simply for passing tests, we can go ahead and have a technological deployment. This is not new; it has been happening since the 1960s in the form of computer-aided instruction. But it gives us the opportunity and the reason to transform our education systems. Technology is surely a part of that and it can be a positive future. There is an ethical issue with this. Let us imagine that AI is owned by a single corporation or collection of corporations. If big pharma also owned the hospitals and the doctors, I suspect it would be problematic. Or we could have an existing multinational textbook company—

The Chairman: I think we must move on from there. We may well come to it later, but we will stick for the moment to the curriculum.

Professor Rosemary Luckin: To summarise, there are three main elements that we need to think about. First, there is how AI can help in education. Secondly, there is how we need to help people understand AI. Thirdly, there is the impact of that AI within the world, what needs to be taught and how we need to teach it in schools. So there are three connected issues that relate to AI and education.

To unpack that, we need to understand a bit more about what we mean by artificial intelligence. For me, it is not just the technology—everybody thinks of it as the technology; it is also about intelligence. The only reason why we think that artificial intelligence is intelligent is that we have an impoverished conceptualisation of intelligence. Human intelligence is immensely rich and varied, and it includes social intelligence, emotional intelligence, self-efficacy and understanding yourself. Machine learning is only one form of artificial intelligence, but we cannot beat machine learning at learning in the way that it is good at. It is like me trying to run against Usain Bolt. He is always going to win hands down. AI will learn subject matter from the kinds of curriculum that we have in schools faster and better than us; it does not get tired et cetera. We need to recognise that and realise that it is a key issue.

The thing that AI as machine learning cannot do is explain the decisions that it makes. Human beings can. They can be taught to understand what knowledge means and how you justify that you know something. That is something that AI cannot do. If we want a knowledge-based curriculum, we need one that helps students to understand what knowledge is and how they justify it, because that is what AI cannot do and will not be doing in the near future. I will stop there. I am happy to expand on any of those three elements.

The Chairman: We may come back to them. I will bring in Baroness Grender first of all.

Baroness Grender: Thank you. That was a really fascinating introduction. Miles, you described the AI provision as scaffolding in the core curriculum. Does the panel think there is enough scaffolding at the moment in the core curriculum?

Miles Berry: There is enough opportunity within the programme of study as it is framed for teachers to deliver a very good grounding in at least our present concept of AI from a technological perspective. There are issues related to the ethical and societal implications that are not really addressed at the moment.

Among the aims of the national curriculum for computing we included the aim that children should learn to evaluate and apply information technology but including new or unfamiliar technologies—it was certainly not the aim to freeze it at that point, in September 2014—and analytically to solve problems. Even for children as young as five or seven, we included a requirement that they be taught to recognise common uses of IT beyond school. We did not specify what those uses were, but they now include various artificial intelligence machine learning systems.

I have a six year-old daughter, who is very happy talking to Google Now in our house. At key stage 2, the curriculum states that seven to 11 year-olds should be taught to use and combine a variety of software to create systems and content that accomplish given goals including analysing and evaluating data and information. The notion of a child, before the age of 11, creating a system that can analyse information—if we just take the words that are in the curriculum literally—has to include working with some machine learning technology, a naive reading of this would suggest.

The curriculum for 11 to 14 year-olds includes undertaking creative projects involving using and combining multiple applications to achieve challenging goals, including analysing data. Again, there is plenty of opportunity for teachers to address machine learning there. Even at key stage 4, where we step back from specifying lots of detail, we say that children and young people should be taught to develop and apply their analytic problem-solving design and computational thinking skills.

At A-level, we are seeing some wonderful opportunities for artificial intelligence projects. If we look at the specification from just one exam board, we see that some suggestions for projects include applications of artificial intelligence, investigating an area of data science such as Twitter

feed data and online public datasets, and investigating machine learning algorithms. So there is plenty of opportunity there. We do not necessarily have teachers who are confident in teaching artificial intelligence, machine learning and data science to the young people in their class; work needs to be done there. We do not necessarily have the resources that those teachers can draw on in providing such a curriculum. There is much else that needs to be done.

Professor Rosemary Luckin: Actually, I am going to disagree with you. I think we have an incredibly impoverished curriculum. What you describe for computer science is fine, but artificial intelligence is more than machine learning; it is about intelligence, it is interdisciplinary, it is about psychology, philosophy, linguistics and computer science. We have become hung up with machine learning, because it has produced most of the technologies now in common usage. It is running into the ground because it cannot explain itself. That is why DARPA is investing billions of dollars in what it calls XAI—explainable artificial intelligence—which is going back to using old-school, good old-fashioned AI methods because machine learning cannot do it. If we include in the curriculum only something in the technology sphere with respect to artificial intelligence, we are dumbing down the population, because we are not helping them to appreciate the broad spectrum of what intelligence is all about. Once again, we then end up thinking that things are intelligent when they are not; they are just very smart at a very particular kind of thing.

Viscount Ridley: If you had told children and students in the 1960s, when I first went to school, that by 2017 your fountain pen would know what word you were writing before you had finished it, that a device on the window sill could tell you the weather forecast if you asked it, that there would be an encyclopaedia that just sprang open at the right page—called Google—you would have described that as artificial intelligence, indeed bordering on magic, as Arthur C Clarke once put it. Yet, as several of you have pointed out, we have hardly changed the way we do teaching. Is that because we think that there are more important things to be done in school, as Professor Luckin implied, or because we are being unimaginative, as perhaps Mr Brown-Martin implied? I did not mean to suggest that you were against each other.

Graham Brown-Martin: I think we are probably in violent agreement. I tend to agree with Rose about the impoverishment. The computer science curriculum as we have it is fine as a sort of 20th-century attempt at what we think a computer science curriculum is because we are locked into an education system of a siloed nature. When we were at school, we all probably had maths, English, and computer science now. Computer science and AI come to life only when they are across the curriculum. We use technology and AI throughout our daily lives in different aspects of what we do.

I do not want to put words into your mouth, Rosemary, but I think that is what you are saying. STEAM, for example, is a buzzword in education right now, but it is very poorly understood. People think it is science, technology, engineering, arts and maths—separate things—but it is not. It comes alive only when you de-silo and get rid of the fake boundaries between the subjects. So I think you are absolutely right. Perhaps we are

looking at this in a very siloed, 20th-century fashion and wondering how we use computer science across the curriculum. Coding, for example, has become something that we are all supposed to be about, but it is just a notation, like the notation for STEAM, in the same way as there is a musical notation for playing an instrument. You would not normally start the notation before you play the instrument. It is a very demotivating way of doing things.

Baroness Bakewell: I wonder whether you are not in danger of creating your own silos. The most inspirational teacher is a completely integrated human being with an enthusiasm for everything. It seems to me that we need those teachers, and we need them to be well briefed. I can see a place for this exact use of AI in marking, correcting, talking, language correction and so on. That is one silo. The other is the emotional intelligence that interprets it, but do they have to be at odds? Can we not use both?

Professor Rosemary Luckin: You can, absolutely. The key thing here is that we can use AI to deliver a knowledge-based curriculum. That is no problem; we have systems that can do that perfectly well. In fact, the current knowledge-based curriculum is based on the same psychology and models of memory that AI systems were originally built on. We let the AI get on with that, which means that the human teacher can do all the rest of the stuff that you were talking about—the integrated, holistic approach to learning—and take students beyond what they are doing now. If we do not do that, we are basically dooming these students to trying to do what the AI is doing but not being able to do it as quickly or as well.

Baroness Bakewell: But is there not an opportunity to give AI the scope to do all the rote learning, freeing up teachers to be the sort of fulfilling teachers that we all want?

Professor Rosemary Luckin: Precisely, and to do the really important human interaction part.

Miles Berry: There is a famous quote in the domain that any teacher who can be replaced by a computer should be replaced by a computer. If we reduce teaching to something that a computer can do and to something that can be done by a dull machine, those sorts of teachers really ought to be replaced by computers. When I challenge my own students with this—our new, first-year teacher training students—they say immediately that there is much more to their role and much more to being a teacher than simply setting and marking work. There is teaching young people how to be a person. I think we are a long way away from a machine being able to take over that responsibility. I am talking about teaching young people how to get on with one another, and inspiring and motivating them. There are some great tools out there if you are already sufficiently motivated to learn. Gifted teachers provide that inspiration and motivation outside the sphere in which the child is already interested.

The Chairman: We are going to move on to our next question from the Lord Bishop.

interesting conversation. There has been an interesting and significant shift from teaching information technology to a more demanding curriculum. The Royal Society reported on that in November, as you know. It highlighted that the teaching profession is not keeping pace with the change in the curriculum. There is a lack of highly qualified teachers and poor availability of GCSE courses across the country. The Government announced new measures in the Budget and in the industrial strategy, but are there sufficient remedies for those issues? What more is needed? And is the national centre for computing, which was announced, a step in the right direction?

Professor Rosemary Luckin: I think you need to separate computing and IT education, about which Miles can talk eloquently—far more eloquently than me—and then artificial intelligence. In terms of artificial intelligence, I think that teachers' jobs are going to change the most—more than anybody else's—not because they are going to be replaced but because they are going to need to teach in different ways and they will need to teach different things. One thing that teachers need to be prepared for is understanding enough about artificial intelligence to help their students to understand how to use it. We need to help those teachers to understand how to use it and we need to make sure that we use AI to make us smarter, not to dumb us down.

The Lord Bishop of Oxford: Just to clarify matters, what you say is really interesting, but presumably you are talking about all teachers' roles changing, not simply those in computer science.

Professor Rosemary Luckin: Exactly. There is an idea that there is computer science and IT and then there is everybody else, but everybody needs to understand enough about AI to use it effectively and to make decisions about their own personal relationship with AI. They need to ask, "Do I feel happy about this system making this decision, and do I feel happy about this system having my data and manipulating that data to tell me things about myself?" Everybody needs to understand that. We need some people to understand how to build AI systems. But we should remember that a lot of these AI systems can be based on copy and paste codes. I am sorry, but that is true. It is not really so much about the coding and the tech when it comes to AI; it is much more about understanding the fundamentals.

The Lord Bishop of Oxford: And what should we be doing to equip all teachers to make that shift?

Professor Rosemary Luckin: As far as I can see, nothing, and that is really worrying. Yes, there is lots going on with computer science that Miles can talk about, but we are not really addressing the shift that every teacher, not just a teacher of IT or of computer science, is going to need to make.

The Chairman: Do either of you disagree?

Miles Berry: I do not disagree, but I would like to return to the question about what we are doing for computer science or computing teachers. It is worth making a distinction between primary and secondary education. In primary education, the class teacher teaches pretty much all the subjects on the national curriculum, so there are some wonderful

opportunities to integrate computing with mathematics, art and history, and to teach in a genuinely cross-curricular way. The problem in primary education is that there are very few primary school teachers with any background in programming or computer science. Thus, we have already had to address their professional development needs with projects that we have done.

The announcement in the Budget and in the industrial strategy focuses much more on the top end of GCSE education, where it turns out that only about 30 per cent of secondary school computing teachers have a background in academic computer science or a degree in computer science. So they find themselves teaching a subject for which they do not yet have the subject knowledge, and that is obviously a challenge.

I do not think that the picture is quite as bleak as the media coverage of the Royal Society report painted it. In fact, close on 70 per cent of year 11 students found themselves in a school where computing at GCSE was offered by the school. The problem is that only 11 per cent of students were choosing to study the subject at that level, even if it was offered by their school. So simply increasing the number of qualified teachers may not be enough to see a massive growth in the number of students who study this subject up to 16 plus. Qualified teachers will certainly help but there are other parts to this picture.

The GCSE qualification is a very difficult one. Perhaps some of the difficulty stems from the fact that students have not been studying computer science for the nine years leading up to the start of their GCSE course, although that will change over time. But it is also a qualification that has been tailor-made to cater for the most academic students. We made the case that this was at least as hard as GCSE physics. If we are serious about computing and not just computer science for all, we need a broadly based, although still academically rigorous, qualification that encompasses computer science and programming, but also its applications, and crucially, for me, the implications of that for them as individuals and for us as a society. There is a need to work on the qualifications base matters as well.

The £84 million is certainly very welcome, but how effectively that money is spent will depend on procurement and implementation. The simple calculation of dividing the money by the number of teachers that we need to qualify suggests about £10,500 per teacher. That is enough to fund the fees on a master's course.

Graham Brown-Martin: Perhaps I may add a bit on the qualification aspect, although I find myself agreeing with what has been said on both aspects. I wonder whether a qualification will be that useful. I have employed probably well over 1,000 software engineers in my career involving fast-growth technology start-ups. I have never looked at the qualifications of those software engineers; I have only ever looked at their work. It is the same today. When I hire software engineers, I look at their GitHub page. This is where they place their work. It shows whom they have collaborated with, whom they have shared with and how they have worked together. It is far more interesting for me as an employer to know what you can actually do. Can you collaborate and solve problems

with other teams of people? That is far more useful to me than a certificate.

Q184 **Lord Giddens:** First, I would like to congratulate you all on your brilliant introductions and say how much I agree with the central idea that there is a fundamental gulf between human intelligence and machine intelligence, because human intelligence is grounded in everyday society, history and institutions. So, for me, it is quite correct to say that that gulf cannot be breached.

My formal question is this. Given the pace of developments in AI and computer science more broadly, how can the curriculum keep up? I would like to add a question to that. As has been said, this is the first generation to have grown up in a truly digital world and we know that it has a very noxious side to it. We know that very young children are exposed to massive areas of addiction, pornography and so forth. How should schools seek to educate the very young generation as they enter school about these things? Do we need what Martha Lane-Fox calls “education and digital understanding” rather digital skills? She says that it is a much broader question, and I must say that I agree with her.

Graham Brown-Martin: It is a really interesting point. One of my issues with the computer science agenda, certainly in the past few years, has been that while it is great to be able to code and to understand how computers work, it is far more important for young people to understand that the digital world is a built environment in exactly the same way the physical world is and that it contains all the biases and other limitations of the physical world, as we can see in the media and so forth. That is not greatly understood by the population in general and specifically among young people. We are seeing issues with social media platforms, where hundreds of billions of dollars are being spent on creating behaviours that appear to be addictive; they simulate addictive behaviours. It would be wise for part of the way we educate our children to be about understanding how those things work so that they have some form of intellectual self-defence in that regard.

Professor Rosemary Luckin: Intellectual self-defence is a really good way of describing it. Understanding the limitations of technology is really important, as is being able to demand from technology rather than being demanded of by technology. That is part of the problem at the moment. Interestingly, another member of the Committee commented in his question that, “The computer will know the word I’m going to use”. No, it will not know it; it will be able to predict it, but it will not know it. It is because we have lost touch with an intellectual understanding of where knowledge comes from and what evidence is that people get sucked into things. They are not making good critical judgments about what the technology is offering them. They do not recognise the limitations of the technology and they are being sucked into it. It is problematic, and we need to put much more emphasis on this understanding.

Even the computer science curriculum, which is a real challenge, will never be able to keep up with the changes in technology that are going on. It will always be out of date on the specifics of the technology. But Miles is right to say that there are some fundamentals of computer

science that it is still really important to understand. Equally, there are some fundamentals of artificial intelligence that need to be included so that people can make more intelligent decisions about what they do with the technology and what they allow the technology to do to and with them.

Miles Berry: For me, that was one of the major shifts between the old ICT curriculum and the new computing curriculum. The ICT curriculum focused particularly on the skills of using technology and the drafters of the curriculum were careful not to specify particular technologies. However, implicit within it were the technologies that would be appropriate. What we have tried to do with the computer curriculum is to step back from the particular details of implementation to consider the fundamental principles of computer science. Yes, there are applications for getting useful things done, and perhaps we need to bring the balance back by doing more work on the tech skills, which are still crucial, but there are implications in this, particularly for the individual. There are also hints about considering some of the wider societal implications.

We introduced computer science into the curriculum from the age of five up. Programming was already there, but it gets a new emphasis in the computing curriculum. However, alongside that we have written into the curriculum a great deal about online safety, or what I would call digital literacy. It is about teaching children to keep themselves safe, and to act respectfully and responsibly when using online tools.

Of course, schools have always had a duty to safeguard and promote children's well-being, but the new computing curriculum has also given them the responsibility to teach children to do this for themselves and make sure that nothing bad happens to them beyond what goes on in school into their wider lives so that they have some sense of moral responsibility. Key stage 4 talks about new ways to protect privacy and identity. Even at key stage 1, which is for five to seven year-olds, children are taught that they should keep personal information private. What that means in this day and age of data mining and the digital footprint that a child leaves behind wherever she may walk has wide implications. There is enough in the curriculum for teachers to be able to provide that understanding of how to use technology wisely and help children learn to do that for themselves.

Professor Rosemary Luckin: One area that is not covered and which we do not acknowledge enough is the bias that exists within many technologies. Because we do not help students to understand how things like AI work, they do not understand that a lot of the systems they are using are deeply biased, and that is really problematic. A big piece of work needs to be done that goes beyond the core technology element. It is interesting to note that the large tech companies are increasingly employing more and more people who do not have a technical background. They are bringing in humanities experts because they are realising that they need a broader perspective in order to make sense of the technology.

Q185 **Lord Puttnam:** I totally agree with Lord Giddens and it is a refreshing conversation. In fact, Graham and I have been singing from the same

songbook for the past 20 years, so it is really nice to hear it.

Before we started this process, I reread CP Snow's "The Two Cultures", which was both refreshing and frustrating. Here we are, over 60 years later, and in paragraph after paragraph you are forced to say, "Oh my God, we've not got past that". What should the appropriate balance be between computer science and other subjects in schools? How can we be sure that the arts and humanities, which may become even more important with an increasing focus on creativity in the workplace, are not crowded out by computer science and other technical and scientific subjects?

Perhaps I may lead the witnesses to one specific area. I have been hearing a phrase recently that I like very much indeed, "the family of jobs" and the fact that employers and HR departments realise that you cannot employ people who will actually deliver a full outcome. Not many anaesthetists in an operating theatre are also surgeons and vice versa. This concept is quite a tough one for schools, yet if we do not begin to take it on and if young people are not encouraged to believe that there are wonderful jobs out there for them that rely entirely on there being four or five other people who do complementary jobs, the question I have just put becomes rather binary and dull. Perhaps I may ask you to comment on that.

Professor Rosemary Luckin: Interdisciplinarity is fundamental. Having made a comment about bias earlier, I should say that I am completely biased. I went to the University of Sussex when it was truly interdisciplinary, and I think it is a wonderful foundation, because you learn about the relationships between subjects and therefore realise that it is not just about the technology or just about the humanities; they have to be blended together and integrated. That is even more important now than it ever was.

Let us consider what has been said about social intelligence. Think about the sorts of intelligence that we as humans rely on, bearing in mind the social foundation of thought, which helps you to realise that there are other people with whom you need to relate and work: understanding what knowledge is beyond AI; social intelligence beyond AI; metacognitive beyond AI; metasubjective—understanding our relationship with the physical world beyond AI; metacontextual—understanding the importance of our physical experience and being in the world beyond AI; and the notion of self-efficacy—an evidence-based understanding about what we can do and what we cannot do, what we are good at and what we are not good at across disciplines rather than specific to disciplines. That is fundamental and goes beyond the technology. We have to start looking at a much more integrated approach that does not say, "It's arts or it's humanities", because that is the way of the future.

Graham Brown-Martin: The Harvard longitudinal study is the longest study of what it takes to live well in the world—to live a long, happy, healthy life. It turns out that it is not quality of education, it is not gene pool, it is not having a good job. It is about relationships. The single key indicator of a long and healthy life is good relationships, with your

partners, with your friends, with your neighbours and so forth, yet very little in our education system works towards that. I do not think it is an AI-related thing, but it speaks to what it is to be human.

The other point is that it is hard to imagine a society that does not have the arts. Again, I am not trying to be binary about this, but it is interesting to note that following recent changes in the British education system with the English baccalaureate, we saw a 24 per cent drop in the number of students taking drama, for example. Research from various parts of industry and organisations such as the Cultural Learning Alliance identifies the kinds of skills that we will need for the jobs of the future as being the things that machines cannot do. Among the human competencies where we beat machines are creative endeavours. That is creative thinking. It does not necessarily mean the arts with a capital "A", but when it comes to scientific discovery and entrepreneurship there has never been a better time to be an entrepreneur.

There is also physical dexterity—hundreds of thousands of years of climbing up mountains, dancing and swimming and so forth. A Dalek finds it very hard to go upstairs. We have that capability as well. There are social interactions—how we engage—which speaks to the Harvard study. My biggest concern is that in our attempt to adhere perhaps to a 20th-century industrial strategy, we are training kids to compete with machines when it is not an either/or. Art, like computer science and drama, needs to be across the curriculum. I would not be here today if I was unable to present myself.

Miles Berry: I fear I am going to have to agree. The breadth and balance of the curriculum is absolutely paramount. Primary and secondary education is too soon for vocational training, whether for engineering or for theatre. It should be a broad-balance, liberal education.

How can we see that in reality and ensure such breadth and balance? For head teachers, much of this comes back to the accountability measures in place. Ofsted is a particular driver of schools. It is great to see it recognise the importance of breadth and balance. We also have the league tables, which focus head teachers' attention. I wonder whether, if we had SATs in subjects such as art and music, primary education might be quite different. If in the EBacc performance measures the same attention was paid to the creative subjects, to art, to music, to drama, as we do to physics, chemistry, biology and computer science, I wonder whether we might have a very different experience of secondary education. It would be worth studying the effect that that might have.

The games industry talks about T-shaped individuals: somebody who has particular strength and depth to one dimension of what they do but who also has a great breadth. The depth is something that we can deal with later, through vocational training and higher education, but the top of the "T", the breadth, is the responsibility of schools. In primary, it is easier to link those subjects, because you have one teacher teaching all subjects. Why would we not link storytelling and programming? Linking art and mathematics, linking history and music; there are so many opportunities

to make these connections and avoid the silos which characterise secondary education.

Baroness Bakewell: Can I ask a squib of a question? I understand that computers can be taught to write Bach cantatas. What is going to happen to creativity? Can it be subjected to artificial intelligence?

The Chairman: You think that is a squib of a question, do you? I think it is a very sneaky question that could take some time to unlock.

Baroness Bakewell: I would just like to enter it into the debate.

Graham Brown-Martin: Are we going to have an examination system that allows us to take our AI into an exam? Would that not change educational experience?

Lord Holmes of Richmond: The AI fugue.

The Chairman: I think we may need to move on rather swiftly from that.

Q186 **Baroness Greender:** Professor Luckin, you introduced the issue of bias. I want to ask in particular about diversity. What concrete measures can be taken to get girls interested in particular in computer science and keep them interested in it right through to career level? What is currently putting them off? Is there a role for anything extracurricular and any organisations that can play a role in this and turn it around?

Professor Rosemary Luckin: It is such an interesting question. We debated this at UCL just a few weeks ago. I think that a lot of girls do not see how what they feel passionate about relates to computer science. That is the fault of computer science, not the fault of girls. As I said earlier, a lot of big tech companies are realising that they need people with a much broader understanding of the world. A lot of it is how it is positioned. There is a problem with role models. I think you have done quite well. I looked through the list of people who have come before you and you are up to 25 per cent female, which is not bad given the state of computer science and artificial intelligence in particular.

Too much emphasis is put on the technology, and certainly, with AI, too much on “It’s machine learning”. It is not; it is about intelligence in its full wonder. We need far more people who look at it that way. We have to present it in a much richer and more diverse way so that girls will think, “Oh, actually, that is something I am interested in and something I can contribute to. I can see how I could have a career in that”. It is about more than the technology; that is probably the simplest answer to that.

Miles Berry: I think we have done a huge amount already. By putting computer science and programming into the national curriculum, we have ensured that it is an entitlement for every child in a local authority school, irrespective of gender, home background or ethnicity and irrespective of whether they are interested in the subject or not, to be honest.

Professor Rosemary Luckin: If they do not take it, it is not accessible to them, Miles.

Miles Berry, Graham Brown-Martin and Professor Rosemary Luckin – Oral evidence (QQ 181–189)

Miles Berry: Well, give it time. We will have to wait and see. The anecdotal evidence from primary education is that girls are already very engaged in the subject. When we look at co-curricular provision such as Code Club, we see a very high take-up of places among girls. I do not have the accurate statistics.

Professor Rosemary Luckin: You said it: “Code Club”. It is not within the formality; it is the extracurricular activity that is richer and more interesting.

Miles Berry: I would argue that in primary education, in the best of our schools, we have that very rich experience of computing.

Professor Rosemary Luckin: In the best of our schools. I am sure you are right, Miles.

Miles Berry: At GCSE, it may well be another story, but I suspect that that is not so much down to the schools as to the awarding organisations. So one option for getting a more balanced intake when it comes to GCSE computing is to make it compulsory or close to compulsory. Mathematics at GCSE is 50 per cent male, 50 per cent female; physics is pretty close to 50:50 male-female, yet because computing, or computer science as the GCSE is now termed, has that niche position, girls are not choosing to study it. Twenty per cent of last year’s entrants were girls; it would be lovely to see that proportion increase. As the cohort of those having a very rich experience of computer education at primary school make their way through, more of them may be persuaded to study it. If we can get the qualification right, if we can get the GCSE spec right and have something that is not just computer science, which includes the programming but also includes media opportunities and the ethical/societal side of this, it will become a much more appealing qualification across the cohort to girls and boys.

Graham Brown-Martin: I think there is also an intersectionality issue. It is around the design of AI. It is vital that the intake to career is intersectional, not just male and female but across a diverse population. We are already seeing within existing technology how machine learning can make assumptions based on the data that it finds. We find the straight white patriarchy reinforced within the digital landscape. That would be an opportunity missed.

Baroness Bakewell: But that must be determined by the data.

Graham Brown-Martin: Exactly.

Professor Rosemary Luckin: I think that compulsion would be a disaster. If you made it compulsory to code, girls would be switched off. I can code. I have taught coding. It is not that hard, you know, and it is not the key to everything. Design is not about coding.

Graham Brown-Martin: And it is coding for what?

Professor Rosemary Luckin: Exactly.

Graham Brown-Martin: Just putting, “Hello, world” is not such a great thing, but if you are making a musical instrument, it is cool.

The Chairman: We have clearly touched a good nerve here, so we will

move on to Baroness Rock's question.

Q187 **Baroness Rock:** That leads very nicely into the question about qualifications. Mr Brown-Martin, you have made your views on qualifications quite clear, but perhaps I can ask the panel this question. Are the current qualifications at secondary and further education levels in computer science but also, importantly, because of the debate that we have been having, in associated subjects now fit for purpose?

We have not touched on the new T-level, which will be coming in in 2020, and I would be interested in your views on that. Do you think that it will provide a more vocational path, as opposed to an academic path, for those wanting to work in the area of AI? Perhaps, Mr Berry, I could start with you.

Miles Berry: When we were phrasing the national curriculum, we had in mind that there would be two parallel GCSEs. There would be one in IT for getting useful things done with computers, and the other would be in computer science. Key stage 4 states: "All pupils must have the opportunity to study aspects of information technology and computer science at sufficient depth to allow them to progress to higher levels of study or to a professional career". That is in the statutory guidance on the national curriculum, and I worry that some schools are only paying lip service to it, if that, and are not providing equality of opportunity for qualifications in IT and computer science, or, strictly speaking, the opportunity to study those things.

With the reformed GCSEs, computer science GCSE became a much harder qualification and then there was the removal of the IT GCSE, with Ministers being minded not to reform the qualification but simply to remove it. Alongside that has been the removal of many of the vocational qualifications in this space at 16 plus. If we are serious about this being part of everyone's entitlement, we need a broadly based, still rigorous qualification that takes on board the foundations, the applications and the implications, a little like GCSE science. Imagine the state of affairs if, instead of GCSE science, we had three subjects and you were expected to do physics, chemistry and biology, but then we removed the chemistry and the biology, leaving just the physics behind. That would be a very impoverished science education, and I worry that in the 16-plus space computing has been reduced to just a single qualification in computer science.

At 18-plus, the picture is much more complex. The A-level in computer science is certainly a very rigorous and very academic qualification but very few students are choosing to study it. The Russell group of universities do not require it for admission to computer science degree courses, and one or two are suggesting that they would rather students did not do computer science at A-level—they would like to have them fresh in the first year. So there are issues with academic computer science.

I think it is too early to give any sense of what will happen around the T-levels. I am encouraged that digital is one of the first to be introduced and we look forward to seeing the detail on that. I think it has the

potential to provide a very good vocational training for those who choose not to pursue the academic route.

The Chairman: Sadly, we have a Minister on the way, so we have to start speeding up the proceedings, I am afraid, but do you have any quick comments on that?

Professor Rosemary Luckin: I do have a quick comment. It is very easy and attractive for somebody in my position to say, “No, the qualifications are not fit for purpose”, but I am not going to go there. You cannot have no qualifications—that is not going to work—but we need to think slightly differently about the qualifications that we have. It would be perfectly possible to automate the kinds of qualifications that we have at the moment because they are built on a progression model that is automatable. We could be doing this on a continual basis without having to stop and test. It is all beautifully laid out and therefore perfectly easy to build a system that can do that for us. So teachers do not have to waste time invigilating exams; we can do away with that all. I think that that would be a good way to go because it would save quite a lot of money, apart from anything else.

Then we need to develop an enriched progression model that acknowledges other forms of intelligence beyond those that are purely to do with knowledge and integrated skills and that looks at the other types of intelligence. However, we need a progression model to sit underneath it, so let us automate the one that we have now, because we can do that, and then let us look for a richer progression model.

Graham Brown-Martin: As I commented earlier, generally in education we have a problem with the measurement industry, which needs to be addressed to get the sorts of transformations in education that this group may wish to see.

Q188 **Baroness Bakewell:** Are the ethical considerations currently taught as part of the computer science curriculum, and should they be? I bear in mind that we have heard a variety of witnesses who are concerned that the ethical considerations with respect to AI should receive more attention. It would be useful to know whether you think that they have already been addressed or whether more consideration should be given to them and, if so, in what way.

Miles Berry: The curriculum, as it is phrased, addresses personal morality in relation to technology. Children should be taught to use it safely, respectfully and responsibly, and to recognise what is acceptable and unacceptable. That is good, but it is not enough. The draft that we submitted five years ago as expert advice to Ministers went much further than that. We included as an aim that children should be taught to develop an awareness of the individual and societal opportunities, challenges and risks raised by digital technology, and that they should know how to maximise the opportunities and manage the risks. Seven to 11 year-olds should recognise the personal but also the social and ethical impacts of technology on their and other lives. Eleven to 14 year-olds should recognise the impact of digital technologies on society and the implications of technological innovation.

Regrettably, in my view, Ministers asked for much of that content to be removed to make room for more computer science elements, such as Boolean logic and binary arithmetic.

Baroness Bakewell: So those elements died, did they? They did not go further. We need some action on that, do we not?

Professor Rosemary Luckin: Definitely. If you want diversity and equality, there should be far more attention on ethics. It is fundamental. Everybody needs to understand something about ethics.

Baroness Bakewell: And will there be any agreement on what that ethical teaching will be?

Professor Rosemary Luckin: That would be a bit of a minefield, but it is a conversation that we have to start having, because it is really important. People will be put in the position of having to make decisions for themselves and for their family, but they will not understand sufficiently what has gone on. We have to look at ethics in some serious detail.

Baroness Bakewell: Where has this ethical teaching come from? Who is going to set the ethics?

Professor Rosemary Luckin: Exactly. We have to train the trainers and educate the educators.

Graham Brown-Martin: It is a societal issue. At the moment within social media platforms we are seeing the results of not having ethics, which is potentially very damaging. You are talking about a question for society to answer in the public domain about what our ethics are. Just because we can do something does not mean that we should do it, and I think we are on that cusp. This is called technological risk. We might do something really stupid if we are not careful, and we seem to have a propensity for doing that.

The Chairman: Thank you very much. The final, wrap-up question comes from Lord Holmes.

Q189 **Lord Holmes of Richmond:** If there was one recommendation that you thought this Committee should make at the end of our inquiry, what would it be?

Professor Rosemary Luckin: I think you should set up a commission on intelligence, not just artificial intelligence but intelligence.

Miles Berry: There are broad recommendations on the curriculum as taught in schools that I would love you to make. One is that no school should be regarded by Ofsted as outstanding unless its curriculum is as broad and balanced as the national curriculum.

Graham Brown-Martin: Reform the measurement industry and allow 21st century technology to be taken into a 21st century examination room. That would change our education system completely.

The Chairman: Thank you very much indeed. That was very crisp. Thank you for a very thoughtful set of discussions. We have really appreciated that. You unpacked quite a number of very important aspects

Miles Berry, Graham Brown-Martin and Professor Rosemary Luckin – Oral evidence (QQ 181–189)

of education, which has helped the Committee a great deal.

Professor Nick Bostrom, Professor Dame Wendy Hall and Professor Michael Wooldridge – Oral evidence (QQ1–8)

Professor Nick Bostrom, Professor Dame Wendy Hall and Professor Michael Wooldridge – Oral evidence (QQ1–8)

Evidence Session No. 1

Heard in Public

Questions 1–8

Tuesday 10 October 2017

[Watch the meeting](#)

Members present: Lord Clement-Jones (The Chairman); Baroness Bakewell, Lord Giddens; Baroness Grender; Lord Hollick; Lord Holmes of Richmond; Lord Levene of Portsoken; The Lord Bishop of Oxford; Viscount Ridley; Baroness Rock; Lord St John of Bletso; Lord Swinfen.

Examination of Witnesses

Professor Dame Wendy Hall, Professor Nick Bostrom, Professor Michael Wooldridge.

Q1 The Chairman: A very warm welcome to all three of you. While members of the public are coming in, I will make a few introductory points. The session is open to the public. There is a live webcast of the session which is subsequently accessible via the parliamentary website. A verbatim transcript will be taken of your evidence and will be put on the parliamentary website. A few days after this evidence session, you will be sent a copy of the transcript to check for accuracy and we would be very grateful if you could advise of any corrections as quickly as possible. After the session, if you want to clarify or amplify any points made during your evidence or have any additional points to make, you are very welcome to submit supplementary evidence to us. Could we kick off by asking you to introduce yourselves?

Professor Nick Bostrom: I am Nick Bostrom, a professor at Oxford University, where I run the Future of Humanity Institute.

Professor Dame Wendy Hall: I am Wendy Hall, professor of computer science at the University of Southampton and co-chair of the AI review by BEIS and DCMS.

Professor Michael Wooldridge: My name is Mike Wooldridge. I am a professor of computer science and head of the Department of Computer Science at the University of Oxford. I have been an AI researcher for 25 years.

Q2 The Chairman: The usual process is we will take a series of questions—I think you have had sight of them beforehand—around a number of Members of the Committee as we go forward. I know that we have to conclude this session at 4.45 sharp. Could I ask a very general question to begin with: what, in your opinions, are the biggest opportunities and risks associated with AI over the coming decade? Could I start with you,

Professor Bostrom?

Professor Nick Bostrom: Over a one-decade timescale there are exciting opportunities regarding basic research. It is really beginning to open up the box of how you produce intelligence that indirectly sheds light on how the human brain works. It is interesting in its own right. There are many economic applications. The first and most obvious locus of applicability is with the big tech companies. If you are Google, you can immediately apply it to make a better search algorithm, and, if you are Facebook, you can serve better ads and so forth.

Ten years may be enough to start to see an impact regarding self-driving cars. It is a little hard to say exactly how long it will take for that to be rolled out on a large enough scale to start to affect society, but it is possible. Surveillance is another area where there are obvious applications. There are massive data sets, text interactions or voice recordings of people and you are able to data-mine that more effectively. Facial recognition software, which is being rolled out in China to monitor the public, would be another strategically relevant application area within 10 years and autonomous weapons might be another. As with any new general-purpose technology, it might well be that the most exciting applications are not obvious at the outset and are only discovered as people start to play with the technology.

The Chairman: I am going to come on to ask about some of the limitations but, Dame Wendy, perhaps you would like to answer the same question.

Professor Dame Wendy Hall: To set this in context, AI has been around for a long time as a concept. I do not know how my colleagues would describe it, but we are in the fourth or fifth wave of AI. I cut my teeth on Prolog in computers in education as an MSc masters for my computer science thesis in 1984, and that would be classed as AI. What is happening now is that we can build systems that learn because of the huge amounts of data that is available, which is what Nick was talking about, and the machines, compared to when I started computing in the 1980s, are so much more powerful in terms of processing power and storage ability. That is why we see this huge surge now.

There have been surges before and AI is already in use in society in many ways. In the financial area, a lot of artificial intelligence is being used in the algorithms they apply to predicting and forecasting exchange rates and so on. Already in the health arena there are a number of applications. We are looking at escalation and acceleration. The genie is out of the bottle. We will see an increasing amount of automation in every walk of life over the next 10 years. How much of it will lead to great breakthroughs such as the automated car, I do not know. Most of Nick's examples were pretty negative—surveillance and weapons and so on—but there will be lots of positive benefits for health and from the discovery of knowledge. We need to get a grip of the downsides because it is happening so fast and there are lots of issues about algorithmic accountability, bias in data sets or in algorithms, and generally, of course, how it is going to impact society regarding future jobs. Clearly,

there will be job losses but there will also be a lot of new jobs, and it is about positioning ourselves for it.

The Chairman: We are going to tease out some of those issues later on. Professor Wooldridge, where do you come in on this?

Professor Michael Wooldridge: I largely agree with what my colleagues have said. The way I would think about AI is that it is a new software technology. That means it is software technology that lets you do things with computers that computers could not do previously. There is a whole range of things. We heard literally just last week that a long dreamed of vision of automated translation of speech is now a reality in a product that you can buy. That is a new capability that was not there before which is entirely due to AI. Whenever you have this capability to get software to do things it could not do before, it means there are new products, services and economic opportunities. There are huge opportunities there. This new software technology creates new economic opportunities for products, services and so on.

If I just take a UK slant for a moment, the UK is in a really unusual position, a position which as an AI researcher for 25 years, as I have already mentioned, I would not have guessed at even a decade ago. We are absolutely at the centre of that AI revolution. This city is at the centre of that revolution. Those opportunities are really in our hand as a nation right now, and we can make use of those and exploit them for the national benefit, or not exploit them. AI is software technology that lets computers do things they could not do, and that provides you with new services and economic opportunities to exploit.

Wendy is absolutely right that one of the most exciting specific opportunities over the next decade is health and personalised health. More of us, as they are becoming gradually more affordable, are buying things such as Apple Watches and FitBits, which are amassing huge amounts of very private data about us on a moment-by-moment basis. That data gets fed into a supercomputer in your pocket: your mobile phone. It is not really a phone—that is the most trivial thing it does—it is a supercomputer. That computer has an incredible capability to process that data and learn things about you that you would never have predicted. With AI you can potentially do things such as predict the onset of dementia and heart disease—all these things—long before you would ever become aware of the symptoms or think to visit a GP. It would be like having a doctor in your pocket. There is a huge opportunity there for us in health over the next decade.

The Chairman: That is great. I am sure we will give you a chance to expand on that later. I want to ask you to say briefly what you think are the limiting factors, Nick.

Professor Nick Bostrom: As of right now there is very limited ability to use reasoning and common sense and truly understand concepts and language. The strengths are in pattern recognition. It is an open question how much advance there will be towards solving those as yet unsolved problems over the next decade.

The Chairman: Do you agree with that, Dame Wendy?

Professor Dame Wendy Hall: I do. We use the term artificial intelligence, which has become a noun recently, and we talk about “an AI”—which I think came from Hollywood—but it is not artificial intelligence in its broadest sense. We do not have something that mimics what our brains can do, but we can now create systems that can do more than our brains can because the systems can learn from datasets over and over again and can be trained to be very accurate in particular circumstances, such as diagnostics in disease. Healthcare could be revolutionised. I am sure we all now look to Wikipedia or Google before we go to a doctor’s surgery. When I was a child you had no idea what the doctor was going to say when you went to the surgery. Now you have a conversation with them when you get there—“I think I’ve got so and so”—and that is only going to increase. Our first interaction with the health service when we have a problem could be online with the right support. If the UK takes that on board, we can really begin to have more efficiencies. It could lead to all sorts of things.

The Chairman: I love it; I ask you a question about limitations and you come up with more positivity. It is wonderful. Michael, just very quickly.

Professor Michael Wooldridge: Regarding technology right now, I absolutely agree that transparency is a big issue. AlphaGo, the famous Go-playing program plays world-class Go but it cannot tell you how it plays. You cannot extract the strategy. All its knowledge is somehow embedded within it in a really opaque way. This will be an issue if this software is used in circumstances where it is making decisions that have substantial consequences for people. Going back to the UK, I want to emphasise we are in this incredible position, but to capitalise on that we need capability. A DPhil student from Oxford who is doing a PhD in machine learning has a reasonable expectation of being made for life—being a millionaire—in a very short number of years after their DPhil. Why? Because there is so little capacity in this space and so few people have the right expertise. If we are going to capitalise on it, we need to build that expertise base. Other nations are busy doing that right now.

The Chairman: We saw you in the opening scene of AlphaGo by the way. Viscount Ridley first and then Lord Giddens.

Viscount Ridley: I wanted to pick up on something here before I move on to the second question. Dame Wendy quite rightly said that what is new is that these things can learn from scratch without having expertise injected into them. In that sense, they can learn anything.

Professor Dame Wendy Hall: They have to be programmed to learn.

Viscount Ridley: To learn, but not to have the expertise, not to be an expert on the subject, not to generate the expertise from the expert; whereas, Professor Wooldridge, in your written evidence you say there has essentially been no progress in general AI. Could I tease out the distinction here between them?

Professor Michael Wooldridge: General AI is the big dream of AI. It is the dream that you see in Hollywood; it is the dream of conscious machines, machines that are self-aware that can have conversations such as we are having right now. That is a very nebulous goal and I would say, possibly provocatively—and this would be the kind of the discussion over

a pint of beer—that there has not really been any substantial progress in general AI. We are beginning to get there with better ideas about the brain, but all the progress in AI over the last decade which is real and substantial and exciting has been on narrow AI. It is very narrowly focused tasks such as recognising faces or doing a very specific task, where you can get lots and lots of data about, “This is how you do the task”—you just repeat and repeat and repeat—“This is how you do it right; this is how you do it wrong”, and the program can learn from that data over time. General AI is the big, nebulous dream of conscious machines. It is my belief that that is not in sight at the moment. It does not mean it is impossible but it is not in sight at the moment.

Lord Giddens: I always agree with everything my friend Viscount Ridley says. I had just picked up on the same issue because I see your assertion that there is no progress towards general AI as rather central.

Professor Michael Wooldridge: General AI, but we need to be clear about what I meant by that.

Lord Giddens: I understand, but is that just a question of lack of advancement or is there a logical issue here? We know that John Searle’s Chinese room argument has been around for 40 years. As a bit of a Wittgenstein, I find it quite powerful. There is a distinction between syntactics and semantics. To be able to master meaning you have to be in the real world; you have to be an agent; you have to have a saturated knowledge of human society. I just do not see how AI could ever get to that. That is my question, really.

Professor Michael Wooldridge: That is not a small question to ask, I have to say.

Professor Dame Wendy Hall: It is quite philosophical.

Professor Michael Wooldridge: Nick probably has something to say.

The Chairman: Nick, fire away: treat this as a seminar.

Lord Giddens: It is a really big issue and structurally important for things such as how far it will intrude into job markets.

Professor Nick Bostrom: One of the things that is different about the current generation of machine-learning programs compared to the good old-fashioned AI is that at that point in the past, you had symbols that did not mean anything. The letter A could stand for “apple” and it could then make simple, logical derivations on that but the computer did not really know what an apple was and it could not recognise an apple if it saw one in front of itself. Now you have these grounded representations, where the representation of an apple is learned from looking at a lot of pictures and you can build artificial agents that run around in an environment acting and behaving and learning and maybe interacting with virtual apples. If you build a physical robot you could have it interact with real apples, so that the representation behind the simple A, as it were, now is linked to motor programs and the representation of visual recognition algorithms inside the AI in a way that was not present in the traditional AI system.

Lord Levene of Portsoken: Earlier you said that some of the

intelligence in the system is now embedded and it would be quite difficult to explain where it had come from. I find that interesting because I wondered to what extent any of you believe that this will be transformed over time in human evolution. The reason I ask that question is this: if I am trying to learn how to use a new piece of software—with great difficulty—and I turn to one of my grandchildren, they do it immediately. If I say to them, “How did you do that?”, they say, “It’s obvious”. It is obvious to them because, somehow or other—the way they have been taught and brought up—their minds are wired differently from ours. To what extent do you think that human evolution will be such that the understanding and know-how, which lesser mortals among us struggle with, two generations down will be self-evident?

The Chairman: That is another big question right off the bat. Does anyone want to tackle that one?

Professor Nick Bostrom: There is certainly human learning in figuring out what the capabilities are. One common error that people make when they see a robot performing a particular task is that they assume that if a robot can do this particular thing it can also do all the other things that a human would be able to do if a human could do that one task. That is not necessarily so. The bundles of competence that we find in biological humans, and even animals, might be very different from the bundles of capabilities we find in AI. We will learn more about that as we get more experience. There is also an active research area that tries to make these big neural networks more transparent and maybe make it possible to extract out simple, explicit logical rules from the big patterns that are learned. I think the sense in which these are less transparent than the old AI systems is that with the old systems, there was an explicit rule saying if you saw X, Y and Z, do a particular action, whereas now that is embedded in a big complex neural network that requires further processing to read out that principle.

Professor Michael Wooldridge: I absolutely agree with what Nick said. This challenge of extracting that knowledge and being able to reason at the language level in the way that a human would be able to rationalise what they were doing is one of the defining challenges for AI of our times, if not the defining challenge. We have had these incredible accomplishments, but, as I say, they cannot explain what they are doing or why they are doing it. To go to your question about human evolution, I think AI is increasingly going to play a part—and already is—in augmenting our intelligence in areas such as intelligent prosthetics. Google Glass did not quite work but something like it will, and it will overlay reality with a version of reality which is augmented by all sorts of information that is produced by artificial intelligence. That is realistic within the next 20 years.

The Chairman: Thank you very much. I have been reminded that all of us, including myself, have failed to declare any interests we have, which we are meant to do when we first ask a question at the very first session of this Committee. I am going to declare my interests and, Viscount Ridley, if you do that on the second question and then we can come back. I declare my interests as non-executive chair, Ombudsman Services; chair, Council of Queen Mary University of London; partner in DLA Piper

Professor Nick Bostrom, Professor Dame Wendy Hall and Professor Michael Wooldridge – Oral evidence (QQ1–8)

UK LLP; and co-chair, All-Party Parliamentary Group on Artificial Intelligence.

Q3 **Viscount Ridley:** In addition to those interests listed in the register, I declare that I was flown to Mexico to debate with Professor Bostrom last year and we had an enjoyable time doing so. I enjoyed it; I do not know if he did.

The second question follows on very much from what Professor Wooldridge was saying in his rather eloquent remarks about how well placed this country is to take advantage. Is this Government doing enough to maximise the opportunities and minimise the risks associated with artificial intelligence? If not, what more could be done?

Professor Dame Wendy Hall: I will take that question in terms of the AI review that we have been conducting. Unfortunately, I cannot talk to you, as you know, about any of the detail. It will be published very soon and I am happy to come back to talk to the Committee when it has been published to go into more detail. I have been talking to the Clerk about that. It is here. When you see it, I think you will see a very thorough review of economic opportunities and the potential for job creation and grasping the initiative and the opportunities that Mike was talking about. From my point of view, I have seen a lot of activity from government in the last four months because it has been very intense, and I think will continue to be intense for the next week or so, but I am afraid I cannot share that with you just yet.

The Chairman: We are going to make the assumption that there will be other forms of activity by government following on from your recommendations.

Professor Dame Wendy Hall: I hope so.

The Chairman: Professor Wooldridge, did you want to add anything?

Professor Michael Wooldridge: I am just going to perform according to stereotype. Clearly, education plays a key role here. If you walk around the DeepMind offices, basically everybody you meet has a PhD in computer science or a mathematical subject. Many of them are from overseas. We need capacity. It is not enough to turn out programmers. That is not it; we need programmers with very specific sets of skills and they are a level above raw programming skills. We need more PhDs with the right kinds of skills. We need to create an environment which is friendly for start-ups. At the moment there is an incredibly vibrant start-up culture in London, again not necessarily something we would have predicted 15 or 20 years ago. It is extremely healthy, but it is fragile and needs to be nurtured. Government need to create an environment which is friendly for those start-ups. I am going to use the word that nobody else has used, but Brexit is an issue. There are not that many British accents in DeepMind and if you want to protect that—believe me, Berlin and Paris would love to have DeepMind based there and would give them all sorts of breaks.

Professor Nick Bostrom: If one looks around the world, there are a number of countries which have caught AI fever and are going in quite ambitiously to support this. We see it in Asia with China and South Korea.

Canada is another nation that has been at the forefront. They have a very strong basis in machine learning and have decided to double down. They saw that their key university people were hired away by American companies and, rather than just sit back and lose their lead, they have recently launched a pan-Canadian artificial intelligence strategy with funding of 125 million Canadian dollars from the Government. Just a month later, there was an announcement from the Vector Institute with another 170 million Canadian dollar funding from the Canadian Government and the Ontario Government as well as a group of 30 industry partners. There are a number of initiatives to try to support the Canadian start-up scene in AI and to support Canadian universities. A senior academic active in this field could easily move into industry and get a half a million dollar a year salary from any of these big tech companies. To make it attractive for those people to stay in academia, one needs something different from business as usual.

Viscount Ridley: In recognition of what you have just said, Professor Bostrom, how does Britain look in that international comparison? Is it trying hard enough compared with, say, Canada?

Professor Nick Bostrom: I do not know all the different things that are going on within the UK. In particular, I have not seen this report. There has been a long investment over decades in computer science in Britain and that is the soil out of which these recent successes have grown, but my impression is that, aside from DeepMind, there would be a risk of losing the lead without doing what all these other Governments do where they see this as a massive national priority.

Baroness Bakewell: The talk is of PhDs for the start-up levels, which all begin in your late teens, at the earliest, and go on through your 20s. Are you aware that the Government are doing anything thorough in starting young people early in the understanding of AI? Is there any evidence they are acting on that already?

Professor Dame Wendy Hall: The Government have changed the curriculum for computer science in schools. We looked at this as part of the AI review. Now computer science is taught throughout, from primary through to secondary. That is still rolling out. That was only changed two or three years ago. It is very important that that include knowledge and awareness of AI. As far as I am aware, looking at the curriculum, that is being addressed.

Baroness Bakewell: Enough, do you think?

Professor Dame Wendy Hall: These roll-outs are very new. It was only about two or three years ago. I cannot remember.

The Chairman: The September before this one.

Professor Dame Wendy Hall: It is too early to tell. It is very important—and we talk about it in the review—that we keep an eye on that. I would also like to make the point that while we need to go aggressively for PhDs and masters in this area to give industry the machine-learning programmers they need, there is a huge amount we can do to upskill all sorts of sectors in the population.

The Chairman: Hopefully we will come on to that. Lord Levene.

Q4 **Lord Levene of Portsoken:** I declare an interest as chairman of General Dynamics UK Ltd. Given the possible impact of artificial intelligence on the labour market, should the Government be considering how to mitigate this now or is it still too early? If they should start planning now, what options might they consider?

Professor Dame Wendy Hall: I have thought about this quite a lot. One could be quite radical and opportunistic because, as with all big technological revolutions, over time there will be more jobs created than lost overall, but that does not help the people whose jobs are going to be lost in the short term, such as the truck drivers. Any job which is repetitive and can be replaced by some form of automation is likely to go in the next 10 to 20 years. We see the fight about taxi drivers in London at the moment, but with automated vehicles you will not need a driver, so that is a short-term issue.

We could be quite radical as part of social reform because I firmly believe that, while automation can help with a lot of things, I do not see in the next 10 years that robots would be able to help care for my mum when she was 96 and needed a lot of help. That is going to need people with empathy. Robots can do the process things and maybe even make a cup of tea, but not the empathy and the care that is needed. What we should be doing in social reform is taking the opportunity to say, "What are the jobs that need people to do them?" and value those jobs and not trying to say, "We are going to make everybody into an AI worker". We need lots and lots of those, but what is AI not going to be able to do and how do we fill those gaps? I think we could be quite radical about that, but I am not a politician.

Professor Michael Wooldridge: I absolutely agree. Wendy makes some very nice points very eloquently. We can either race to the bottom or to the top in terms of jobs. I would much rather see us nurture those start-ups and build the next tech giants and create wealth and jobs through that to offset these losses, which will come. I am more bullish than Nick about driverless cars. Within the next 20 years, I think that will just be the norm. The driverless car revolution is right here, right now. The technology is solved. It is just a matter, essentially, of rolling it out and tackling the big social and legal issues, of course, that go with that. Within the next five years driverless cars and vehicles on our motorways is entirely plausible. Within 20 years, I think it will be the norm. Within 50 years, our grandchildren will laugh at the idea that we actually used to drive our own cars.

Lord Levene of Portsoken: In that case, the Government should start retraining truck drivers and taxi drivers to do something else.

Baroness Bakewell: What do you think of the idea of a lorry driver caring for your mother?

Professor Dame Wendy Hall: There are all sorts of jobs that need doing to support people, and, actually, why not?

The Chairman: Nick, where do you stand on all this?

Professor Nick Bostrom: It makes sense to try to position the UK to take advantage of the AI economy. Over a timescale of 10 years,

although there will be impacts on the labour market from automation, it is not clear that it will be a dominant driver of worker displacement. Rather than specifically targeting AI as a source of unemployment, there should be a general labour market policy. There will be other factors—offshoring, Amazon displacing retailers and all kinds of other things.

The Chairman: Other technological changes.

Professor Nick Bostrom: And other economic factors that may be bigger over a 10-year time scale than the impacts from AI.

The Chairman: That is interesting. Thank you very much. We move on to Lord Hollick.

Q5. **Lord Hollick:** I declare my interests as set out in the register of interests. Data is one of the key feedstocks for AI, and so the ownership of datasets—your personal data—and how you transact with bodies and organisations that want to use your data becomes a really critical issue. These concerns came to the fore when we saw that the Royal Free allowed DeepMind to have 1.6 million NHS records. What, in your view, is the right approach to this? What are the right terms of trade? How do you protect privacy and the public interest? Do you allow this, once released, to be available for everybody, which possibly means that it will go straight to Silicon Valley and they would benefit from some of the economic consequences of this? What do you think about that multi-faceted relationship to protect and to ensure that this country, the public sector, benefits as much as possible from the use and exploitation of the data?

Professor Michael Wooldridge: The first thing to say is that this country since 1984 has had in place robust and quite far-sighted data protection legislation. That legislation has been updated twice since then, in part through European initiatives. It is pretty robust, sensible and pragmatic. It covers quite a lot of the bases and already deals with quite a lot of the issues that are thrown up by AI. AI-specific legislation is not the right way to go. I would look at our existing data protection legislation and ask what AI adds into this mix that we need to start thinking about. For example, suppose from just looking at your Facebook feed, it could predict that you are coming down with dementia, to use an example from earlier; do they have some kind of obligation to release that information to you? This is not information that you provided; it is information that they have derived from information that you have provided and its smart services have come to a conclusion. You can think of endless similar examples. Potentially, we need to look at existing legislation and think about its impact. As I say, we have excellent data protection legislation in this country—if we apply it. It is not, I have to say, being rigorously applied throughout the country at the moment.

Lord Hollick: Can you give an example of where it is not rigorously applied?

Professor Michael Wooldridge: There are issues about where your data is stored. I suspect that we would not have to leave this building to find examples of people using data storage services and it being difficult to find out where the data was—in California, Montana or another obscure

place. If you use a service such as Dropbox—and there are dozens of those things—it would be very hard for you to identify where your data is actually stored. For example, if I stored an email that I exchanged with my students on a server that was outside the EU, potentially I am in breach of that legislation. Can we all put our hands on our hearts and say we are completely compliant with that legislation?

Professor Dame Wendy Hall: And of course we have GDPR coming in very shortly, which is even tougher. First, the issue of access to data is really crucial, and a level playing field for small companies and research institutions in accessing data that currently only the big Silicon Valley-based companies have is something we address in the review. I cannot talk about it here but I am happy to come back when the review is released and talk more about the ideas we have. We have some quite interesting ideas about how we can tackle that.

I subscribe to the view that society needs to move to a point where we take ownership of our data. It is still unclear what the path to that is because you have market forces, government regulation, societal issues and personal views to take into account on this path, but at the moment your data is spread all over the place. You talked about the Facebook example and you are absolutely right, but I also have data in Google, Twitter and Amazon. There is data about me all over the place and the digital me is completely distributed. I do not think that is sustainable over time. There are lots of people working on personal databases. Tim Berners-Lee's big passion these days is the whole issue of how we retain our personal data and then have a contract with the world through whatever legal framework to allow other people to access that data.

Baroness Greener: I am going to take the very specific example of the NHS because it is unique in the world, in the sense that it is such a large organisation with so much potential data. In 100 years' time, would we look back and almost describe it as a "natural asset" that we sold on the cheap at a particular moment in time by releasing data on a piecemeal basis and by not ensuring that we recognise that this is an extraordinarily deep mine of data that we could commoditise and realise a significant asset from in the future?

Professor Nick Bostrom: I am not sure what the answer is. In terms of data being valuable as a one-off, to have a big base of training examples to train an algorithm that you can then apply might have a limited shelf life, in as much as there are many other places where medical data exists. China just put up online for free a dataset with 1 million to 2 million health records. At some point there might be enough in the public domain. If all you want to do is train your algorithm to recognise chest X-rays to see if there is some pathology, the incremental bit of data that the NHS is sitting on might not be valuable. There is a different model whereby if you want to serve the British people, obviously you need access to their health data to be able continuously to make recommendations to particular people. In that model, this data retains value and would create some bargaining strength that could be used to get the deal with private providers that would get value for the taxpayer or the patients in the NHS.

There is another important way to provide value to the citizen: giving trusted partners access to this data to enable new kinds of services, to enable the NHS to provide quicker diagnoses and better paths through the medical system and better indications of what might be needed. Rather than trying to tap every last penny of the economic value of this, if you could provide better medical care and advanced medical research, maybe requiring that companies publish some of the findings in scientific literature and so forth, that would be another path to try to eke out value.

Q6 The Lord Bishop of Oxford: I declare my interests. I am treasurer of the All-Party Parliamentary Group on Artificial Intelligence. My eldest son, Paul Croft, is co-founder and creative director of a cluster of games development companies of which the largest is Mediatonic. I would like to ask about ethical questions. How could the ethical issues associated with artificial intelligence best be addressed? Should some form of regulation be considered, either now or within the next decade, and, if so, what should this regulation seek to achieve? Would an industry-led voluntary approach, changes to corporate governance, et cetera, be preferable? Could you also comment on the need to educate the next generation of computer scientists in ethical practices and what is happening in universities? It seems to me we are still reaping the effects of a generation of bankers not learning ethics as part of their formative education. I would like some assurance about what is happening in computer science.

The Chairman: We never said the questions were going to be easy.

Professor Dame Wendy Hall: There are lots of questions there. I was president of the British Computer Society once and I can say absolutely that if your computer science degree is accredited by the British Computer Society, you have to be teaching your students about ethics.

The Lord Bishop of Oxford: That is really reassuring; thank you.

Professor Dame Wendy Hall: I think that most degrees in the UK—I do not know about Oxford—are accredited by the BCS.

Professor Michael Wooldridge: You asked that because you know we are not.

Professor Dame Wendy Hall: I was not sure, actually. I am sure you do ethics. It is a major issue. We absolutely need to instil not just in the students that we teach but in companies which are growing in this area some sort of accountability. I know there is a move in the EU Commission to look at algorithmic accountability. As Mike said, it is very hard, once these algorithms are let loose, to unpack exactly what they are doing. In the old expert systems you could in some way ask them, "How did you get to that conclusion?" The way we are building these things now it is very difficult. All companies should be aware of their responsibilities in this area and of the ethical issues in what they are doing, and that it is very easy for bias to creep in.

Professor Michael Wooldridge: To underline what Wendy says, if you do a computer science degree in this country you will almost certainly get training on ethical issues under the BCS code of conduct. The Association

for Computing Machinery has a more elaborate and detailed code of conduct which you can boil down to, “Don’t do bad things when you write computer programs. Think about the consequences of what you are doing”, and so on, and, as I am sure they do in Southampton, we expose our students to various case studies and talk through the issues that arise. It is a similar situation to the one I mentioned earlier. Those existing codes of practice cover a lot of the bases. Certainly, it would be worth going back and looking again at those codes of practice and seeing whether they are AI-specific. My suspicion is there would be different case studies, but the kinds of issues that arise would be very similar. Those codes of practice are out there. I am not convinced that AI-specific ethical guidelines are particularly necessary.

The Chairman: Or regulation?

Professor Michael Wooldridge: AI law, I think no. We could look at specific areas such as healthcare and data protection legislation. The insurance industry is waking up to the value of all that data that FitBits and so on are gathering about us and they can do some things which make me slightly squeamish when I stop and think about it. Looking at specific areas would make sense, but not a general AI law.

The Chairman: Nick, do you agree with that?

Professor Nick Bostrom: Yes. I would add something. My anecdotal observation is that there is a lot of appetite in the machine-learning community, within the culture, for AI being used for the social good. A lot of technical people are not primarily focused on maximising their income or providing a military edge to whatever their country is. I am going around the world to those conferences and there is a sentiment of, “We want to be responsible people”. There is still a lack of clarity as to exactly what that means, but I think the foundation is there, especially as one considers the longer-term prospects for AI—these more radical visions of what will happen if one really succeeds. It might not be too early to try to articulate a vision for how this could be for the common good and begin to bake in that commitment. It is not relevant right now to AI systems—it is not a law we are going to put in place next year—but we are ultimately committed to AI and this transition to the machine intelligence era being bigger than just something that will benefit one corporation or even one nation. It should benefit all of humanity.

Professor Dame Wendy Hall: For information, the Royal Society produced two very good reports and I am sure you are aware of them: the machine-learning report, which we used a lot for the AI review—it talks about the whole scope of this—and the data governance report which recommends a stewardship body for data. There was also a commitment in the Conservative manifesto for a data ethics commission. These ideas have not landed yet, but they are being talked about and I am sure someone will come to talk to you about them.

Q7 **Lord Holmes of Richmond:** Good afternoon. I declare my interests as set out in the register.

To turn to the international collaboration question, to what extent should the UK be working with other countries to shape the development of

artificial intelligence? Should the UK be looking to lead that debate or, as I imagine, do other countries have lessons from which we can learn?

Professor Dame Wendy Hall: We have already mentioned a bit about international collaboration. Mike said that we are in a very privileged position here at the moment and I think we are. It is a bit unusual, but because of the situation in the US there is an opportunity to grasp here. There really is nobody running science in the US at the moment. The Obama White House produced two very good reports on AI which we have used, but there is nothing really happening there. We need to collaborate. All scientists want to collaborate to move the subject forward, but it is also very competitive. I think it was Mike who said the demand for AI skills is huge from the companies. This is the new big-salary job. This is where people will go to earn a significant amount of money and it is very competitive. You talked about the success of Canada, but Singapore, China, companies based in America, ourselves and other European companies are all looking for the brightest and the best. It is very competitive and we need to be able to attract the brightest and best from everywhere to come here and train, and then to try and keep them here. It is really important. I do not know if this fits into this question, but I saw a graphic recently that London has as many start-ups in AI as every other European country combined, and we need to keep it that way.

Lord Holmes of Richmond: This is absolutely the appropriate point to say that.

The Chairman: Did either of you two want to add to what Dame Wendy had to say?

Professor Michael Wooldridge: I do not think I could have said it any better.

The Chairman: You were pretty upbeat earlier on.

Lord Swinfen: I have a declaration of interest, in that I run a charity that works entirely on email to provide specialist medical advice to people overseas.

Are there parts of the world that we should not be working with for security reasons? I am thinking, of course, at the moment of some of the Communist countries.

Professor Dame Wendy Hall: I thought you were going to say China. China is the obvious one at the moment. I was out there the week before last. I was in China and Singapore and I am on the AI advisory boards of both Xingwei University and the Singapore AI activity. The Chinese absolutely have set themselves the target of being the top in the world in computer science in the next 10 years, and that includes building up AI. The issue with collaborating with China, of course, is some of the security concerns with sharing data, but they have huge amounts of data to train algorithms on. It is an amazing country and this is a big area for them. That is the country I first think of in that sense.

Viscount Ridley: To follow up on the international collaboration point, in the written evidence from the Future of Humanity Institute, it says, "The UK Government could begin by making a commitment to fostering AI

research and development for the common good. Such a commitment would signal the UK's leadership in AI governance ...", et cetera. There is an analogy here with what we did in reproductive technologies with the Warnock report, setting out, "Right, we have sorted the ethical question; we know what we are doing; we are doing it the right way. Come and work here if you want to". Is that an appropriate way to look at it? Can you steal a march if you get that kind of thing right?

Professor Nick Bostrom: Not exactly, inasmuch as I do not think with AI you can work out all the ethics and put them in place. There might be something analogous we can do which is to occupy the moral high ground. In the UK we are in a very strong position to do that. Several world-leading academics—maybe I should declare a conflict of interest here—are active in research on ethics and social impacts. There is a group at Cambridge. DeepMind is very engaged with these issues up to the senior levels. It might be a comparative advantage inasmuch as, if there is this appetite among a lot of the leading research talent to feel that they are working for responsible employers and good purposes, then being able to have this parallel track—not just the technical expertise but making sure we are going to do this right, responsibly, safely and ethically—would enhance the ability to make progress on the technical side.

Professor Michael Wooldridge: One area for international co-operation involving the UK would be autonomous weapons. I am not talking about autonomous weapons that look at a human being and think, "That is Mike Wooldridge; he deserves to die," but much dumber things than that. Essentially, it is things such as flying landmines. Imagine a drone, the kind of thing that you can buy for £20 from a local shop, with a lump of high explosive on it, a camera and a Raspberry Pi. The AI software simply has to recognise a human being. It will fly around the streets of London and, when it sees a human being, not identify the individual, just that it is a human being, fly down there and explode. Just as the international community has, to a certain extent, stepped up with respect to landmines, we could do something similar and ask ourselves how we feel about British companies manufacturing these and selling them on the open market. There is a debate to be had there that should be had.

Professor Dame Wendy Hall: That is an extremely good point. The AI community is already trying to self-organise in this respect.

Viscount Ridley: A Raspberry Pi is a computer in this context.

Professor Michael Wooldridge: The capability to do that is absolutely there now. A smart masters student could do that.

The Chairman: We have two questions to fit in in about two minutes, so I am going to get Lord St John to ask his question and then I am going to ask Lord Giddens to ask his question and there will be a coda at the end from the three of you in a couple of sentences, hopefully. It will be quite tricky. Lord St John.

Q8 **Lord St John of Bletso:** I will be very brief. I declare my interests as in the declaration of interests and, until recently, I was an adviser to Silicon Valley Bank.

Professor Bostrom, you mentioned the funding initiatives in Canada; and Dame Wendy, you mentioned that we are in about the fourth or fifth wave of AI, and the limited amount of funding. You have all referenced what is happening in Singapore. Professor Wooldridge, you spoke about the need for more support in AI start-ups. The UK needs a much clearer strategy for AI. How should that be and how should it be promoted?

Lord Giddens: I have no interests to declare apart from being an academic, economist and sociologist interested in the digital revolution.

My question—and I will not ask the one directly on here because it has more or less been covered—is, how is it possible for people like me to work with people like you? It seems to me this is a divided field; you have the experts in the technological areas and other people working outside, and it is pretty hard to get effective collaborations between them. It seems to me really important that we try to do that. I think it is a really informative session and I am sure everyone is very grateful to you.

The Chairman: I will ask you simply to say what you think is the most important thing that should be done.

Lord Giddens: That was not my question, though.

The Chairman: No, but I think it still needs asking at the end of the session.

Professor Nick Bostrom: In developing some kind of national strategy, one could look at examples such as Canada and see what one can derive from that. In addition to government funding, I am not sure of all the details but in facilitating these industry/academia collaborations, some of this funding is coming from industry. We could enable, say, the government datasets to be more accessible in a responsible way and maybe have people in government or in the NHS who have some competence in this area so they can help facilitate the application of these machine-learning advances to provide better services. For that relationship to work, there needs to be people who know their stuff on both sides—on the industry side, of course, and somebody who can carry the public's interest and has some technical expertise. Finally, I would emphasise that, in addition to securing the lead in technical capability and entrepreneurship, there is this parallel track of doubling down on being a leader in the moral space with a long-term vision of AI as being for the public good of all of us.

The Chairman: You have wrapped up the answers to all three questions in your answer. Dame Wendy.

Professor Dame Wendy Hall: I have three points. First, yes, the Government need a strategy, and I hope when you see the review you will think there are the beginnings of that strategy there, and I am happy to come back and talk to you about that. Secondly, having the most money is not necessarily the best solution. It is like the tortoise and the hare. Spending that money wisely is the most important thing. I hope, again, we have put some discussions into the review about that. Finally, in terms of translation, you are not going to put people like us in front of the general public and say, "This is AI". Of course we can all do our

Professor Nick Bostrom, Professor Dame Wendy Hall and Professor Michael Wooldridge – Oral evidence (QQ1–8)

outreach lectures, but there are some very clever journalists who are coming up next who can help with translation.

The Chairman: That is a very good segue into the next session. Professor Wooldridge.

Professor Michael Wooldridge: My wrapping-up comments would be that we need to leverage the UK's extraordinary position. We are in an extraordinary position right now. We should be doing everything that we can to get the most out of it for the country. Concretely, what would I recommend? I said earlier that we need to build capability. That is not the same as having people who can program; they need higher-level skills than that. Machine-learning start-ups are not based just on programming skills; it is people with master's degrees, PhDs and so on. We need to support and nurture the start-up culture which we have already talked about today. As I say, we have an extraordinary start-up culture in the UK. We should do everything we can to nurture that and not drive it away because it is very, very fragile. There is a whole set of different debates, which I am sure you will discuss, around data and leveraging the data we have in a responsible way. That is potentially a real source of economic growth as well. In short, we need to leverage the UK's extraordinary position.

The Chairman: Thank you very much indeed. I am afraid it has been an incredibly high-speed set of questions, but you have given some very high-speed answers, so it has been altogether very satisfactory. I hope, Dame Wendy, you will be able to catch your plane as a result. Thank you very much indeed; we really appreciate your coming today.

Graham Brown-Martin, Professor Rosemary Luckin and Miles Berry – Oral evidence (QQ 181–189)

Graham Brown-Martin, Professor Rosemary Luckin and Miles Berry – Oral evidence (QQ 181–189)

[Transcript to be found under Miles Berry](#)

Dr Mercedes Bunz, UK Information Commissioner Elizabeth Denham and Dr Sandra Wachter – Oral evidence (QQ55–64)

Dr Mercedes Bunz, UK Information Commissioner Elizabeth Denham and Dr Sandra Wachter – Oral evidence (QQ55–64)

Evidence Session No. 7

Heard in Public

Questions 55–64

Tuesday 31 October 2017

[Watch the meeting](#)

Members present: Lord Clement-Jones (The Chairman); Baroness Bakewell; Baroness Grender; Lord Giddens; Lord Hollick; Lord Holmes of Richmond; Lord Levene of Portsoken; The Lord Bishop of Oxford; Viscount Ridley; Lord Puttnam; Baroness Rock; Lord St John of Bletso; Lord Swinfen.

Examination of witnesses

Dr Mercedes Bunz, Dr Sandra Wachter, Elizabeth Denham.

Q55 **The Chairman:** May I extend a very warm welcome to our first set of witnesses today? We have Dr Mercedes Bunz, Dr Sandra Wachter and Elizabeth Denham. Thank you very much indeed for coming along today. I have a little rubric that I need to read out at this stage before we come on to the evidence-taking. The session is open to the public. A webcast of the session goes out live and is subsequently accessible via the parliamentary website. A verbatim transcript will be taken of your evidence and put on the parliamentary website. A few days after this evidence session, you will be sent a copy of the transcript to check for accuracy. We would be grateful if you could advise us of any corrections as quickly as possible. If, after this session, you wish to clarify or amplify any points made during your evidence, or have any additional points to make, you are very welcome to submit supplementary written evidence to us.

We will probably have some votes this afternoon, so I should note that if there is a Division in the Chamber, the Committee will adjourn as soon as the Division Bell rings and resume after 10 minutes, although sometimes the congestion is such that it might stretch to a quarter of an hour, I am afraid. We will do what we can to get through the Lobbies at high speed. Would you introduce yourselves for the record and then we can begin with questions?

Dr Mercedes Bunz: My name is Mercedes Bunz. I am Senior Lecturer at the Communication and Media Research Institute of the University of Westminster. I was the *Guardian's* technology reporter before that. My research focuses on digital technology and how it transforms our society. My last book was on the internet of things, which I co-published with my

dear colleague Graham Meikle. I am also a member of the British Standards Institution working group on the internet of things.

Elizabeth Denham: I am Elizabeth Denham. I am the Information Commissioner for the United Kingdom. I have been in this role for 15 months. I come from Canada where I held a similar role for a dozen years as a regulator of information rights.

Dr Sandra Wachter: I am Sandra Wachter. I am a researcher in data ethics and a lawyer as well. I work at the Oxford Internet Institute at the University of Oxford. I am also a research fellow at the Alan Turing Institute and I deal with ethical and legal questions of big data, machine learning and robotics.

Q56 **The Chairman:** Thank you very much indeed. Let us kick off with a very general question. Who should own data and why? You might look in particular at the question of whether personal ownership of all data generated by an individual is feasible and, if so, how?

Dr Mercedes Bunz: I noticed when I researched our panel that our perspective of data is coming from different sides that, at the same time, are not positioned against each other. My focus is on collective data and the role that datasets play at the moment. Who should own this data? I see a very specific role for the state. Collective datasets have acquired a new value. Artificial intelligence is not just about managing information any more; it includes the creation of new information. If we look at organisations which have large, valuable collective datasets of big data, we normally look at Google, Facebook and multinational corporations, but the state is also an organisation that has this data. The data the state has are fewer in number and have fewer users—if we might call citizens such—but they are much more valuable. The quality of the data in transport, health and areas of public interest is really good. What are we going to do with this data and who should own this? One answer that I would give is the state should have a responsibility in owning this data for creating something in the public interest here.

The Chairman: It is not so much individual ownership as state ownership, or have I got you wrong?

Dr Mercedes Bunz: No, you have got me right. Individual ownership was discussed a couple of times in previous sessions. But from my perspective, collective data is of much more interest. With NHS data, for example, there would also be a privacy problem on a technical level, if we want to give the money to the individual but would like to make it anonymous at the same time. Technically, that would be quite difficult. We need to understand that a lot of the time, individual data is not of high value. It is the collective data that is of high value, and, therefore, we should not always discuss them together or link them in the wrong moment.

The Chairman: You have made that position very clear. We may well come back to you on that. Information Commissioner.

Elizabeth Denham: It is a challenging question about ownership. Usually, in data protection law we think about the rights of the data subject and we think about the responsibilities of organisations which

collect and process the data. It is more about who is controlling the data, who has custody of the data and, from an individual perspective, it is about what kind of rights people can assert on the use and sharing of their personal data. We think less about data as property and more about the balancing of the rights. The question is also complicated between the public sector and the private sector. My colleague is using the example of the NHS, and that is a publicly paid healthcare system, so it is important to think about the collective rights and ownership and the use of that data in the public interest. However, if you go to your physician and you provide information to them, the physician, arguably, owns the record that is created, but individuals should be able to assert rights over how that data and information is used. That is why I think ownership is a difficult concept in discussions about data.

The Chairman: You come on to interesting points that we spend many happy hours on in the Data Protection Bill, talking about what is public interest. We will probably start unpacking that later this afternoon. Dr Wachter.

Dr Sandra Wachter: I would very much agree with that. The idea of data ownership is very foreign to data privacy laws in general. The notion of the GDPR and what the European Court of Justice thinks of as privacy and data protection are much broader than protection. We always say “data protection” but we are not protecting the data; we are protecting individuals. It is about how you collect the data, what can be done with it and what can be inferred from that data. It is not so much about making sure that the data is safe; it is more making sure the individual is safe. Data protection is only one fraction of the right to privacy. The right to privacy is much broader. If you look at the jurisprudence of the European Court of Human Rights, this has to do with identity, having relationships, forming a family and having a profession. This all comes together when it comes to privacy. Data protection is only one element that allows you to do all that.

The Chairman: We sometimes have to provoke our witnesses to disagree with each other, but not on this occasion. Viscount Ridley.

Viscount Ridley: Can I push you a little further on this point and try to get clarity? Medical information that comes from me starts as my property. It then becomes the property of the NHS and therefore, as you rightly say, of the state. It goes off to Google and gets turned into something else. Value is added at each stage and is available, hopefully, back to me to cure my pancreatic cancer or whatever it might be. Why stop at the NHS stage, Dr Bunz, for example, and say that that is where the ownership belongs; or why stop at the individual stage; or why stop at the corporate stage? It seems to me that all three have a controlling influence, as the other two have suggested.

Dr Mercedes Bunz: I would not understand it as stopping at the NHS stage. I would say the NHS, or the state, should process the data or create the data in such a way that it can be harnessed further. Of course, the NHS is an institution that has rich value when it comes to data. It has really interesting sets of data if you are a computer scientist or a data scientist in the health sector. We have seen that with DeepMind. They got

1.6 million patients' data from the Moorfields Eye Hospital to train at pattern recognition on eye diseases. We want more of that. It is not ownership; the NHS—and the state—has a responsibility to create processes whereby we can make this data accessible for research but maybe also for business.

Lord Ridley: Am I right in thinking that none of you likes the analogy of property here? You think it is a misleading word.

Elizabeth Denham: I like “custody and control” rather than the legal concept of ownership.

Q57 **Lord Swinfen:** Do you think that the Government’s artificial intelligence review addresses the key issues relating to data and AI? Does it go far enough?

Elizabeth Denham: The independent industry AI report was focused on the need to increase capacity in the industry and move the UK forward to a leadership role in artificial intelligence across all industries. The focus was very much on capacity and standards and skills and what needs to be done in terms of an overall industrial strategic approach to artificial intelligence. There was very little in the report—and it is related to the terms of reference—about data protection and data privacy, save a recommendation that the Information Commissioner work with the Alan Turing Institute to standardise things such as algorithmic transparency. That is a very positive recommendation, but the missing piece of the jigsaw is that the UK has a lot to offer in terms of leadership around data protection in new technologies. That would be worth the Government exploring. We can bring that to the table as well as innovation, to lead the balance between protection of individuals and innovation. It is innovation and privacy, not innovation or privacy.

The Chairman: Do you like the data trust concept yourself?

Elizabeth Denham: I like the concept of data trusts and data hubs. Obviously, the devil is in the detail in the governance of data hubs and data trusts. It would also help if we all worked from the same definitions. My office very much supports the approach that more anonymised data should be made available for SMEs and researchers, and a levelling of the playing field.

Dr Sandra Wachter: The review addresses very important topics and is large scale and comprehensive. I agree that it is very much about upscaling and developing skills. One aspect that is evident is the call for increased research in the STEM fields. Social sciences and the humanities are underrepresented. We should also focus on funding those subjects, especially when it comes to AI. We are having a discussion about biases, negative impacts, risks and ethics. Those are questions that have been answered by the humanities—by philosophers, social scientists, legal theorists and political scientists. If we do not get them on board with that, if we do not have an interdisciplinary approach to these questions, we are going to leave out very important issues.

Dr Mercedes Bunz: I come back to the issue of data trusts. The report that Dame Wendy Hall prepared is an excellent first step. We should not forget that she calls it an industry-focused review, so we should not

criticise it for the things it does not want to be. In terms of data trusts, it is really important to make sure this is internationally tackled. The EU is working on something similar already, and that is a really important next step, but the state could also make sure there is no bias in data in areas of public interest, which I find quite interesting; there is a chance there.

Another thing that I find a bit difficult is that the review distributes the responsibility for data among a few institutions, all of which I value very highly, such as the Information Commissioner's Office—the ICO—and the Alan Turing Institute, and it comes up with an AI council, but there is no request to co-ordinate that. If we want to have a strong AI answer, we need central stewardship, as it was called in the review of the British Academy.

Q58 Viscount Ridley: Is the Data Protection Bill, as currently being considered in Parliament, fit for future challenges? What about the GDPR in that context? I know, Dr Wachter, you have written about the so-called right to an explanation and how this, potentially, is unfeasible in practice. I would be interested to know your general view of the Bill and the regulation as they are going through Parliament.

Dr Sandra Wachter: I have looked at the Data Protection Bill and one thing that is evident, and I think excellent, is that the UK has taken the opportunity to close some of the gaps in the GDPR. For example, Article 22 of the GDPR allows automated decision-making based on member state law. It does not say what safeguards need to be in place to meet this law; it just talks about "proper safeguards". In the current Bill, it has been proposed that after a decision has been made, the individual has to be informed about the outcome, which is new and better than what the GDPR currently offers. It also states that data subjects should have the right to ask that the decision be reconsidered, or that the decision not be made by an algorithm. Both those things meet and exceed what is currently envisaged in the GDPR, and that is excellent.

One of the open questions with regard to the GDPR and automated decision-making in research is that we see there are many exemptions when it comes to research and big data, which makes sense on the one hand, because of course we want to facilitate research, but we need a definition of what research means. Traditionally, research is associated with universities. Universities usually have ethics committees that review projects, look at the methodology and see if it is valuable and for the public good. Now we see companies taking up research, and the question is whether the research they do is also considered research and whether the same privileges should apply to them. That is an open question that we need to address.

Elizabeth Denham: It is important to remember that one of the benefits of the new law is that the GDPR—which was four years in the drafting—addresses profiling and decisions made by automated means and gives individuals new rights. That is a really important aspect of the GDPR, and I agree with my colleague that the Data Protection Bill takes that further. However, there are other tools in the GDPR that help with individuals' rights and push forward in an ethical and legal way decisions that impact significantly on people's lives. There are mandatory data protection

impact assessments of high-risk profiling. There is a requirement for companies and public bodies to do DPIAs, which are subject to review by my office. There are also new provisions in the GDPR that provide for codes of conduct and certification, so if an industry is moving forward on artificial intelligence in an application, our office can certify the approach the industry is taking when it comes to fairness in automated decision making. There are new tools in the law that I think it is worth this Committee considering as adding to the safeguards and protections, in addition to explanations of decision making and scrutiny.

The Chairman: This goes beyond Article 22 and Clause 13?

Elizabeth Denham: It does. I am saying it is comprehensive across the landscape of the law, and, because we want to ensure that there is fairness and transparency in decisions made by machines, these are the tools that will help us in the way forward, and they did not exist in the previous law.

Viscount Ridley: Can I press you on one of the supplementaries, and maybe the others as well? Is there a risk that the new data protection legislation might hinder progress in AI research? If I am going around the world looking for somewhere to set up my innovative AI business, is there anything in this new legislation that would make me say, "I am not sure I want to set up in the UK; the data rules are going to get in my way"?

Dr Sandra Wachter: I would say no. First, it is very important that we make sure that our data protection standards are in compliance with what the European Union wants. If we want to be part of the digital single market, in some sense, we have to make sure that we have the same standards, otherwise we are not going to be able to compete with that. At the moment, we can see countries around the world, in Asia, the US and everywhere, looking to Europe to take the next step in data protection laws. If we are not at the same level as them, we are not able to compete with that. The Data Protection Bill is an excellent step forward. I have also read the amendments that have been made and it just shows that we are taking it seriously. Fairness, transparency and accountability are issues that all Governments around the world are dealing with, and if we take this first step forward we can set the bar for the rest of the world.

The Chairman: Dr Bunz, do you have anything to add to that?

Dr Mercedes Bunz: An important aspect is that all companies that develop artificial intelligence services want to create a product that they can sell to a wide market and not just the UK market. Having a law that is in line with the law of the EU and mirrors what is going on helps more than it hinders, I would say.

Q59 **Lord Giddens:** How can we improve access to data for researchers and SMEs? How do we prevent monopoly ownership of data by large corporations? I hope the Chairman will not mind if I bracket with that a request for a comment on this huge wave of reaction against the giant platforms in the US, which looks unstoppable. They have gone from being approved by everybody to the tide turning against them. A story in the

news today says that 126 million users have seen content uploaded from Russian sites on Facebook. How would you think your way through what is, in terms of traditional corporations, clearly something close to a system of monopoly?

Elizabeth Denham: It is a vexing problem and I agree that people are becoming more and more concerned about information monopolies and large platforms that offer many different services and collect and link all kinds of personal data. What I hear is the public expressing concerns over what used to be neutral platforms; people are now looking for them to act like information fiduciaries. These new, big platforms have a responsibility to make sure that the data processing is fair, especially if it affects our voting system and our news system. There is a concern about monopolies.

Data protection does not necessarily give you a way forward—as long as the companies are collecting and using the information in accordance with the law. Data protection does not stop mergers and acquisitions where the purpose of the acquisition is to exploit more personal data. That could be an area to examine. The competition authorities have a role to play as more and more data is put together—what that means for the individual. There are some regulators that could look at this, but it is a vexing problem.

Dr Sandra Wachter: I very much agree. The question of monopolies is a question we need to answer with competition law. Data protection law is probably not aimed at that. However, what the GDPR will give us, for example, is the right to data portability. Article 20 will now mean that an individual can easily take their own data from one data controller, one company, to another. This could enhance competition in a very healthy way. I would have the opportunity to go to Google and request all my data, get my data deleted if I do not want it any more and go to a competitor—it could very well be an SME, for example—which could very helpfully increase competition without the need to regulate all that. There is a potential for individuals to have a say in how their data is being handled in the future.

In terms of having proper access to data for researchers, we could think about having public access or public data repositories, something similar to the biobanks, where you have a trusted third party which holds data and grants access to bona fide researchers, for example. It has a lot to do with ethics. You would need to lay out your research project, explore your methodology and go through an ethics committee. If you pass all that, you would get access to the data and afterwards you would be required to publish those results so that the greater research community could benefit from that. Such a model would work with data in general and enhance collaboration between researchers.

The Chairman: Dr Bunz, did you want to add to that?

Dr Mercedes Bunz: To start with the last point, there are several research networks already in place that help us to access public data that we can use. For example, there is the Administrative Data Research Network that helps researchers who are interested in accessing public data. We could think about whether we need to make those networks

stronger and how important they are. There are also several accelerators that help SMEs and businesses to access the data of the NHS, to create small start-ups that help the NHS with certain problems, some of them using artificial intelligence, such as DigitalHealth London or MedCity. These are projects that are worth looking at and worth having more of. One should not be shy.

There is a danger of monopoly of ownership of data for artificial intelligence. We can see that clearly in natural language processing. The competition there is nearly decided. It is really hard to create natural language processing databanks. There have been many attempts and lots of money has been spent by Google, Facebook, Microsoft and IBM—Watson, for example. We know those programmes now. We can not compete with the research there any more, but we can compete with other research in transport and health. There is no ownership, no easy monopoly of data ownership. There is a good chance for us with more interesting data.

The Chairman: Lord Giddens, did you want to ask a supplementary on that point?

Lord Giddens: This is a neutral-sounding discussion, but the issues involved here, to quote Donald Trump, are absolutely huge for democracy, and for so many things, because you are creating platforms which claim to be neutral but are plainly not neutral in their impact. That is quite different from traditional publishing, for example, where you have monitoring authorities. You simply do not have them at the moment. That is what has allowed Russia to possibly sway the outcome of the American election. These are global corporations. We cannot just talk about the UK. The issues that we have to confront seem to me to be very deep and structural and are somehow integral to all this. We are all struggling to master them, but we must because the problems they have created are so fundamental. It is all so vast and all so global.

The Chairman: I am going to ask the Bishop to add to that and then maybe you could comment.

The Lord Bishop of Oxford: On a similar line, I am looking for a sense of how urgent you think addressing this question is. I do not get a real sense of urgency yet from what you have said. We seem to be saying that the way we have been handling data is leading to these large data monopolies and excluding SMEs from accessing datasets in a way that is critical to their life and flourishing, but I do not hear you saying, "This is dreadful, we must address it really urgently for the health of the world and the common good", which seems to be the way public opinion is moving.

The Chairman: Would you like to address what Lord Giddens and the Bishop have said, Information Commissioner?

Elizabeth Denham: We launched an investigation in May of this year on the use of data analytics and micro targeting of voters in the referendum and other campaigns. We are looking at data analytics companies, political campaigns and social media companies, because it is a bit of a black box for the public to understand how data analytics works by using platforms such as Facebook or Google to deliver political messages and

political ads. Where I can go is to see how personal data was used by the platforms to target messages or micro target voters and individuals to persuade them to vote in a certain way, to send a certain message. This is a really important investigation for my office. It can pull back the curtain to show how the data is used and how data analysis is happening with these three big players in the marketplace.

The Chairman: Do you think that should involve them making public some of the data they have, subject to anonymisation and so on and so forth?

Elizabeth Denham: We will be looking at how the data is used, and there will be a public report at the end of our investigation, but we are talking about two different things here. We are talking about how SMEs and researchers can get access to data so that we can grow the economy, versus what is happening behind the black box in some of the platforms that affect our lives. Are we talking about two different things here?

The Chairman: Information Commissioner, you have many things on your plate, not least the Data Protection Bill currently going through and the implementation of GDPR and so on. How concretely could you address that particular area, which is clearly of huge importance? Around the table we all have a great interest in this area.

Elizabeth Denham: We have dedicated a large team in our office to carry out the investigation that I just described to you, but our office does other work: giving advice, setting up codes for anonymisation, auditing anonymisation, giving advice to Government in the setting up of data hubs and data trusts. We are involved in all those things but right now, with the GDPR and the changes in the law, it feels a little like changing a tyre on a moving car.

The Chairman: That is a good analogy. Lord Levene.

Q60 **Lord Levene of Portsoken:** This is something of an overlap with our first question on who should own the data. How can data be managed so that it is used for the public good?

Dr Sandra Wachter: That is an excellent question. The one thing that we need to start with is defining what “public good” means. That is not really clear yet. We do not have a definition of the public interest or the public good. Potentially, the whole of society would benefit from it. I can see applications for climate change or healthcare, for example.

One thing that the new framework—the GDPR—will mandate is the data protection impact assessments, which have already been mentioned. It is a very good step forward because it forces anyone who is doing big data analytics to think about the possible consequences of the applications before they are even allowed to do it. In some cases, they might even have to go to the data protection authority and ask for permission to use that if they fear that the outcome of the application could be harmful for the individual. This is a very good step forward from a regulatory

perspective—thinking about the ethical and legal implications of a product before it is deployed. Ethical and legal foresight analysis is the right way forward, rather than reacting after something happens.

Lord Levene of Portsoken: You are saying, rather than the company or the individual who has compiled the data taking their own view on how far it should be protected, there should be some overarching body to whom they are responsible, because otherwise an individual might say, “I think that will serve the public good” and others would not. Do you think there should be some kind of authority/control over how it should be used? If so, what?

The Chairman: For example, a stewardship council or something such as that.

Dr Sandra Wachter: The idea of stewardship is very good because if we have one central body that has the authority, expertise and resources to manage all that, we can ensure that those rights are being protected. None the less, individual rights have to be valued. The whole data protection concept is very much based on consent. Consent means, “I can give you my data if I believe you will handle it ethically for research, for example, or for any other applications, but I shall always retain the right to withdraw consent if I am not comfortable with it any more”. Having something that protects my rights is good, but I still think that control for the individual should be maintained.

Lord Levene of Portsoken: Of course, once a piece of information has been released, shall we say irresponsibly, you are lost.

Dr Sandra Wachter: Yes. I guess that is the problem that we have in the digital age anyway, that digital information cannot be deleted. This is why ethical foresight is a more important step forward because, if I am being informed about the possible risks, I will take this into consideration if I want to take that risk. We propose to make the data protection impact assessments publicly available. It is not very clear in the framework whether this is mandatory or not. The Article 29 working party has published guidelines recommending that those impact assessments should be made public, and I very much agree with that, because it would give the individual an easy way to understand what is going to happen with the data and then take a calculated risk, because you cannot always protect privacy; of course, sometimes systems will fail.

Dr Mercedes Bunz: I would come back to a more collective view of datasets. If I ask a question of how data can be managed for the public good, one problem that we have at the moment is that we do not talk about data in a nuanced enough way. We try to come up with one solution when we use the term data, but there is no one solution. It needs to be looked at in sections. Health data is of a different value and should be processed in a different way from the transport data that my Oyster card leaves behind. If we would like to give somebody access to this data, there are different procedures. One sort of data might be interesting for the highest bidder. Health data, of course, should go to the most interesting research and not necessarily to the highest bidder. They are different interests and the context in which the data is produced is different, and it is really important to make sure that it is used in the best

way for the public good. We need to start thinking about it in a more complex way.

Lord Swinfen: Can you tell us what the public good is?

Elizabeth Denham: There is jurisprudence on what the public interest is that comes from freedom of information legislation. We can start by looking at some of the jurisprudence that exists for defining the public interest involved in the decision to release data or to withhold it. In that context, there is some jurisprudence. I agree that if we are going to talk about public interest and the use of data in the public interest, we are talking more about the public sector. If you turn to the private sector, arguably the use of that data is always going to be focused on the commercial good of the company balanced against the reasonable expectation and rights of the individual. These are slightly different contexts. The GDPR gives us a huge step forward in requiring companies and public bodies to think and to focus on what they are doing with machines, machine learning and artificial intelligence, and to consider the rights of individuals, document that and stand ready to account for the decisions that they have made. The Information Commissioner has the ability to look at those decisions. Individuals have the right to challenge those decisions. We have taken a couple of giant steps forward.

If we start talking about how you make that judgment in a commercial entity, some companies have implemented ethics groups within the company with outside experts to be able to help them make these calls. This new data ethics body that we are talking about could help to certify ethics groups within companies to make these judgments. I could see a framework where, "This is what 'good' looks like", when it comes to DPIAs and ethics reviews within the private sector. That is a different concept from the public sector and looking for the good of society in the use of data.

Baroness Bakewell: You speak of jurisprudence, and the Bill that is going through applies to Great Britain, but AI and data is global, so how does this match with what is going on in places such as China and the East? How do these different systems of control operate? Do they have a chance? Do they overlap? Do we have international agreement?

The Chairman: Information Commissioner, I think that is one for you.

Elizabeth Denham: I think that is one for me. That is a really challenging question because, you are right, data knows no borders, data is travelling around the world and yet our data protection laws are mainly national. I see a convergence in the western world, with the exception of some large countries, of standards for data protection; and arguably, the GDPR from the European Union and the UK's Data Protection Bill are coming to the highest standard. That said, Australia is going through a reform of its data protection laws. Various US states have upped their game when it comes to the standards in law. I am not saying the country as a whole has, but individual states have. Canada, Japan, South Korea—all these countries are upping the standards, and I see some convergence. That said, the long-term solution is probably not one region of the world making judgments on another region of the world to enable

data flows, but, rather, an international treaty so that we can respect and protect data flows.

Baroness Bakewell: How far off is that?

Elizabeth Denham: I think it is years away, but we should work on it. The UK is in a very good position to advocate for this because of the role that we play between Europe, the US and the rest of the world.

Lord Hollick: When the Royal Free NHS Trust decided to release patient data to DeepMind, did they consult with you?

Elizabeth Denham: No, they did not. We were disappointed that we were not consulted. In the new world of GDPR and the Data Protection Act, there would be a requirement to do a data protection impact assessment with that kind of arrangement, but in the world we are in right now, with the Data Protection Act 1998, there is no legal requirement that required them to consult with our office. If they had, we would have given them advice on how to construct that arrangement to comply with the law and would have asked about their decision not to anonymise 1.6 million patient records. We would have looked at the contract terms and asked for a privacy impact assessment, which did not exist. There were some gaps in that arrangement, as has been made publicly available.

The good news in the undertaking that we have required the Royal Free to make with Google DeepMind is that it sent a message to the rest of the trusts and the rest of the country on how to construct these kinds of arrangements so they comply with the law.

The other piece that I heard much dialogue about is how you make this data available to other researchers, companies and groups that could take advantage of the solutions in the machine learning and algorithms.

Lord Hollick: Picking up on that last point, which addresses the earlier discussion about monopoly behaviour, would you envisage that you should have the power—I do not know if you have the power—to require the information to be made generally available to research organisations and other public corporations, assuming that they had the right skill set to do it? As a corollary to that, would you see the role as being to decide whether the exploitation is indeed in the public interest? Secondly, is the benefit accruing to the public, either in monetary terms or general information about a particular sector and how to deal with a particular disease, adequate to justify the release of that data?

Elizabeth Denham: That was three questions.

Lord Hollick: Sorry.

The Chairman: I had noticed that, too.

Elizabeth Denham: If I leave one aside, you can remind me. The first question is whether or not my office has the power, or should have the power, to order data to be made publicly accessible in an anonymised way; we are talking about health data. We do not have the power to order data to be published, but the Government have a framework for encouraging and committing to open data. The Government already have the regulatory framework that encourages more open datasets to be

available for research.¹ The framework is there. Maybe the impetus to get the data out there is not, but I also think that anonymisation is a really important tool and should be taken up by many more organisations. Google DeepMind is a good example of people saying, “Wait a second here. What else could we have done and which other researchers could have benefited from accessing that kind of rich dataset? How can we build this differently?” That is another good news story about the lessons learned in Google DeepMind. You will have to remind me of your third question.

Lord Hollick: Do you make a judgment about whether the public benefit is adequate for the release of data to an organisation, which is presumably going to use that data commercially?

Elizabeth Denham: To the extent that the ICO makes decisions on the fairness of the collection and use of data, we consider the public interest. But I also think it is a broader ethical question that the UK, especially with a publicly paid healthcare system, should ask itself and engage with the public on the proper exploitation of the rich data stores that we have in this country in healthcare and in delivering healthcare.

The Chairman: Thank you very much indeed. That is a very interesting conclusion to that. Lord St John.

Q61 **Lord St John of Bletso:** What technical approaches are there to preserve privacy while also ensuring the benefits of AI are realised? Could I ask specifically for your views on how blockchain is a potential solution?

Dr Mercedes Bunz: Differentiated privacy is definitely a method that has recently been very promising. We know that we can never fully anonymise data, but we can get there in a way that we can publish it. Methods have already been developed. It is really important to invest in research in this area if we want datasets to be released and if we have an urgency to make them publicly available for SMEs and for research. That is an important point to make.

Elizabeth Denham: Anonymisation is a really important technical tool. I get frustrated when people say there is no such thing as perfect anonymisation; there is anonymisation that is sufficient. A provision in the Data Protection Bill on sanctions for intentional re-identification was set up to encourage anonymisation and deter intentional re-identification. I would like to see more anonymised datasets, but the other technical tools that we have already talked about—data protection impact assessments—are really important here to protect that data. I am afraid I am not an expert on blockchain.

Dr Sandra Wachter: I would agree with everything that has been said. It is very important to understand that privacy is not a binary thing; it is moving, and you cannot always guarantee privacy, but you can get there. You can do the best you can. It is very important that we have realistic expectations when we talk about data privacy. Anonymisation is obviously something that is very good and should be encouraged, although we have to keep in mind that technology is evolving and there

¹ Note by witness: data.gov.uk/about

will be techniques to reverse engineer practically everything that we develop today. Ideas such as differential privacy are very good because they show that systems can be built with a different mindset, because differential privacy allows the querying of a data base rather than granting access to the actual data. This is a different technical approach, a privacy-preserving approach, and such tools should be preserved, and funding should be made available for that.

Encryption is a very important tool that we need to focus on, but there is a lot of tension currently as to whether encryption should have mandated backdoor solutions, for example. This should be discussed because if we make encryption intentionally weaker, we cannot preserve privacy in general. It is not only a technical question but a question of ethical data handling and ethical sharing practices: making sure that if I give my dataset to someone else, that person is going to use it in accordance with data protection law. It is not just with one company; it is with everybody who shares that data.

Lord St John of Bletso: I appreciate I have thrown in a bit of a curve-ball by asking about blockchain, but blockchain has had such a revolutionary impact in providing transparency and accountability. One of the reasons why I mention the blockchain solution is that it can specify that data will be available for a set time.

Elizabeth Denham: And it is all audited.

Lord St John of Bletso: There is a complete audit trail. That is one of the potential solutions.

Elizabeth Denham: It is decentralised and audited.

The Chairman: Thank you very much. We are going to have to keep moving. Lord Holmes.

Q62 **Lord Holmes of Richmond:** Should the use of black-box algorithms, such as deep learning, be restricted in domains of high-risk decision-making—military, health, finance, law? If so, why?

Dr Sandra Wachter: It is a very sensible approach to think about AI not just in general but as context specific. It makes a lot of sense to have different kinds of approaches to sensitive areas such as finance, health and military. This is why I think fairness, accountability and transparency as general principles are very helpful for all AI applications, but we need to think about what they mean on the ground in specific applications. In my opinion, making things more explainable and granting the right to explanation could be a very good step forward in making sure that we understand what is going on inside the black box. If that is not possible, we need to have different solutions. We can think about auditing mechanisms. If we come to the conclusion that a system is inherently opaque and not understandable, we need to have other safeguards in place. This could be auditing after the fact, auditing during data processing or inbuilt processes that could detect biases, for example. We could also think about certification schemes where you certify your algorithm before you deploy it. In high-risk areas, yes, of course, we should have safeguards in place that ensure transparency, fairness and accountability.

The Chairman: Information Commissioner, do you feel an amendment to the Data Protection Bill coming on?

Elizabeth Denham: Yes. I would agree that there is no black box that should be completely free from some kind of oversight and accountability mechanism. Whoever is going to look at military black boxes, it needs to be done. I agree that you need to look at this in sectors and at what kinds of exemptions are in place for intelligence services, financial services, et cetera.

Dr Mercedes Bunz: The interesting thing is that the technical development at the moment goes in the direction of creating more explainable AI. The Pentagon and DARPA are developing that, for example. There is also a young computer scientist called Been Kim who developed a method of looking into huge datasets of millions of data, in order to understand what an algorithm sees as a prototype and what it understands as the minority of this prototype that will not be seen as following the rule. There is technical development. We can see that there is a law—GDPR—and there is technical development following it and catching up. It is not that we need to accept that artificial intelligence works with a black box and has to work as a black box. In a few years from now we will not accept it working as a black box so easily. I am sure we can all remember saying that artificial intelligence was not prejudiced, and then we found out that it was trained on the wrong datasets.

Lord Levene of Portsoken: So far we have been talking about data being used in the public interest and, hopefully, there is a general understanding of what the public interest might be. If we are looking at military applications, you have two sides which are clearly against each other. If you are going to use the data there, one side is going to say, “Yes, that was very good use of data”, and the other side is going to say it was catastrophic. How do you deal with that?

Dr Sandra Wachter: That is one of the most complicated questions.

Lord Levene of Portsoken: I dare say.

Dr Sandra Wachter: It is very hard to answer.

The Chairman: Our Ministers normally say they will write in response to that sort of question.

Dr Sandra Wachter: Military is one of the most sensitive areas for the use of AI. We should consider whether we really want unsupervised, solely automated systems in the military. There have been a lot of calls recently to ban completely autonomous weapons. It is very fair to say that AI might not be the best solution in all sectors. When it comes to a decision about taking someone else’s life, maybe AI should not be the last step. Maybe we should always have a human in the loop to make the final decision.

Lord Levene of Portsoken: In that case, who is going to stop it being used?

Dr Sandra Wachter: It is always the case that some countries will have different opinions on how warfare works. That does not mean that we have to have the same standards. There will always be controversy with

that, but the more countries get on board with that, the more likely it is we are going to have a framework in place—an international treaty, for example—whereby we say, “These are the principles that we value”; something like the United Nations has.

The Chairman: These have not been quite added to your duties yet, Commissioner, have they?

Elizabeth Denham: Not autonomous weapons, but if you are asking a question about the use of data for national intelligence or defence, I still think the trust of the public is essential, and appropriate checks and balances and transparency are important to take the public with you. The public in the UK do not expect that the use of their data will be entirely taken on trust. There needs to be checks and balances and oversight even with that, but I am not in the place to make decisions about autonomous weapons.

Q63 **Baroness Greener:** We touched on prejudice in the last question and I would like to ask you about that very issue. What can we possibly do to mitigate against unintended prejudices, which are already seeping into this area? Even though it is a new area, it is the same old prejudices. I appreciate, Dr Bunz, that you write that government-led guidelines might be a solution, but I put the question: are guidelines enough? If they are not, what are the standards and the mechanisms? How can we overcome this? I speak with some impatience, which I share with one or two other members of this Committee, when we talk in the gender context alone, and there are others of course.

Dr Mercedes Bunz: In certain areas, such as policing, this needs to be tested and contestable. One good answer for me is that as long as the Government develops or have a hand in the development of datasets, they can do something. If they leave that to Google, Facebook and the big companies, this will not be the case. Developing, looking at and taking responsibility for datasets could mean asking the question why those companies did not find out that the app was already tagging people as animals. This definitely cannot happen after the app is released. The Government’s responsibility in having data and publishing data could be one way of looking into that and of making sure the data has no gender bias. I think that is important.

Baroness Greener: May I come back—and forgive me if I interrupt the other members of the panel? Let me push you on that. Do we trust government in this area, even though we know about equal pay and other things, and so far there has not been much progress? There must be more that can be done other than what you have just described to us. What is the really proactive thing, because it sounds to me like quite a passive approach?

Dr Mercedes Bunz: Are you asking me about the law and regulation direction, or are you interested in the technical direction?

Baroness Greener: Either. Give us a solution or guide us towards a solution in this area.

Dr Mercedes Bunz: As I said, technical methods are being developed as we speak to test datasets and to understand what datasets understand as

a prototype—how the classification works—and that is a really important step. Artificial intelligence will not come for free. Where we put money into research makes decisions and will develop things in certain areas further than in others. There is a really important decision to be made there. We should also have a certain amount of regulation. We do not allow other publications to be completely biased. We have certain standards and these standards should be kept with artificial intelligence services as well, naturally.

The Chairman: Information Commissioner, do you or Dr Wachter want to add to what has been said?

Dr Mercedes Bunz: May I say one last thing? A lot of times we have laws and frameworks in place but we do not enforce them enough, and I think this is one of those cases.

Elizabeth Denham: I would make a short addition to that. It is still people who are determining what data is going into the system. It is still people/humans who are writing the algorithms. It goes back to the decision-making by individuals. We have a law that requires fairness, and now a law that requires more accountability and provability and evidence of fairness.

Dr Sandra Wachter: I would propose something a bit different, I guess. AI forces us to look into a mirror. AI shows us the world as it is; the biased world that we live in. That is something that we have to deal with. The problem is not the technology but the world that we live in. With AI we now have the opportunity to make better decisions than human decisions. To do that, we have to remove the biases. There are a couple of things that you can do. Of course, there are the things that everybody talks about. We have to have a diverse coding community. If we have only white men coding those systems, there will be a gender bias built in. You can do that and also educate coders so they have a better understanding of their ethical responsibilities. It has not been mentioned that it is very important to have other disciplines working on this. That is why I stress so much that the social sciences should have a say in that. Social scientists have been researching the underlying causal models of our world. This is what they do. They run experiments and try to explain the world as it is. They can detect biases. They can give us a reality of the world.

Let me give you an example. If a decision about a loan is based on my employment history or on my salary, we would say this is fair because we can measure the probability of defaulting on a loan but, because of the social scientists, we know that salary and employment histories are gender biased. Women make less money than men, on average. Often, they have less work experience because they take care of their children. If you do not know that, you are going to discriminate against people without knowing it. You need to have a causal model of the world to understand the complex socioeconomic structure of our world, so you can make sure you do not make the same mistakes with AI that we did in the past with humans.

The Chairman: Thank you very much. That was terrific. I am going to ask Lord Hollick to ask the final snappy question.

Q64 **Lord Hollick:** What, in your view, is the single most important recommendation that we should be making?

Dr Mercedes Bunz: I would come back to the creation of datasets. The Lord Bishop asked whether there was some urgency here, and I would say there is. We see those datasets at the moment only with the five big companies that have taken the lead in AI. If we want anything to happen, whether in research or in SMEs and business, we need to create procedures and frameworks whereby that data can be accessed and worked with.

Elizabeth Denham: The trust and confidence of the public in new systems that use artificial intelligence is the key. If you do not have trust and confidence and people are spooked by new technologies, you will not have take-up and people will not play. You will not take the public with you. Data protection and data ethics have to be built into the process and baked into the process all the way along, and not bolted on at the end of the day.

Dr Sandra Wachter: If I could make one recommendation, I would say to support a couple of the amendments that have been made to the Data Protection Bill. I am very pleased to see that there is now a proposal to make the right to explanation legally binding. That would be an exceptional way forward: to increase the trust of people who are concerned about those decisions, and getting people on board so they can challenge the decisions if they think they are being discriminated against. It is also going to increase accountability and fairness. I will have to lay out the reasons why I made a decision, I can be challenged about it, and I have to justify my actions. I would support that.

The Chairman: Thank you very much indeed, especially for your forbearance in that final answer, because some of our witnesses made three recommendations. Thank you very much indeed for your evidence today; we have really appreciated it.

Eileen Burbidge MBE, MMC Ventures and Project Juno – Oral evidence (QQ 46–54)

Eileen Burbidge MBE, MMC Ventures and Project Juno – Oral evidence (QQ 46–54)

Evidence Session No. 6

Heard in Public

Questions 46–54

Tuesday 24 October 2017

[Watch the meeting](#)

Members present: Lord Clement-Jones (The Chairman); Baroness Bakewell; Baroness Grender; Lord Hollick; Lord Holmes of Richmond; Lord Levene of Portsoken; The Lord Bishop of Oxford; Viscount Ridley; Baroness Rock; Lord St John of Bletso; Lord Swinfen.

Examination of witnesses

Eileen Burbidge MBE, David Kelnar, Libby Kinsey.

Q46 **The Chairman:** A very warm welcome to Eileen Burbidge MBE, of Passion Capital, David Kelnar, of MMC Ventures and Libby Kinsey, of Project Juno, AI. Thank you very much indeed for coming today. The session is open to the public. A webcast of the session goes out live and is subsequently accessible via the parliamentary website. A verbatim transcript will be taken of your evidence, and this will be put on the parliamentary website. A few days after this evidence session you will be sent a copy of the transcript to check for accuracy, and we would be grateful if you could advise us of any corrections as quickly as possible. If, after this session, you wish to clarify or amplify any points made during your evidence, or you have any additional points to make, you are very welcome to submit supplementary written evidence to us.

In view of the fact that we have votes today, if there is a Division in the Chamber while we are sitting, which I am afraid is likely, this Committee will adjourn as soon as the Division Bells have rung and resume after 10 minutes. It is a very heavy vote today, so 10 minutes might be conservative. We may need 15 minutes. We all know when there is a vote; there will be a great hullabaloo and ringing, et cetera. Perhaps you would like to introduce yourselves for the record and then we will begin with questions.

Libby Kinsey: Thank you very much for the opportunity to speak today. My name is Libby Kinsey. I spent 10 years doing technology investing for NESTA and then I went back to university to study machine learning at UCL. In fact, Dr David Barber, who was here earlier, was one of my teachers. For the last three years I have been a machine learning consultant and am co-founder of an entity called Project Juno. We exist to promote European machine intelligence start-ups.

David Kelnar: Good afternoon. My name is David Kelnar. I am an investment director and head of research at MMC Ventures. We are a UK-based venture capital firm investing in early-stage technology companies, particularly those that leverage software and data science. I lead a team at MMC called the Insights team, and that is a research group. My goal is to identify emerging areas of value creation, including technologies such as AI, to understand those technologies deeply and help us identify and invest in the very best early-stage companies in those spaces.

Eileen Burbidge: Eileen Burbidge. Despite the accent, I am now a British citizen. I am very interested in investing in this area for the UK. I am a partner at Passion Capital, which is not too dissimilar to MMC. We are an early-stage venture fund, so we invest in the very earliest rounds of start-ups which want to be the internet structures of the future, including AI but not exclusive to AI. I am also chair of Tech City UK, which is the DCMS-funded organisation to support and accelerate the digital economy across all of Great Britain. I am also a Treasury special envoy for fintech, as the technology enables and accelerates in financial services. Finally, as a point of disclosure, I also serve on DeepMind Health's independent reviewers' panel.

The Chairman: Thank you very much indeed.

The Committee suspended for a Division in the House

Q47 **The Chairman:** I am going to ask a very broad question to all three of you. What, in your opinion, are the biggest opportunities and risks in the UK over the coming decade in relation to the developments and use of AI? It would be helpful if you could put that in the context of Dame Wendy Hall's review and any other developments in that area you might care to mention.

Eileen Burbidge: It will be hard to follow. There will be lots of facets brought up which have already been brought up in the previous sessions. The one I would focus on is an answer to both sides—opportunities and challenges—that relates to skills and jobs, which I am conveniently lumping together as one topic but is, clearly, multifaceted. The greatest opportunity is for the UK to continue to lead in skills leadership in this space and sector, where it clearly has a lead now, but a real challenge is maintaining that skills leadership and, also, trying to minimise the skills divide that would be highlighted or exacerbated by what will happen in the sector and outside the sector as a result of the implementation of artificial intelligence, where jobs might be lost or replaced by incorporating the output of AI.

The Chairman: You are seeing the opportunities but you are very well aware of some of those downsides as well.

Eileen Burbidge: Perhaps to put it a little more succinctly, I think the opportunity is for the UK to continue to lead in this space. The challenge and the risk, though, is that, by continuing to lead in this space, we increase the divide between artificial intelligence leadership, subject matter expertise and research leadership with what we might be leaving behind in terms of not enough reskilling and not enough highly-skilled talent to come in and fill the pipeline. Perhaps, as a result of AI being

incorporated in automation or other types of services and industries, we might be leaving displaced jobs or other workers that are not in our artificial intelligence industry.

The Chairman: Do we need more of a route map than we have at the moment?

Eileen Burbidge: I do not know that one is possible but I think anticipation and trying to look at domestic skills development—retraining for jobs that might be most vulnerable to automation—would be important. Continuing to work with the higher education sector and attracting highly-skilled talent from outside the UK is very critical.

David Kelnar: The opportunities presented by AI will be profound. The greatest opportunities, though, are threefold. I think of these as benefits for society at large, benefits to the economy and new commercial opportunities. I will describe those briefly and talk about some risks that are the most significant. Regarding benefits to society, AI will have the greatest impact in datacentric sectors. Those include finance, retail, transport, manufacturing and healthcare. AI will offer broader access to better and cheaper healthcare; improved transport and mobility with fewer accidents; improved management of financial assets and risk, increased manufacturing capability and agricultural output; and more efficient and satisfying retail experiences.

Secondly, AI will offer new sources of economic growth, taken as a whole. PwC has estimated that UK GDP will be up to more than 10 percentage points higher by 2030 as a result of AI, the equivalent of an additional £232 billion. The impact over this period will come from product enhancements—better, cheaper, more personalised products that stimulate demand—and increased productivity through augmentation or automation of some roles in the labour force. Finally, some of the most exciting opportunities around AI over the next decade will be shifts in the commercial landscape. Sector value chains will shift. Some existing business models will obviate and be replaced by new ones, and companies' competitive positioning will change. These commercial changes create opportunity for early-stage companies in the UK today—some of the kinds of companies that we back—either to enable or disrupt incumbents and create significant new UK winners.

The Chairman: Thank you. I have a question, in that context, which I did not put to Eileen. Does the Dame Wendy Hall review have any relevance to your activities and estimate of what the opportunities are, and so on?

David Kelnar: It does. Let me put it this way: a number of the areas on which Professor Dame Wendy Hall focused are extremely well guided towards some of the greatest risks and challenges that we see. Two that stood out for me are increasing ease of access to data that she highlighted as a priority in programmes such as Data Trust and others that could be solutions, but also development of greater AI skills, so industry-funded master's programmes for AI, as well as some of the other steps she suggested. Without going into the detail, the recognition of the challenges around data access, talent and more start to address some of the significant challenges.

Libby Kinsey: As David has already outlined, the opportunities are myriad. Some of those opportunities are more important than others. There is the potential that we could transform how government, healthcare, education and care are delivered. That would be important for society as well as for venture capital returns. The big risk I worry about is that we will squander the advantages we have here. The UK is so well positioned. It has a breadth and depth of expertise across a range of AI technologies in both academia and industry. It has markets on our doorstep; it has investors and, at the moment, it has talent. London is the most attractive place for entrepreneurs to come in the UK. To speak to Eileen's point again, talent will be one of the major constraints going forward. There is already some concern when start-ups are hiring about uncertainty as we go through Brexit. There will always be uncertainties, I guess, around that sort of thing. We also need some leadership. It is great to have all these ingredients but we need to ensure we are targeting support where it is required. That comes, again, to David's points around needing to help start-ups through accessing public data and ensuring that the value in public data stays in the UK. We can help with channelling public resources, such as computational power, towards start-ups.

The Chairman: Digressing slightly, did you hear what Marko Balabanovic had to say?

Libby Kinsey: Yes.

The Chairman: That ties in very well.

Libby Kinsey: I agree wholeheartedly.

The Chairman: Thank you. Does Dame Wendy's review help with your worries about the "squandering" of our advantages?

Libby Kinsey: Yes, it does, for the same reasons as David has pointed out. Data is going to be very important, and some of her recommendations were good. Talent and having industry supporting master's students will be very important. Did she mention finance? Did she mention the UK being an investor?

The Chairman: I do not think there was a huge amount in the report on finance, if I recall, I must say. There were many areas that, for us, were not covered. Eileen, I wanted to make sure that, wearing your Tech City hat, if you had comments on the review, you were able to make them.

Eileen Burbidge: I thought that the review was a really strong synopsis, to be honest. As is evidenced by the fact that you are having 14 sessions, this is a big, hairy, ever-evolving sector and there are a lot of different considerations. Professor Dame Wendy did a magnificent job of trying to synthesise all those and tried to identify priorities for policymakers, which I felt was her remit. I thought that was quite important. On your question about whether or not it might have influenced our perspectives, from either a Tech City point of view or as a private investor, because we are constantly looking ahead—I would like to think our horizon is probably further ahead than policymakers—I do not think there was anything new, unsettling or surprising in the report. It was more a case of validation in

terms of “These are the right things to emphasise to policymakers today”, and confidence therefore.

The Chairman: So it was a consolidating document, in some respects.

Eileen Burbidge: Exactly, and affirming and validating.

The Chairman: I hope we will stir it up slightly more, in that case, when we make our recommendations.

Lord St John of Bletso: A very brief question. You each spoke about the need for more skilled AI staff and the challenge of maintaining the UK’s leadership in the AI space. On a political angle, to what degree is Brexit uncertainty a threat to keeping those skilled staff and bringing in new talent?

Eileen Burbidge: Personally, and certainly with my Passion Capital hat on—I probably cannot wear my Tech City hat or my Treasury hat when saying this—I think it is a massive risk. It is a tremendous threat. I am not making political statements about whether or not I support Brexit; my point is more about this transition period and the uncertainty before there are clear guidelines about the nature of labour movement between the European Union and the UK. That uncertainty alone raises issues and gives people pause before they consider coming to the UK, or challenges for companies trying to recruit outside the UK. Unfortunately, the sentiment or the noise amplified through the media is unhelpful when there are either radical groups or minority groups saying certain things about why they may or may not support Brexit and how welcoming or unwelcoming the UK might look as a consequence, as a perception. I do not think that is necessarily accurate but certain groups or voices are louder than others, and it might look as if the UK is not as welcoming of diverse talents as it might have been.

Lord Swinfen: This may have been partly answered. What is being done to produce enough young, home-grown talent to supply the market?

Eileen Burbidge: That is definitely a question I would hope your report will ask or amplify. That is a question we ask all the time. From a Tech City point of view, the audience or the constituents we serve are individuals who are largely outside of higher education and probably already in the workforce. We try to offer digital skills to make people more digitally literate and to convince them to be digital entrepreneurs, but not specific to artificial intelligence. From a private investor point of view, we see things such as introducing computing curriculum into the youngest school ages as a positive step but one that will take years and years to bear fruit. The best poll measures will be role models, examples, more commercial successes—whether they are acquisitions of DeepMind, Magic Pony or Vocal IQ—but local UK businesses, whether they are to foreign acquirers or not, as types of ambitions that younger entrepreneurs or younger entrepreneurs-to-be will want to mimic. That will try to pull more people into the sector. As Libby also said, we have leadership now in skills. Cambridge is widely regarded as the centre of the universe in artificial intelligence. The amount of investment coming from the Government but, also, commercial, private sector companies, from Apple to Amazon, in setting up research centres in Cambridge

because of that is fantastic investment and will help to bear fruit. We need to be doing more across the board.

The Chairman: Thank you. That confirms our decision to visit Cambridge, very shortly.

Eileen Burbidge: You absolutely must do that.

Q48 **Baroness Rock:** You have touched on quite a lot of this already, but I want to focus a little on start-ups. We have talked about access to talent and skills, and access to data. We have also heard about the importance of supporting start-ups. There are three points. What do you think is the status of start-ups in the UK at this precise moment? Obviously, you are investing in quite a lot of start-ups. What are the challenges beyond what we have already touched on? Perhaps finance might fall into that. All three of you have talked about how we are leading the field at the moment, but in the future how difficult is it going to be to compete against the larger businesses in America, China and elsewhere? David, perhaps you could start.

David Kelnar: Certainly. Would it be helpful if I gave you a quick overview of the state of AI start-ups in the UK and then touch on the three greatest challenges we see in entrepreneurs when we speak with them? There are over 400 early-stage AI software companies; that is, software companies that have AI truly at the heart of their value proposition. AI entrepreneurship in the UK is thriving. The number of AI start-ups founded annually in the UK has doubled since 2014, and since then, on average, a new AI start-up has been founded every five days in the UK.

We are also entering a second wave of AI entrepreneurship. The first wave related to AI research or the development of core AI technologies, particularly in the area of computer vision or language with cross-domain application. Today, while a lot of that activity continues, we are firmly into the second wave, the wave of applications. Over 80 per cent of those 400 start-ups today are business-to-business suppliers that are providing solutions to address a specific problem in a sector, or a given business function. Just 10 per cent of UK AI start-ups today are developing core technologies—wave one technologies—applicable to a broad range of areas, and 10 per cent sell directly to consumers.

Entrepreneurial activity is also very unevenly spread. For companies that focus on a sector there is a lot of focus on the finance sector and providing solutions to finance. Activity is also extensive in healthcare and retail. For companies that focus on business functions, more UK AI start-ups than any other—about one in seven—address marketing and advertising business function. Activity is also extensive in general IT and business intelligence applications. Entrepreneurs are focusing on these areas for sensible reasons: they all offer numerous prediction and optimisation challenges well suited to the application of AI; all have extensive and relatively accessible datasets to enable the training of algorithms; and, finally, they offer measurable results, which is pretty important for entrepreneurs trying to build businesses. In select sectors, such as manufacturing, and in some business functions, such as the

finance department, interestingly, activity appears a little more modest, relative to, perhaps, market opportunities.

Finally, a brief note to put this activity in context relative to Europe. UK AI companies comprise about half the European total. AI is well-represented in the UK with a slightly higher proportion of start-ups in the UK focused on AI than in Europe, or even in the US. This sector is nascent, though. Two-thirds of UK AI start-ups are in the very earliest stages of their journey, with seed or angel funding. Again, interestingly, the sector is maturing quite rapidly. UK AI start-ups are a little less embryonic than their European counterparts, offering a little bit of advantage in competitive procurements.

I will touch on the three biggest challenges that we hear from AI start-ups. Those are about shortages of AI talent, as Eileen and Libby have described; the limited access to training data; and, finally, the difficulty of productising AI. Regarding shortages of talent, “Access to talent and its competitiveness is the biggest challenge”. That is from David Benigson, the CEO of an early-stage AI company called Signal Media. Similarly, “The number one challenge when developing AI is recruiting the best human brains”. That is from Fabio Kuhn, CEO of another early-stage company called Vortexa. This challenge is unlikely to lessen in the short term given limited supply of AI talent and fierce competition, not just from start-ups but from global technology companies, from banks and others.

Secondly, access to training data is a classic chicken-and-egg problem. Early customers and their data are hard to acquire for these companies without existing reference clients. Again, this problem is likely to persist in the short to medium term, but leading entrepreneurs are trying to mitigate this challenge by developing data access strategies from early in the lives of their companies. Finally, “Taking what works well in the lab and getting it to work in a diverse population is a big challenge”. That is from Chris McCann, CEO of snap40. Again, this problem is unlikely to ease materially in the short term but may lessen somewhat over time as AI technology continues to improve and data becomes more widely available.

The Chairman: Is that third one what you call productising?

David Kelnar: Productising AI.

The Chairman: I am learning new language every day. Thank you.

Baroness Rock: Eileen, I do not know whether you have anything to add?

Eileen Burbidge: No, I think that is very comprehensive.

Libby Kinsey: I would give a positive, which is that Eileen has already mentioned the number of start-ups that have sold. What is interesting, to me, at least, is that in the past when companies sold at an early stage, often those teams would go to the acquiring country, but in these cases DeepMind, Vocal IQ and Magic Pony have stayed in the UK. That means we now have teams productising in a corporate world, and they are taking that expertise outside. We are seeing engineers from those teams

starting up their own companies. There is a virtuous circle that we have not seen so often before.

David Kelnar: Could I add one thing?

The Chairman: I am going to bring in Lord Swinfen and you can come back as part of the supplementary.

Lord Swinfen: Can start-ups get the funds they need to start relatively easily?

Eileen Burbidge: The good ones always can. As Libby was saying, with the virtuous circle, now it is happening with the AI community. The best of the best know each other, so the successful companies that have already exited or done well or are most prominent in their field, their colleagues and peers will be known to them and they provide references for others. There is plenty of investment appetite for the sector. Yes, is the short answer.

David Kelnar: To give you an example, about two per cent of all start-ups in the UK are what we might classify as AI start-ups. Over the last 12 months at MMC about 53 per cent of all the capital we have invested has been into AI first companies. There is capital available for some of the best AI companies.

The Chairman: Do you want to come back on the earlier point? I will add my question: you are saying there is capital for start-ups there. Does that apply all the way up the Sherry Coutu gap, so to speak? Has that now been cured, effectively?

Eileen Burbidge: In this sector there is not a scale-up gap at all. There is a huge amount of appetite. Given how instrumental and transformational the sector can be for every single industry, the experiences we have had for companies in our portfolio, and it would seem from Tech City, is that there is no shortage of capital at every stage of a life cycle.

David Kelnar: I wanted to give a real-life case study to illustrate the reality of what Libby and Eileen were describing in terms of this virtuous circle. There was an acquisition of an early-stage AI company by a large company headquartered overseas in the last 36 months. As a direct result of that transaction, we witnessed new investments into our venture fund, new investments into early-stage AI companies, the creation of new start-ups by founders who had achieved that successful exit, and managers who went on to form senior leadership teams or serve as advisers at other early-stage companies. I would not underestimate the importance of that virtuous circle. We are seeing the impact is real and as that continues over the years ahead, acquisitions, whether by companies overseas or in the UK, are very valuable.

The Chairman: That takes us, very neatly, into Lord Levene's question.

Q49 **Lord Levene of Portsoken:** Libby, perhaps we can start with you. You have been telling us how the UK has got itself into a relatively strong position, vis-à-vis its size with the rest of the world. Do you think more should be done to prevent the acquisition of AI start-ups in the UK by larger, foreign corporations? If so, what?

Libby Kinsey: I have been looking forward to this question. We should be thinking about this because AI is strategically important for nations. The way I would look at it is to try to understand why start-ups are selling earlier than, perhaps, they might, before they IPO or build businesses. There are two reasons. I am going to be a bit hand-wavy here. European investors tend to have smaller funds and less appetite for risk than US investors. They may be younger, they may have less of a track record, so they are looking for a liquidity event and looking to exit. The same is true of the founders. Many of them will be earning less than they might in the private sector, and so that crystallising event when they exit is compelling. The other reason is that we have talked about finance being available for start-ups but by the time you need to scale you might be spending a great deal of money on talent and computational power, and you might need access to very large datasets. All of those things are very compelling reasons why you might go to another acquirer that has those things. If we want to keep these companies here we have to look at all of those aspects and make it easier for start-ups to access the data they need, the talent they require and the computational resource. We need to think about patient capital, and significant funds of patient capital.

Lord Levene of Portsoken: Eileen, to what extent are funds available in this country with those who are prepared to put money in without having to look offshore for them? I am not necessarily talking about taking the companies out but, at least, to take a share in them and provide them with the capital they need to keep moving forward?

Eileen Burbidge: As I was saying earlier, I think there is great appetite from UK-based investors for companies at all stages in artificial intelligence. Also, as has been touched on and implied in earlier remarks, foreign investors are interested in the UK's artificial intelligence ecosystem and not necessarily insisting that these companies move from the UK at the time they invest, to support further growth, or even after they are acquired. I do not disagree with much of what Libby said, but I think on the underlying thesis I might have a different viewpoint. As to whether the UK should be comfortable, for example, with 61 per cent of the national grid being owned by a Chinese company, I do not think artificial intelligence poses any greater risk, threat or concern than that very general meta-question. I do not think the sector should be protected any more than the rest of the UK infrastructure or what we might perceive to be our assets, whether it is energy, water, even Harrods, Hamleys or PizzaExpress. All of these have large foreign ownership stakes, the majority owned by foreign investors. I do not think artificial intelligence warrants any greater protections than other sectors. A lot of what Libby was saying is valid, but also speaks to investment, as you rightly asked about, not necessarily exits and acquisitions.

Lord Levene of Portsoken: David, do you agree with that? If the companies and the talent are going to stay here but the capital is coming from outside, presumably it does not really matter. Do you agree with that?

David Kelnar: Somewhat extreme examples of national security notwithstanding, I think efforts to stop or inhibit companies being

acquired by overseas companies would do far more harm than good. Acquisitions by large corporations, whether headquartered in the UK or overseas, are incredibly valuable for the UK technology ecosystem. They incentivise entrepreneurs, they attract investors, they create capital, as I have described, as recycled into new start-ups, and establish networks of talent and support that catalyse new companies. For those early-stage companies, as Libby described, often joining forces with the world's largest technology companies can provide additional technical and non-technical resources, access to global client bases, the opportunity to realise their global ambitions, and so on. I think we are solving for the wrong thing if we think about inhibiting that. However, for companies that wish to become independent, global leaders more can and should be done to support them. Greater access to growth capital in the UK, increased guidance during the internationalisation process would be of value.

The Chairman: I must bring in Viscount Ridley and then Lord Hollick.

Viscount Ridley: It was a footnote to that. Perhaps Lord Hollick is going to ask it, because I was going to ask Lord Hollick's question from the last session. There is foreign and foreign. Would we have been comfortable with a Russian firm buying DeepMind? I think most of us would have real qualms about that. How do you test the limits of that, in terms of security?

Eileen Burbidge: The last panel answered by saying the "what next" question. Will the commitment be to retain the talent here in the UK? Would there be continued investment in the UK to grow the team? That is critically important. Going back to what I was saying a few minutes ago—whether there is a concern about Russian or Chinese interests, or other interests, owning UK physical infrastructure—the risks are similar. I do not think there should be a different level of risk. I do not know if Libby or David have a different point of view.

Libby Kinsey: I do not have a better answer than that.

The Chairman: Do you want to come in at this point?

Lord Hollick: The previous panel laid quite a lot of emphasis on the need to have substantial UK-based, operated in the UK, businesses as another tier of funding. You have talked about the investment funding, and I think there is a difference of opinion as to whether there is enough growth capital here. Having substantial companies here is an important part of the mix. The substantial companies we are growing tend to be snapped up elsewhere, which is very nice for the investors, and one can understand why they would take it. When we met DeepMind—I do not think we are giving away any secrets—they explained they had real concerns about being owned overseas, if I can put it that way, but they entered into a number of arrangements and understandings which kept the business here, funded, and everything like that. Let us put a tick against that as good, but it is not a UK-owned business based here; it requires money and capital coming from outside. Eileen, you do not seem to share those concerns which were expressed by the previous panel.

Eileen Burbidge: I was not here for the entire panel, so for the part I heard I agreed with the whole "what next". To the extent that foreign

acquirers are going to be comfortable continuing to invest in a UK team and a UK operation, I think, net, it is a positive thing, which David articulated much more fluently than I did. In the case of DeepMind, for example—I mentioned as a disclosure point that I am on the independent reviewers' panel for DeepMind Health—it is very obvious, and the founders are very vocal about this, that if they were not British, because of the work they are doing with the healthcare system, given the lack of maturation, or commercialisation, at the NHS, they would abandon the UK market as a commercial market, and would go straight to the United States or, potentially, Singapore in order to do the work they are doing for the healthcare systems. I think we have benefited from their nationalistic pride in wanting to stay in the UK, invest in the UK and improve the NHS. Perhaps this is tangentially related to your point, but I think we need the public sector to invest more in enabling public sector departments and datasets to be more receptive and appropriate for artificial intelligence innovation and exploitation in a positive way.

Lord Hollick: Are you attracted to going down the route described by the previous witnesses that Canada is making efforts not only to bulk up and beef up the public sector but ensure that they have a number of companies based in Canada which can exploit the research they are doing?

Eileen Burbidge: Yes, and I do not think it is limited to Canada. France is being extremely aggressive; I think Singapore is hugely aggressive and so is Germany. There are many countries that are looking with envy at the current UK position and working very hard too.

The Chairman: Thank you. I make no apologies for having gone on, on this, because for you three witnesses this is a crucial question. I am going to ask Lord Levene if he wants to come back in and, also, you, David, if you had something to add because you wanted to add something in response to Viscount Ridley. We have probably moved on from that. Are you happy with that, Peter? Fine. We move on to Lord Holmes.

Q50 **Lord Holmes of Richmond:** Good afternoon. It might even be good evening now. Who the hell knows? I want to ask about collaborations. What barriers are there to collaborating with the higher education sector to turn AI research into innovative products? What can be done to turbocharge such collaborations?

Libby Kinsey: The incentives in UK universities around publishing and being cited are sometimes a little at odds with working with the commercial sector, which might need NDA agreements and not publishing. I think there is a requirement for a university that combines vocational education and fills the pipeline with master's students but also works on research and products that have an industrial focus and can result in spin-outs. That could be done within the existing universities, but I feel the incentive structure probably makes that quite hard.

David Kelnar: I would like, first, to acknowledge and thank David Grimm from Albion Ventures who shared with me his expertise on this issue. There are at least two challenges with commercialisation of technology developed by researchers at universities. The first is that a lot of

universities' venture investment programmes are structured still according to quite a traditional life sciences model. Universities, typically, seek quite substantial ownership stakes in spin-outs in return for assets, such as patents, the substantial support they offer and the expectation of significant dilution of ownership that will occur over time due to the spin-out's large capital requirements. In the era of AI, though, researchers' primary assets are more likely to be a little different—it is more a case of their expertise and capability rather than those existing assets. Their capital requirements, typically, are lower over time than life science companies. To better align AI founders and universities, universities can experiment with iterative models. Founders' choice, for example, is a pilot programme operated by Imperial Innovations at Imperial College, London. It enables founders to retain a greater share of founding equity if they require less support. That is one thing to look at.

The second challenge, briefly, is that researchers have limited access to commercially experienced management teams in AI, because this is a nascent area. There are not very many commercially experienced AI leaders to support a significant need. Efforts to broaden and strengthen AI talent networks, connecting successful entrepreneurs and operators in the field of AI with university researchers would be of value. Finally, closer collaboration, of course, between IP commercialisation companies and early-stage venture capital firms, such as ourselves, would be beneficial to broaden sources of support for nascent AI companies as they emerge.

The Chairman: That is the relationship with higher education, and so on. What about the relationship where it is publicly-funded AI research? What kind of return should taxpayers expect, or should they simply say, "This is pure research. We don't expect to get a return on it"? What about where public datasets are used? Is that an area that concerns you at all in making sure that those public assets are preserved or have returns on them?

David Kelnar: I do not know if any of us have a particular point of view on this. There is a balance to be struck. We need to recognise that the assets we have with regard to data in the NHS, with regard to the work coming out of universities, and so on, are of significant value. Forgive me, I am not sure which gentleman it was, but a member of your Committee in the last session, I think, gave the good example of the scan of a pancreas and identification of pancreatic cancer.

The Chairman: Viscount Ridley's pancreas will become legendary.

David Kelnar: Absolutely.

Lord Holmes of Richmond: It has already crossed the table this afternoon.

David Kelnar: That is a powerful example. We should recognise that while it is extremely important to drive financial value in the short to medium term from our assets, ultimately it makes at least as much benefit and difference to our daily lives to enjoy world-class services and products that have not been possible to date but which AI can enable. At a high level, I would encourage us to take a long-term view to recognise that we are ultimately all members of society and consumers of products

and services, not just or at all investors, and to recognise that benefits can be non-financial as well as financial.

The Chairman: Do Libby or Eileen have a contrary view?

Eileen Burbidge: I do not have a contrary view. The only thing I would say in addition to that, because you folded it in with your question, is that everything is relative. I would say on an encouraging note, because I am an investor so I am an optimist, that compared to material sciences as a sector, for example, I think we are doing better. It is not a very high bar. You asked about university collaboration. If we are talking about datasets and what the public sector can do, graphene is an unfortunate example of products that can be stalled or not realised, even though it was invented in universities here but not commercialised. That is a lesson learned. I think we are doing far better in artificial intelligence.

I completely agree with what David said on the university point, but the Cambridge innovation funds are doing better, even in the last two years, than they might have been doing previously, at being slightly more commercially minded. Research fellows at Cambridge, for example, are starting to be motivated to split time between the papers they are writing and the other incentives they used to have and work with small teams to spin out of the university.

On datasets, I would go back to what I said earlier on DeepMind. It is such an asset. We are fortunate that the founders negotiated so that they could stay in the UK and continue building in the UK. They wanted, first, to try to introduce pilots with the NHS. They have successfully done so, but they are finding that far more effort is needed in relation to data that is ready to work with than with other jurisdictions. That comes even before the question of how much we expose or share, and how much privacy we might encroach upon; it is literally whether or not the CIOs, CTOs and department heads in public sector organisations recognise how they have to store, harness or organise their data so that they can then choose whether to work with commercial enterprises.

The Chairman: That is a very interesting point. Generally, a good news theme is coming through here.

Eileen Burbidge: We are optimists.

The Chairman: We have better financing compared to some of the life sciences spin-outs and better commercialising compared to some of the life sciences spin-outs. It is interesting that we are building on some of that experience, perhaps.

Q51 **Lord Hollick:** Data is the feedstock of AI. Therefore, it has considerable value. Have you looked at investing in companies that partner with the National Health Service and with other institutions that are rich in data to help them exploit that data successfully? In the old analogue world, the model between publishers and universities was very much along those lines and proved to be very successful. Have you seen such investments? If not, why are you not out there seeking them?

Eileen Burbidge: We have looked at those kinds of investments. We call it the ingesting of data, the implicit data that might be available out there

and that people are not exploiting. I anticipate a question later in the session—it was asked of the last panel—on the regulation point. The current regulation in the framework for data is either too restrictive or simply too ambiguous for there to be the confidence to put a great sum of money behind a company that would be all about trying to expose or utilise data. Instead, at least from our point of view at Passion, we are investing in companies that might build the tools that work with the data, anticipating that at some point we will be clear about how we can use the data and what the policies on that will be.

The Chairman: We have kept you prisoner beyond your allotted time. Are you all happy to stay, say, for another quarter of an hour? We are very much enjoying the discussion.

Q52 **Lord Swinfen:** In your view, do investors have a duty to ensure that artificial intelligence is developed in an ethical and responsible way? If so, how should they do this? Should such development be regulated?

Eileen Burbidge: I thought this was an incredibly insightful question when I saw it on the papers for the session. The stark and objective answer, strictly speaking, is that I do not think investors have a duty to ensure ethical and moral behaviour. Most investors sign up to a code of conduct and are authorised persons by the FCA because they have fiduciary responsibilities as a first and foremost point. That is simply the objective current situation.

However, I also think that most investors today apply a moral and ethical compass to their investment activities. Most of them recognise, when investing in such nascent technologies and the potential that those technologies have on markets and sectors, that there will be ramifications from this and that responsibility needs to be applied to it. We are relying on human nature in that regard, which is why our business might not yet be robo-based investment; it is human-led, conscience-led. At our firm, for example, we believe that while there could be businesses that do very well and are commercial leaders in their space, if one is ethically and morally-minded or conscious their results could be even greater. We therefore feel that it is an additive and worthwhile consideration when making investments. To be quite clear, our obligation to our investors is to generate as strong a financial difference as possible. We would like to think that ethically-minded investments will have a stronger commercial upside in results, but that is a hypothesis on our part.

Libby Kinsey: The only way to require investors to invest in an ethical and responsible way is for their limited partners, for their investors, to make that requirement and prefer management teams that have that at the forefront of their decision-making in the way they manage their investments. The question is broader than just AI, is it not? Do we require investors in financial technology companies to manage them in an ethical and responsible way? Should we, perhaps? The current legal requirements are such that that is not an obligation.

David Kelnar: The letter of the law and regulation notwithstanding, and I appreciate they are not requirements in that regard, I do think we have quite a significant duty to the entrepreneurs, the management teams and the employees who we support to encourage ethical and responsible

behaviour. It is about a combination of awareness and action. To my mind, awareness is about investors, as for us all, understanding the risks and potential harms of AI from job displacement, entrenchment of bias, increased social inequality, greater concentrations of economic power, trade-offs between privacy and security, and challenges regarding explainability. These are all issues, and we need to be mindful of them. The actions we should take include, in no particular order, participating actively, thoughtfully and honestly in debates regarding AI, such as today's; ensuring that venture capital firms reflect the diverse populations that we serve; engaging with academic institutions, think tanks and corporations to develop best practice frameworks for assessing and thinking about the risks posed by the companies in which we invest; screening candidates to assess their ability and willingness to be responsible stewards in the age of AI; challenging portfolio companies regarding the strategies they are adopting to mitigate the risks associated with AI, and monitoring portfolio companies' adherence to mitigation plans associated with the AI they have developed. I think there is more I and we can do.

The Chairman: Do you see your company taking a leadership role in exactly the respects you have mentioned?

David Kelnar: It is very important to us.

Lord Swinfen: If regulations were to be drawn up, they would be drawn up, I assume, by civil servants. Do civil servants have the knowledge to draw up the right regulations, in your view?

David Kelnar: I am not an advocate for general, somewhat broad-based regulation in this regard. I do not have a good understanding of how that would work. AI is no different from other areas of society and commerce in that regard, so I am not sure how that would work. As with most important things in life, it is down to the individuals.

The Chairman: Codes of conduct, guidance, whatever it may be, it is not regulation, Eileen?

Eileen Burbidge: No. That is what I was going to say. I do not think it should be regulated. As I think David was saying, I do not think it is any different from other areas of corporate social responsibility which we ask FTSE and other companies to abide by. It will be more social pressure and market pressure. It is demonstrated, proven and measured that consumers and businesses prefer to work with businesses who demonstrate or reflect some level of corporate social responsibility, whether that is related to climate, carbon usage, diversity in the workforce, inclusion or anything like that.

Along those same lines, the AI companies or the companies employing AI technology, to the extent they demonstrate they have ethics boards, review their policies and understand their principles, will be the ones to attract the clients, the customers, the partners and the consumers more readily than others that do not or are not as transparent about that. That is the way it shapes the sector and gives us all a moral compass, not from regulation.

The Chairman: Of course, there is already some regulation, is there not,

on automated decision-making, and so on? We are in the foothills already with algorithms.

Q53 **Baroness Grender:** How possible will it be in the future to commoditise personal data and benefit from that commoditisation? Can you cite any models that give the Committee a way forward on this? I am thinking not of any of us here but of future generations, when personal data becomes such a valuable commodity. At the moment, they give it away for free to anyone who asks them.

Libby Kinsey: Quite an active area of innovation, and an interesting investment hypothesis, is looking at the start-ups that are working with ways to give individuals control over their own data. Who uses it, why, for what purposes, and can they withdraw their consent for their data to be used at any time? Similarly, there are a number of start-ups looking at ways to encrypt or anonymise data so that you cannot find out who the underlying person was. Both those approaches will be quite interesting and useful.

You must always recognise, though, is that if you restrict the data there is a trade-off on the accuracy that these machine-learning products will be able to achieve. It should be an informed-consent question. You should know what your data is being used for and, if you withhold your data, what you do not get. There has been a wholesale transference of our data without any informed consent at all. I hope that innovation—again, I am an optimist—will come to our rescue on this.

Baroness Grender: Do you think there may be a point in the future when somebody can literally earn income or reward or some gain from their personal data?

Libby Kinsey: That is certainly the business model of a number of the start-ups that I have looked at, yes.

Eileen Burbidge: I am not sure how that takes shape or form, but I think it is certainly possible.

Baroness Grender: Possible or essential?

Eileen Burbidge: I do not think it is essential. One possibility, which you intimated—this is a contrarian view—is that younger people today give their data out freely. We could be going towards a scenario where data in future generations is commoditised but they are simply not bothered by the fact that their data is being used, to the extent that there is a different definition of unintrusive. Younger people today do not mind that they are served coupons for McDonalds because they went there last week. They do not see that as an intrusion of their privacy; they are quite glad that the app is aware of things like that. I do not think they have suffered ill-effects from it or that they think there could be massive ill-effects. There is a change in culture on that. It is completely different if we are talking about financial services data, healthcare data, and the like, but in general I am not sure that we will have the business model that you have described for general demographic data.

Q54 **The Lord Bishop of Oxford:** If there is one recommendation you would like to see the Committee make at the end of our inquiry, what would it

be?

Eileen Burbidge: I will be greedy and ask for two, but they are related to one topic, skills, which I have been going on about. They are recommendations that I would love to see. While there is a lot of work to do to continue to invest in domestic talent, for the near term, the next two to five years, where there is an acute need I would love to see a recommendation for at least the Home Office Migration Advisory Committee to consider adding artificial intelligence-related roles, perhaps machine learning, neural networks and artificial intelligence expertise, to the tier 2 shortage occupation list.

The second recommendation is related to the same thing and is on the tier 1 route for the digital sector: the tech nation visa, which is administered by Tech City. It is that the quota is increased, perhaps specifically in order to address artificial intelligence as a sector.

The Chairman: It is surprising how many of our witnesses never go for fewer than two or three at the end.

David Kelnar: I have limited myself to two, which are extremely similar to Eileen's. Access to world-class talent and capital are the key enablers for AI start-ups and scale-ups. As the Brexit process unfolds, welcoming rhetoric and expansion of a friction-free visa programme for skilled workers in this area and the expansion of the British Business Bank to increase flow of capital to UK AI start-ups will help to maintain entrepreneurs' access to talent and capital, which are the foundations for the sector.

Libby Kinsey: Thankfully, because I am last, most of the things I would recommend have already been said. It is about ensuring that the UK remains an attractive place for entrepreneurial and academic machine intelligence and is about talent and finance—finding a way to replace the funds lost from the EIF—perhaps investing in larger, more ethical patient funds and marshalling public resources to enable start-ups to reach commercial validation and build businesses. That means computational resource and the data assets that we have.

The Chairman: Thank you very much indeed. We managed to finish within quarter of an hour of when I said we would. It has been a very good session with some really interesting points of discussion. We will read the transcript with great interest, but if you have comments or further points to make, please come back to us. Thank you very much indeed.

Rory Cellan-Jones, Sarah O'Connor and Andrew Orlowski – Oral evidence (QQ 9–17)

Rory Cellan-Jones, Sarah O'Connor and Andrew Orlowski – Oral evidence (QQ 9–17)

Evidence Session No. 2

Heard in Public

Questions 9–17

Tuesday 10 October 2017

[Watch the meeting](#)

Members present: Lord Clement-Jones (The Chairman); Baroness Bakewell; Lord Giddens; Baroness Greender; Lord Hollick; Lord Holmes of Richmond; Lord Levene of Portsoken; The Lord Bishop of Oxford; Viscount Ridley; Baroness Rock; Lord St John of Bletso; Lord Swinfen.

Witnesses

[I](#): Sarah O'Connor, Employment correspondent, *Financial Times*; Rory Cellan-Jones, Technology Correspondent, BBC; Andrew Orlowski, Executive Editor, The Register.

Examination of witnesses

Sarah O'Connor, Rory Cellan-Jones, Andrew Orlowski.

Q9 **The Chairman:** A very warm welcome to our media witnesses. We have Sarah O'Connor, Rory Cellan-Jones and Andrew Orlowski. Thank you very much indeed for coming today. I should put on the record again that the session is open to the public and a webcast goes out live and will be subsequently accessible via the parliamentary website. A verbatim transcript will be taken of your evidence and this will be put on the parliamentary website. A few days after this evidence session you will be sent a copy of the transcript to check for accuracy, and we would be very grateful if you could advise us of any corrections as quickly as possible. If after this session you wish to clarify or amplify any points made during your evidence or have any additional points to make, you are very welcome to submit supplementary written evidence to us. Perhaps we could begin by asking you to introduce yourselves, starting with Andrew Orlowski.

Andrew Orlowski: My name is Andrew Orlowski. I write for a site called "The Register". I have been a technology journalist for around 20 years, 7 of which were in Silicon Valley.

Sarah O'Connor: My name is Sarah O'Connor. I am the employment correspondent at the *Financial Times*.

Rory Cellan-Jones: I am Rory Cellan-Jones. I am technology correspondent for the BBC. I have spent a lot of my career covering business. I came towards technology in the late-90s during the dotcom bubble and have been doing it full time for the last 10 years.

Q10 **The Chairman:** Thank you. We are going to kick off with a fairly big question. What, in your opinions, are the biggest opportunities and risks associated with AI over the coming decade? Shall I start with you, Andrew?

Andrew Orlowski: I have agonised over how to nuance this question. Dame Wendy left and said she hoped we would give a fantastic success story of Britain and AI. One of our jobs, as journalists at least, is to ground things in reality. I think we have been very complicit in creating a kind of dream world or fantasy world. In 20 years of writing about technology this is a very unusual story, in that, almost every time, the technology industry pushes stories to us with enormous fanfare and hype, but this one started with opinion formers wanting to talk about employment. I am not directly answering the question but I am trying to route it. The expectations of AI we may have heard earlier today far exceed what we could be seeing in five or 10 years, at the very least.

Sarah O'Connor: I am not a technology journalist; I am the employment correspondent, so I came to AI, and come to it, looking through the lens of what it means for the labour market. Clearly, the opportunity and potential for AI is that it could mean a step change in productivity. I am sure you all saw the Office for Budget Responsibility's latest updated forecasts. We have terrible productivity in this country. It is a massive problem. Basically, if you do not have productivity growing at a decent clip then you cannot have sustainable increases in living standards. Being able to do things more efficiently and create more economic output with the same amount of people is really good news. That is the opportunity. As for the risks, I think there are more than one. The big question that people in the economics and labour market world are thinking about is: how will those gains be distributed? If indeed AI leads to vast increases in efficiency, using fewer workers, does that mean that all the wealth that is created from that will go to the people who own the AI—the intellectual property—and the data that feeds into it? If so, what does that mean for the people who might be displaced out of jobs? Will there be new jobs to replace those old ones? If there are, will they be of the same quality? I think that is the main risk.

Rory Cellan-Jones: I slightly echo what Andrew has said, although perhaps not in so extreme a way. There are trends in technology and technology journalism. I think it is Deloitte that produces a hype cycle whereby technologies go through a wave: they reach a peak of hype, they fall into a slough of uncertainty and then they rise again into fulfilment. We are possibly at the peak hype stage. I look at my inbox every day, and five years ago it was all about big data; three years ago it was all about the cloud and now it is AI. Just picking out a couple: "AI wearable solutions for ageing population" has popped into my inbox. "Lloyds of London sign first ever AI deal in a decade". "A significant part of the insurance industry will be powered by AI." "Welsh artificial

intelligence firm has produced a report into North Korea's bioweapons capability". All sorts of extraordinary claims are being made for AI.

That is not to say that this is not a time of great advances. People who have been in the AI business for a long time will talk about AI winters. They came out of a winter about five years ago and began to find ways of harnessing the huge stores of data to which they had access into new algorithms, and have produced some interesting results. You think of all sorts of possibilities that could come from this: more effective ways of treating cancer; organising the way cities work to minimise congestion; improving the delivery of education. I think there is possibly a bit of overexcitement. I heard the distinguished scientists disagreeing about, for instance, how quickly we are going to get autonomous cars on the road. That is the area of risk. I feel partly responsible for the hype, in that three years ago I did an interview with Stephen Hawking in which he said that AI was likely to kill us all in 40 or 50 years' time.

The Chairman: This is what we are trying to do: attribute blame.

Rory Cellan-Jones: We probably do not need to worry about those kinds of issues in the near term. We need to worry about things such as bias in algorithms being built in and about what sort of infrastructure we need to make driverless cars work, for example. We have been far too optimistic in thinking how quickly we will get there.

The Chairman: That is very useful because, in a sense, you have gone from talking about the hype side to the limitations. Andrew, do you agree that one of the reasons for your being rather sceptical is the limitations with AI currently?

Andrew Orłowski: I think so. It is important to distinguish. AI has spent most of its history ignored and derided. As Rory says, you have great long winters, then almost a summer's day and then everyone drifts off again and they are not interested any more. There is one enduring paradox coined by roboticist Hans Moravec 30 years ago. After AI had already been researched for 20 or 30 years, he pointed out that it is relatively easy for a computer to play adult-level games but it is extremely difficult to replicate the perception and mobility of a one year-old.

To me, the biggest discovery, looking at this current AI hype, is that until the early 1970s robots were an expression of AI. Today the machine-learning community and the robotics community exist in a state of antagonism. The machine-learning people say, "Your real-world data is dirty; we want nice clean data", and the robotics people say, "Your algorithms are too slow. They are too messy. They do not stop us hitting an old lady crossing the road". This is why I conclude that for huge leaps in productivity from AI, we need huge qualitative leaps in automation. We are getting great advances little by little every year, but it does not involve the one big innovation of the last five years, which is machine learning.

The Chairman: Are you conscious of the hype aspect, Sarah.

Sarah O'Connor: I am very conscious of it, and the way that I know it is true is partly, as Rory says, the sheer volume of press releases that come

into your inbox. Five years ago companies might have described what they did as big data or data analysis, and they have simply rebadged what they do as AI because they know it is sexier and that journalists are more likely to open their emails if that is the subject heading. The other way I know it is overhyped is that if you ever write an article that has robots or artificial intelligence in the headline, you are guaranteed that it will have twice as many people click on it. People are really interested in this topic. That is not to be sniffed at. If people are interested, we, as journalists, have a responsibility to tell them about it, but for some organisations—and I would not include the *FT* in this—there is probably a temptation to chase the clicks, so the more you can write and probably the more dramatic things you can say about it, the more clicks you will get and, therefore, the better you will do.

The Chairman: We are warned.

Lord Hollick: You referenced productivity. Why is it that the very significant investment in digital technology and computers over the last five or 10 years has not led to a higher level of productivity?

Sarah O'Connor: That is a really good question. That is a puzzle that is taxing the best minds in technology and economics right now. One possibility is that it simply lags and that if you look back through history at the introduction of big impressive technology, the point at which it is adopted and people figure out how to use it best in order to boost productivity takes some time. It could be that, or perhaps we have become more efficient at some things and we have found other pointless things to fill up the rest of our time.

Q11 **Baroness Rock:** You have talked about the press releases from companies and chasing the clicks. I would like to look at the nature of your audiences and how they respond to the pieces you write on AI. Particularly, do you see a difference in the understanding of the risks of AI between the more mature and the younger audiences?

Rory Cellan-Jones: I think it is difficult to know. We do not put out a piece on the "Ten O'clock News" or write a piece on the website and then send people an exam to see what they have and have not understood. It is one of those things that journalists, rather arrogantly, think they know. We think we know what is interesting; we think we know what is a well-told piece that will fascinate people. I find it more difficult, particularly on high-profile outlets with big audiences, to find an excuse to tell the long-run story. Daily news outlets are very event driven. You might go down to the editor of the "Ten O'clock News" and say, "Listen, AI is going to have a huge impact on healthcare" and that person will say, "Yes, but what happened today?" Our job is to find ways of getting arresting stories that tell the longer-term story in accessible ways without either hyping or underestimating the impact of the technology.

The Chairman: You have a slightly different sort of journal, Andrew. People will seek you out to read about AI, will they not?

Andrew Orlowski: Yes. We are almost where the British technology community goes. They are a very experienced audience and naturally sceptical, because I think the right way to approach artificial intelligence

is through scepticism. Let me give you an example. I am not aware of one self-driving truck that can park. To me, if all a truck can do is drive along a nice clear road in nice weather for a few miles, that is exactly what cars in the marketplace do now. That is not autonomous, it is semi-autonomous, and the vehicles we have today are semi-autonomous. They have lane warnings for the driver, and so on. I think the desire to anthropomorphise is everywhere in AI; it is there in the term “learning”. One of the great British computer scientists, Christopher Strachey, really objected to the word learning being used for this. He said they are optimising themselves slightly but they are not learning. Learning assumes that it will reach a certain level then it will be able to do everything that it was set to do.

The Chairman: Sarah, do you have a perspective on this?

Sarah O'Connor: We cannot segment the response of our audience by age, as such, but I would say that when I read letters that have been sent in after I have written about this, or I look at the comments that come in under the articles, probably more than half of our readers are really worried about this. Rather than, perhaps, the technology audience being quite excited, a lot of my readers worry about the implications and, particularly, the economic implications. There are lots of people, as soon as you write about AI, who will instantly start commenting, saying, “This is why we need a basic income. Capitalism is going to eat itself”. There are people who are looking very far ahead and reaching slightly scary conclusions.

Rory Cellan-Jones: I would say that about a lot of the BBC audience. We did an AI week a couple of years ago and the most popular feature was a rather inventive graphic which allowed you to put in various professions and ask, “Will a robot eat my job?” Politicians came out okay and barmen pretty well, but luckily the accountants are doomed.

Sarah O'Connor: As are the journalists.

Rory Cellan-Jones: As are the journalists. Interestingly, the data that went into that graphic came from an Oxford study which has been very widely quoted over the years, which said that more than 40% of jobs are potentially at risk over the next 10 to 20 years. That data has been questioned, I think, by the OECD, who brought it down to a figure of 7% or 8%, so perceptions are changing pretty rapidly here.

Viscount Ridley: I am trying to reconcile the extreme scepticism of you three with the extreme excitement of the previous three. I know I am exaggerating but I am a journalist, like you, so I do that kind of thing. I am thinking of Amara’s law. Do you know Amara’s law? Roy Amara was a computer scientist who said, “In the short run we overestimate the impact of new technology but in the long run we underestimate the impact”. Is it possible that that is what we are seeing here: you guys are going to be right in 10 years but Professor Wooldridge is going to be right in 20 years?

Rory Cellan-Jones: That is entirely possible. I have come across as very sceptical. As technology journalists who have been bombarded with this tidal wave of hype, we recoil slightly. On the other hand, we also like to excite our audiences, so we are split on this. Is the term presbyopia?

That applies to technology. We have seen, for instance, the smartphone revolution over the last 10 years, which has happened much faster than many predicted. Perhaps we are wrong about how rapidly this one will happen. What strikes me, though, is how rapidly software is advancing and how slowly hardware is. In fact, we are going backwards. Ten or 15 years ago I could fly to New York in three and a half hours on Concorde. In the Victorian age they built a railway in five years. There is a big divergence here.

Baroness Bakewell: I have a quick question. Newspapers are chasing audiences now. The BBC, as we know, is desperate to recruit young listeners. I wonder whether you have arenas in which you can use extremely simply-worded explanations intended for younger generations—"John Craven's Newsround", for example. Does that kind of speech occur in your profession?

Rory Cellan-Jones: To be honest, we have the opposite problem, in that there is a great appetite for and, perhaps, an awareness of this amongst our younger audiences. This is one of the few areas of our coverage where that is not so much a problem; this is seen as something that reaches out to young people.

Baroness Bakewell: You mean they know as much as you do?

Rory Cellan-Jones: Yes. Probably more.

Baroness Bakewell: I wondered about the simplification of things so that a mass audience can understand. Is that possible?

Sarah O'Connor: I think it is possible. In a way, that is our job. We sit between those guys and the general public, and our job is to understand enough of what they say to translate it into plain English but in a way that does not simplify and lose some of the reality. That is a difficult job, whether you are writing about economics or any complicated topic. That is what we should do. Perhaps we do not always get it right. At the *FT*, clearly, we do not have many very young readers.

Rory Cellan-Jones: Jacob Rees-Mogg is a reader.

Sarah O'Connor: I always try to imagine, "Could an intelligent 18 year-old understand this; someone who has a decent grasp of the English language who, perhaps, does not know anything particularly technical yet?" That is what we aim for.

The Chairman: People clearly react to your articles too. You mentioned earlier that they come back at you on that.

Sarah O'Connor: Absolutely they do, yes.

The Chairman: Which is good. Lord Giddens.

Lord Giddens: From what you have said so far, the digital revolution has really transformed journalism in the most extraordinary ways. Lots of newspapers are finding it close to difficult to survive. You have massive platforms, some of them organised through bots rather than people. This is very widespread and has had a massive impact on politics. What you are saying, to me, sounds a bit sanguine against that backdrop, because this is surely changing the climate in which journalism works and

transforming the whole business in ways which we do not, as yet, fully understand and are only some way through, probably.

Rory Cellan-Jones: Certainly the impact on journalism is where AI comes in. Yes, the digital revolution has obviously transformed journalism, and, for many people, not in a good way. There is a long hunt for sustainable business models. There is the recent arrival of the fake news phenomenon and what some might see as the malign power of giant not-media companies such as Facebook and Google.

Lord Giddens: If I can interrupt, 24-hour breaking news is pretty recent and we do not know what that does to people's levels of consciousness.

Rory Cellan-Jones: Yes.

Lord Giddens: It is available everywhere in the world immediately.

Rory Cellan-Jones: I suppose we, as journalists, do not necessarily think day to day about our impact on the psyche of our viewers, except to hope for heaven's sake that they are watching and listening.

The Chairman: We are in danger of straying well beyond AI, at this point. Andrew, just a quickie.

Andrew Orlowski: I would make a very quick point. I do not think anybody can say that the market online is healthy. At the moment, there are far fewer journalists than there used to be and the incentives are not there to do long-term pieces. One of the ways we can act, where the market is particularly poor, is in titles such as ours, where we act as a bridge. We know the technology community very well but it is not picked up by the wider audience because they have so many other things to do. One example: Geoffrey Hinton, who is the father of machine learning—the one big breakthrough we have seen—about a month ago said that we need to rip it up and start again to achieve what is now expected of us. Two or three titles—the technology reviews, the business titles—have picked it up in the States, but it has not got into the mainstream yet.

Q12 **Baroness Bakewell:** This builds on what I asked before. Should efforts be made to improve the public's understanding of and engagement with artificial intelligence? If so, how? Could you answer beyond journalism? How are we to do broadcasting? Where are we to take a government effort to expand public understanding, or indeed just social effort or academic effort? Where can it come from?

Sarah O'Connor: To begin with, journalists clearly have a responsibility to do that. We have a difficult balancing act, because we do not want to fall into the hype trap but, as one of you was saying, we do not want to underestimate what could potentially be really important. We need to try to get that right, and that is not always easy. One way I have tried is to give readers a sense of what is actually happening now. Rather than simply report on these various studies that Rory mentioned—"40% of jobs might go"; "No, it is 50%"; "No, it is 10%"—basically, nobody knows and there is only so far you can go with that. I try to go out and meet people who are using AI now and figure out, what does it mean right now? How does it work? What are the pitfalls?

On Lord Giddens' point about the impact on journalism, last year I challenged an AI journalist to a journalism battle.

The Chairman: Was it on the scale of AlphaGo?

Sarah O'Connor: I would not put myself in quite that intellectual category. The idea is that it is fun and engaging, and people want to read about it. Also, it is an attempt to get to grips with where we are right now and how much further we need to go before we should really worry. For me, the lesson of that battle was that I definitely won and that for the time being, I do not think we have to worry about AI replacing journalists.

Baroness Bakewell: Do we want the public at large to be better informed?

Sarah O'Connor: Yes, of course.

Baroness Bakewell: How?

Sarah O'Connor: By explaining it better; by having more intermediaries who can speak the language of the technologists and, also, speak human; and perhaps get the academics out and about more and teach them how to communicate in plain English.

Baroness Bakewell: Do you think the BBC has a special responsibility?

Rory Cellan-Jones: Yes. We are always trying to find good communicators among the technology community, which can be a huge challenge. You had three pretty good communicators on before us, who we try to put on air a lot. We see it as a public role to educate but, also, to entertain along the way. We have to be careful about being preachy about this; saying, "You have a duty. Sit up straight at the back; you have to learn about this." We have to be honest about it.

Baroness Bakewell: Sugar the pill?

Rory Cellan-Jones: A more important role, certainly for government, is in preparing people, from an educational point of view, for these challenges. I have been reading an extraordinary book called *Janesville*, about the impact of the closure of a motor manufacturing plant on a town in America. They found that quite young car workers made redundant in their 30s were completely incapable of starting an educational programme because they had never used computers. There are still quite a lot of people out there who are uncomfortable with computers because when they were at school that was not a big thing. A public education programme sounds slightly questionable.

The Chairman: Anything more, Andrew?

Andrew Orlowski: I hate to be in this position of a contrarian, once again, but my children go to an outstanding primary school in north London where they are taught algorithms every week but they are taught history once or twice a term and art, maybe, once or twice a term. There is an opportunity cost; there is only so much time for educating people. I question the value of teaching them algorithms. That is probably part of a balanced curriculum, but if they do not know culture and history how can they account for the world? How do you account for China's relationship

with Korea? It is complicated, cultural and historic. You need those things probably more than you need to know how to use a computer.

The Chairman: You are slightly echoing something that Dame Wendy said earlier, that she thought some of the creative skills were going to be needed in the future as well as those technical skills. Thank you. Lord Swinfen.

Q13 **Lord Swinfen:** I run a charity that uses a secure system to provide specialist medical advice to doctors and other medical workers in some 77 different developing countries. Do you believe the media are generally covering developments in artificial intelligence in an accurate and responsible way? How could this be improved?

Rory Cellan-Jones: That is a biggie. Given how broad our media are and how vibrant they are still, whatever the challenges in this country, they range far and wide—from “Killer robots are going to eat your lunch” to learned dispositions in the pages of the *FT* about complex studies on the patterns of employment. I would say this, but I looked at the stories we do at the BBC—I did a quick search at the weekend—and we had a story a day from various parts of the world: a story about Chinese peach farmers who brought in AI experts from Beijing University to help them sort peaches more quickly; stories about the impact on women, as opposed to men, of advances in automation. I think we are doing a pretty broad and, generally, sensible job, with the occasional bout of alarmism.

The Chairman: There was almost a concept of exciting scepticism, which we had earlier from you, Rory, which was quite interesting. Sarah.

Sarah O'Connor: I would pretty much agree with Rory. The quality of the reporting on AI varies from media outlet to media outlet, much as it does on every other topic. If you read the tabloids you will not get a massively subtle view of the impact of AI, but you will not get a subtle view on anything else either, probably. There is a lot of good reporting out there, but one thing I am conscious of is that even when you want to do a good job and be responsible, you have to take the time to really understand how it works and get into the nuts and bolts of it. That is time-consuming, and journalists do not have a lot of time. Sometimes it is easier to skim over the details and write up the big top lines. I have been trying to make an effort to sit down with people who do AI and who create AI systems to make sure that I really get what I am talking about.

Andrew Orłowski: I alluded earlier to the fact that we were creating a fantasy world and not reflecting the very real limitation that robotics will progress at a very steady rate; we are not going to see a giant step change. Rory referred to doubts cast on an employment report. I do not know if we have an equivalent in Britain of this wonderful American phrase “check kiting”, which refers to a kind of cheque fraud whereby the cheques are constantly signed but the funds are never in the account. The primary source for that—I think it is the most-quoted report on employment on Google, with 1,300 citations—is a 2011 book by Brynjolfsson and McAfee. Brynjolfsson and McAfee are the source for things improving exponentially. Their new book quotes the James Martin report. It is beginning to look very, very circular.

The Chairman: This is a tricky question. Who do you think is leading the field in terms of anticipating the issues and communicating them?

Rory Cellan-Jones: Andrew Orłowski, definitely.

The Chairman: We wanted a bit of a love-in.

Rory Cellan-Jones: Are you talking about journalists?

The Chairman: The media, broadly.

Rory Cellan-Jones: To be honest, the best communicator in this whole area is Demis Hassabis from DeepMind, who is a very engaging, brilliant and interesting communicator. A documentary has been made about the AlphaGo experiment.

The Chairman: Some of us have seen it, yes.

Rory Cellan-Jones: Whatever you think about the pace of this innovation, if you were to ask the public—and not many people, I am sure, would quote any name—I would have thought he was the best communicator around the subject.

The Chairman: Andrew, how do you respond to that flattery from your left?

Andrew Orłowski: I would agree that Demis is a phenomenal showman, but we need more than showmanship for this. In the end, we must remember these are games. As we were saying 30 years ago, playing games is fantastically easy and impressive, but if we are talking about productivity gains we need a robot—and it is a fine aim—that can process the dirty data of the world and make sense of it. We make much less sense of data. The names I scribbled down were Jaron Lanier, who knows his stuff. He writes very well. There is a tremendously realistic sceptical piece by Gary Marcus, who is a science writer—a neuroscientist—and was head of AI for Uber. Ian Bogost writes very well about technology. He is not a technologist but he is a professor in Georgia, I think.

The Chairman: Sarah, do you have any favourites in this area?

Sarah O'Connor: In the world of work, which is where I focus, I would say the people who are doing interesting things are some of the think tanks sitting in that public journalism-facing space. The Resolution Foundation has done some very thoughtful quantitative research. The Royal Society for manufacturing and the arts has just put out a very interesting report. Andy Haldane, the chief economist at the Bank of England, is another very good communicator. He is starting to think about that now and there are some economists at the Bank who are doing some decent work on the impacts on the labour market. Those are the people, I would say.

The Chairman: Great. Thank you very much.

Baroness Bakewell: Let us talk about other media. What is the impact of these great fantasy films and television series? They are entirely fantastical but they encourage people to think about AI, do they not?

Rory Cellan-Jones: You are probably right. I would imagine films such as “Ex Machina” and—what was the Channel 4 series? Was it called “Humans”?

The Chairman: “Electric Dreams”?

Rory Cellan-Jones: There was a more populist drama about people acquiring a domestic robot—“Humans”. I am sure, for good or ill, they have a greater impact on public understanding/fears/interest and excitement in this subject than we can have.

Q14 **Viscount Ridley:** It is not your job to write government policy, but we did ask the other panel this question and it is worth asking you. Are the Government currently doing enough to maximise the opportunity and to minimise the risks associated with artificial intelligence? If not, what could be done?

Rory Cellan-Jones: As you say, it is not my job to advise the Government, but I would look at two areas, including ensuring that we continue to be a place where this research happens. We have heard how strong we are. You think of the people behind SwiftKey and the people behind DeepMind. Also, thinking about the implications of the design of our cities if this driverless car revolution happens—about which I am pretty sceptical—there will need to be a lot of work done on rules about insurance and how people drive. At what stage are we going to allow people to take their hands off the wheel? It does strike me that we are doing pretty well at that. We are thinking about it, and the fact that you are having this inquiry suggests we are thinking about it pretty strongly and, possibly, faster than we, on this side of the table, might think quite necessary.

Sarah O'Connor: I get the feeling that the Government are at the stage of trying to figure out, as we all are, how seriously to take this and how far along the track we are. The Government probably do not know any more than the rest of us exactly where this is going. My feeling is that, obviously, you need to start anticipating and thinking about what the impacts might be, but you cannot predict the future and you do not want to put in a solution too quickly that is not the right one.

One thing the Government can know now, from the labour market point of view, is which jobs are definitely going to continue to exist. Which jobs are we going to need in 10 years’ time and 20 years’ time? It is easy to do that. I can tell you that one of the ones we will definitely need is social care workers, health assistants—people to look after people. Robots will never be able to do that. It does not matter how good they are. Nor, I think, should we really want them to. If I were the Government I would be thinking, “Okay, these jobs are going to be more and more in demand because the population is ageing. AI is not going to solve that problem for us.” If more and more people are doing these jobs, how do we make them good jobs that people can build a life on? Otherwise, we will end up with an awful lot of people doing jobs that cannot sustain them in the long term.

Andrew Orlowski: Briefly, as you know, my list of risks is different and slightly more subtle. In my risks column I would have liability issues,

where nobody carries the can for something. Very specifically, the Silicon Valley ideology is that we cannot own our own data. These are quite subtle things. How would a Government plan for that? As Rory said, you have already thought about this a lot. I should caution that in the 1980s, Japan put enormous amounts into artificial intelligence only to find it was a complete waste of money. One suggestion around that is red teams. You may come across the idea of red teams, where you have, almost, a professional naysayer within a group or a project who points out everything that is wrong with it. It is probably the world's most annoying person ever. Journalists should be doing that, but I think we love this dream world we have constructed.

The Chairman: You are writing yourself a job here, Andrew?

Lord Holmes of Richmond: Sarah, it is a commonly held view that, and you seem certain that, AI—robots, or whatever—will not be able to do certain roles, particularly in the caring field. I am not suggesting that it will, but what evidence do you base that certainty on?

Sarah O'Connor: Two things: the first is where the technology is now. To look after someone you need to be extremely dexterous; you need to be able to help someone get out of bed; you need to be able to hold them gently without bruising them. We are talking about robotics now, rather than AI, but robots are not capable of these things. There is a reason we have low-paid migrants plucking soft fruits; it is because robots crush them. They cannot do those sorts of things. That is the physical side. We are really far away from that. Robots cannot tie shoelaces. They cannot do these things; you need human hands. On the emotional side, I think part of looking after someone is being a human being and showing them some compassion. AI is nowhere near being able to replicate that.

Rory Cellan-Jones: Can I come in on that and agree? I went to an event last year where a Chinese roboticist was showing off a fantastic robot with an incredibly lifelike face and gestures. He said, "This is going to be used in care homes and is going to be particularly good for treating people with Alzheimer's because it can listen to the same story a hundred times and not lose its rag". I found this a rather horrific vision—that we would think of putting that kind of device in a home.

Lord Holmes of Richmond: If it then turned out to be Donald Trump.

Q15 **Lord Hollick:** In the last session the witnesses described themselves as reasonably sanguine about the issues of privacy and exploitation of the data. Professor Wooldridge said that he felt that the law was adequate. Do you share that view, or the view of many that the concerns around privacy and the exploitation of datasets unduly favour digital mega-corps?

Andrew Orlowski: These are well-founded concerns. As I mentioned in my last answer, if there is one consistent kind of political thread that comes from Silicon Valley it is that they want to exploit data and virtual or intellectual property. It has got to the stage where the incentives are so lined up with large Silicon Valley companies that they do not think of it as a property-based market where that data is traded. When you get into

the area of health data, it is not an easy subject; it becomes a prickly subject. If we remember that the individual is sovereign here—they can have contractual relationships with the private sector or with the state, but the individual's sovereignty is paramount here—that is probably the basis on which to proceed. If we go in any other direction, I see disaster.

Let me give one example which is slightly concrete. In software, it is possible to use data on some licensing models quite freely—a teenager in their bedroom can play with the data—but as soon as it starts to be exploited commercially and has value to the user, then they pay a royalty. That seems quite a conventional idea that we have never tried with data. Perhaps Wendy is looking into it.

Viscount Ridley: It sounds a little like the argument in the 1920s that if an aeroplane flew over your garden, you had a right to a fee from it. That was never going to work. Likewise, are there not unrealistic aims around personal sovereignty that we have to grapple with?

Andrew Orłowski: There are. There is also the idea that information does not have value unless an algorithmic company has processed it; it does not have value, therefore that company should exploit it. It is a case where the individual ought to be able, not to force a plane off its route but to say, "You are not going to use this data in this context because I do not want my insurance company knowing this". Where do you stop? If you have the most optimal system in which the individual's sovereign rights are negated somehow, there is absolutely no privacy. We have to find a balance between the two.

The Chairman: Thank you. We will be exploring those. Just a supplementary on this point, do the Government engage with you on issues relating to AI? How much to-ing and fro-ing is there in terms of engagement?

Rory Cellan-Jones: I am trying to think whether I would welcome a lot of to-ing and fro-ing with government on this subject. Not much, but I have not sought it out, I have to say.

Sarah O'Connor: I was invited by the Treasury last year to take part in a labour market conference day where AI and the impacts on the labour market were one of the things being discussed. That is probably the most engagement I have had with government on this.

The Chairman: We have settled that one. Over to Lord Giddens.

Q16 **Lord Giddens:** I am supposed to ask you about the impact on labour markets but you have been talking about that already, so I will put it in my own inimitable fashion. We have had these amazing swathes of change before: 70% of the adult population worked in agriculture and now 1%—amazing; 40% of the adult labour force used to work in manufacturing, now 9% in this country and 8% in the US—amazing. Now you have mostly a white-collar service economy, the question is, how far will it be invaded by AI especially? What do you make of the consequences for future government policy? What should government be doing now? To me, this issue is very real.

You mentioned Carl Frey's study in Oxford, but the 47% figure was a job

breakdown analysis; it was not a prediction of the proportion of people who would lose their jobs. To me, this is really “don’t know” territory. The previous people we were talking to said that new jobs will be created. It is possible, but it is “don’t know” territory. It is moving so fast. What exactly should government do? What kind of policy recommendations would you make, for example, for different sets of the labour force, because you are obviously going to get a much higher level of job churn than you probably ever have before in service and white-collar jobs? It is bound to invade professional jobs too.

Sarah O'Connor: Your statistics about the proportion of people who used to work in agriculture are useful in reminding us all that normally, there are these technologies that change everything and new jobs are created. The employment rate now is as high as it has ever been, even though no one works on farms and not many people work in manufacturing. It is still plausible that new jobs will be created, but you are totally right that it is “don’t know” territory. We just do not know. Part of the reason, to be honest, we are all sitting here is that people in blue-collar jobs have been disrupted for decades but the thought it is going to hit the white-collar middle classes, I think, has scared a lot of people. It is already happening.

Lord Giddens: These are the only jobs left.

Sarah O'Connor: Yes, partly that and partly because it is us now. Let us be honest, I think that is partly where this sudden set of fears has come from. It is starting to happen. I went to see a law firm recently that has developed an AI to, effectively, automate the due diligence stuff they do when there is a big M&A transaction. Rather than having 20 junior lawyers sit in a room for two nights drinking Red Bull, reading through every page, the AI will scan through, look for things it thinks are potentially problematic and then give them to more senior lawyers to look at. That sort of disruption is already beginning, and you can see that it will continue.

For government, my big question is: what do you do about training and allowing people to cope? We cannot predict the future but we probably can predict that, as you say, there will be more churn and people will need to adapt more quickly. Do we need to think about giving everyone some kind of personal training budget that they can spend on reskilling? Those ideas have been tried out in the past and they have not always been successful. Rather than trying to predict where the economy is going, which is always a mug’s game, trying to think about how we equip people to be as resilient as possible to the changes is probably the best way to go.

The Chairman: Thank you. Any addition to what Sarah has said?

Rory Cellan-Jones: Only that, as well as these forward-looking studies, there have been a couple of looks back at what has actually happened over the last 20 years. We have found not entire professions wiped out but certain tasks being wiped out. The law is a perfect example. We might like to imagine that lawyers will be wiped out by automation but it does not appear to be happening.

The Chairman: I am not sure I approve of this.

Rory Cellan-Jones: The dull work is being done, increasingly, by machines but they are finding new things to do. There is the classic line about bank tellers in America. I am not quite sure where it comes from but apparently, there are more bank tellers today in America than there were 20 years ago, despite the arrival of the cash machine, and so on.

Andrew Orlowski: Sarah's point is an extremely important one.

Lord Giddens: That is not the same as net new jobs, obviously, with such a high level of job destruction.

Andrew Orlowski: I think Sarah's point that the story took life because middle-class jobs were under threat is an important one. Essentially, that has created where we are today with AI—the AI panic of today. It is interesting that nobody has mentioned universal income, which has been proposed by Silicon Valley leaders. They are almost saying, "We are creating this problem for you but, oh look, we have this solution". I am relieved not to hear it because I think it is a very premature debate to have. There are huge implications for social mobility, and so on, if we expect everybody to be on the dole for their lifetime.

The Chairman: Thank you. A final question from Lord St John of Bletso.

Q17 **Lord St John of Bletso:** Before my question, I make the comment that many organisations which are currently using robots would argue that they are doing it more to enhance their service offering than to replace jobs.

My final question, wrapping up, is: if there was one recommendation at the end of this inquiry that you would like to see us make, what would it be?

Andrew Orlowski: I would re-emphasise the point about red teams. It is really exciting to have that robot that can do what we are being promised today but it is not there, possibly even with today's scientists, techniques and machine learning. If we are going to do research, let us take real risks with really strange ideas and let them run for years, rather than looking for instant payback. One of the problems with this debate is almost everybody has skin in the game; even the universities want spinoffs and venture capital investment in their latest big idea. Often, this stuff takes years and is very painful to do. A little courage from the funders is needed.

Sarah O'Connor: I would repeat the point that, rather than trying to second-guess which jobs will go, bear in mind which jobs will stay and make sure they are decent jobs that people can make a decent wage on.

Rory Cellan-Jones: I would make a general point about education; not about teaching all kids algorithms but about having a more flexible attitude to what kids learn, combining creativity with digital skills. The growth industries, such as the games industry, employ all sorts of people; they employ artists, designers and some mathematicians and physicists. Too often we seem to be producing very rigid schemes of education. We are going to need people to be more flexible in a more flexible world.

The Chairman: Thank you very much indeed. That concludes the second

Rory Cellan-Jones, Sarah O'Connor and Andrew Orłowski – Oral evidence (QQ 9–17)

half of our evidence. That was a good session.

CIFAR – Oral evidence (QQ 172 – 180)

Evidence Session No. 18

Heard in Public

Questions 172–180

Tuesday 5 December 2017

[Watch the meeting](#)

Members present: Lord Clement-Jones (The Chairman); Baroness Bakewell; Lord Giddens; Baroness Greender; Lord Hollick; The Lord Bishop of Oxford; Lord Puttnam; Baroness Rock; Viscount Ridley; Lord Swinfen.

Examination of witness

Dr Alan Bernstein [video link].

Q172 **The Chairman:** Good afternoon, or should I say good morning, Dr Bernstein?

Dr Alan Bernstein: You can say either.

The Chairman: Excellent. That is very flexible. Can I welcome you to this evidence session of our Select Committee? We are delighted that you are able to be with us. I should explain that this is the 18th formal evidence session for our inquiry into artificial intelligence. This session is intended to help the Committee discuss the international aspects of artificial intelligence and the policies that other comparable countries are pursuing with regard to artificial intelligence. This is a public evidence session and it is being broadcast. It is being streamed. A transcript will be taken of your evidence, which you will be invited to check for accuracy. If, after this session, you wish to clarify any points made during your evidence, we would very much welcome supplementary written evidence. I will not ask you to introduce yourself because we have your particulars as president and CEO of the Canadian Institute for Advanced Research, and we would very much like to hear your response to our questions, of which I think you have had a copy.

Dr Alan Bernstein: I do, yes.

Q173 **The Chairman:** Thank you very much. Allow me, therefore, to kick off with the first question. What are the biggest advantages and disadvantages that AI could bring to Canada over the next 10 years? What is distinctive about Canada's approach to AI? We need to set that in the context of the advantages and disadvantages that Canada has, in particular, when capitalising on AI.

Dr Alan Bernstein: Thank you very much for this opportunity to appear in front of your Committee. I appreciate this. These are international issues not unique to the UK or Canada, so it is important that we have a

dialogue across countries. With respect to your question, some of the advantages and disadvantages are not unique to Canada but shared by many countries, probably including the UK. Let me dive into them. One of the first advantages of AI, of course, will be to increase productivity—that is why the private sector and Governments are interested in AI—to, perhaps, change the nature of work to make it less rote or routine and more focused on conceptual and people skill issues. It is also, of course, an opportunity for what we call in Canada innovation or the creation of companies and increasing the diversity of the Canadian economy. In brief, those would be some of the major advantages.

The disadvantages or the challenges—the word I would prefer to use—include issues focused, in general, around innovation more broadly, which is that there will be some haves and have-nots when innovations are introduced into our societies. There will be those people who will selectively benefit and those who will be selectively disadvantaged. I am thinking of job loss not uniformly across the economy but selectively in certain areas. Again, that is not unique to Canada. It raises issues that require Governments, especially, to think about policies around skills development for this new world of AI we are venturing into. There are also challenges with respect to inclusivity, bias, privacy issues, and so on. I could go on, but perhaps I will stop there.

The Chairman: We will come on to some of those later. What does Canada, in particular, bring to the party on AI?

Dr Alan Bernstein: We bring a number of things. First, it is important to note, and we are quite proud of this, that deep learning, which is universally recognised as currently the most powerful form of AI, was developed under CIFAR's auspices in Canada by Geoff Hinton and his colleagues. As a result of that, there is considerable talent in Canada in deep learning and other forms of machine learning and AI, largely focused around Montreal, Toronto and Edmonton. Talent is, of course, key to any hot area of innovation and technology, and we have a lot of it here. That is one of the things that is most unique about us. Secondly—and this is an advantage and a disadvantage—we are close geographically and, to some extent, culturally to the US. It is relatively easy for people, capital and ideas to move back and forth between the largest economy in the world and the smaller economies in the world. That is both a challenge for this country and an opportunity. We have very liveable cities. Our major cities are constantly ranked among the top five in the world. Talent of any kind, but especially in science and innovation, wants the opportunity both to pursue their ideas and to live in liveable cities.

The Chairman: Terrific. I will stop you there because we will move on, but that is a pretty good list, if I may say so. Is it your view, on balance, that the impact on employment will be positive?

Dr Alan Bernstein: My view, on balance, is it is hard to predict. We need a finer-grain look at employment. In other words, we need to look at particular job classifications and industries rather than employment as a whole. It is fair to say there will be some job losses and job gains, but what the net result will be is crystal-balling, and I prefer to see data.

Q174 **Lord Swinfen:** Good afternoon. What issues does Canada have in attracting and retaining skilled AI researchers and developers? How is Canada attempting to tackle these problems? What role will the Innovation and Skills Plan have in this?

Dr Alan Bernstein: That is an excellent question. We have the usual challenges of, as I said, being close to the United States and Silicon Valley, so the young people being trained at our universities, of course, are very attracted to both the opportunities and the high salaries in the United States. The Federal Government in Canada has recognised this challenge. One reason they asked CIFAR to run the pan-Canadian AI strategy—the \$125 million investment announced in the last Canadian federal budget—was in large part to attract and train talent and keep it in Canada. The way one keeps people anywhere is through a decent salary, opportunities to pursue what they want to do and the opportunity to live in attractive, dynamic cities. We are addressing all the above.

Lord Swinfen: Thank you very much.

Q175 **Baroness Bakewell:** Dr Bernstein, good afternoon to you. What role is the Government of Canada seeking to play in the development and utilisation of AI in Canada and the world? I see that in a recent interview you said that Canada was still playing catch-up. Is that the case now?

Dr Alan Bernstein: It is hard to measure. I do not know what the metrics are that one would use. What I meant was that there are significant investments by the private sector and by Governments—here I am thinking of China, in particular—in AI, and Canada certainly cannot match the private sector investment that Silicon Valley companies are capable of bringing to the table. To go back to your question, I think that the Federal Government have viewed their role from the beginning as being to fund the fundamental research that has led to the breakthroughs that we are all so excited about and talking about. That goes back, in our case, to the 1980s when, through CIFAR, the Canadian Federal Government funded a programme that we called artificial intelligence robotics in society, in the early 1980s. That ended and a new programme started, with Geoff Hinton, who is universally acknowledged as the godfather of modern AI, as director. Again, that was funded in part by the Federal Government. Governments have an important role in funding the fundamental research that the private sector does not fund anywhere.

Baroness Bakewell: You have spoken already about the \$125 million that have been invested. Will this be sufficient?

Dr Alan Bernstein: That is a good question. Perhaps the way to answer that would be to talk about how it is going to be used, and we could go from there. A large chunk of that money will be used to fund Canada CIFAR chairs in AI, largely in three centres: Montreal, Toronto and Edmonton, and the affiliated universities in those three cities. Those chairs are to attract AI talent to Canada, to retain existing AI talent and to train the next generation of AI scientists.

The next tranche of money in that \$125 million will be used to fund three rather unique institutes that have been created in those three cities. They

are bifocal in their vision. They are looking inward towards the universities and funding academic research but they are also meant to attract and partner with the private sector. Part of the \$125 million will be used to support the running of those three institutes.

The third aspect will be to develop a national brand and national activities around AI to bring shareholders and the private sector together to develop a coherent view of AI and its impacts on the economy. The fourth component will be to address some of the broader issues of AI and its impacts on society. We are hoping to do that on a global scale.

Baroness Bakewell: Is the Canadian Government considering measures to tackle the growth in inequality that may be brought about by AI—social problems?

Dr Alan Bernstein: I am sure they are. In fact, the term that is used in Ottawa is “inclusive innovation”. By that, the Federal Government means innovations that are a rising tide lifting all boats to the largest extent possible. As you have alluded to already, that has to do with skills development, taxation and education policies, and all kinds of other areas in the economy, not focused directly on AI but more broadly.

The Chairman: Can I add something, Dr Bernstein? Do you believe that Canada needs to think about AI-specific regulations or the creation of an AI-specific watchdog?

Dr Alan Bernstein: To my knowledge, at least, there has not been discussion of an AI-specific watchdog. If you parse that, some of the issues are not necessarily unique to AI. Social scientists here in Canada, the United States and the UK have talked about issues around bias, privacy, consent, non-inclusivity, job loss and the overconcentration of capital and data in a small number of companies. None of those issues per se is unique to AI; they have been brought to the fore, I believe, because of AI. Yes, they all need to be addressed but the public need to trust AI if we are going to use it widely and benefit from its predictive power.

Q176 **Baroness Rock:** Good afternoon. You have very kindly covered some of the queries I was going to ask you about the institutes and the three areas you are focusing on. I want to ask you more about the start-up sector. First, how do you see the start-up in AI happening in your country? Secondly, what support is provided to that sector?

Dr Alan Bernstein: That is an excellent question. It is interesting to see what is happening in Canada because there is a change going on in terms of our start-up and scale-up activity, which I will come back to. We have not had, historically, a good track record of venture capital investments and start-ups. That is for all kinds of structural reasons in our economy. That is now changing. I believe it is changing for a lot of reasons that we could talk about, but it is changing. There is a fair bit of start-up activity in Canada, particularly in the AI field, again, because so much great talent is being produced, largely in those three cities. The challenge for us is not so much start-up but scale-up. The two challenges that I am aware of with scale-up are: first, you need more capital to go from a small company to a medium or large-sized company; and, secondly,

larger firms with very deep pockets, when they see a successful start-up with a successful business model, are inclined to buy out those start-ups. Many of the larger companies, of course, are not based in Canada. The Federal Government has created a venture capital catalyst initiative of \$400 million largely focused on helping companies not to start up but to scale up.

Baroness Rock: That is very helpful, thank you.

Q177 **Lord Giddens:** Good afternoon, if that is still appropriate. How does Canada's tech industry see itself in relation to the US? You have commented on that a bit. Which other countries does Canada's AI sector have strong links with, and why? We have heard evidence from a German leader in the academic field, so you might be able to mention Germany.

Dr Alan Bernstein: We certainly have strong links with the US, as I said, largely for geographic, language and cultural reasons. We also have very strong links with the UK, for similar reasons and, I would add, historic connections. Indeed, you see that in the patterns of investment from companies. Google's DeepMind, which is based, as you know, in the UK, has invested majorly and its largest investment outside the UK is in Edmonton, Alberta. That reflects all kinds of things in terms of personal connections, the very good science going on in Edmonton and our shared history, to some extent. Similarly, we have connections, because of Quebec, with France. We have strong connections with the AI community and the computer science community in France.

Lord Giddens: Presumably, you have a lot of concerns about the sheer size of US companies and their proximity to Canada. Can any kind of strategy be utilised to, at least, protect to some degree against such takeovers and, perhaps, pirating of Canadian talent?

Dr Alan Bernstein: The \$125 million pan-Canadian AI strategy was designed to deal with the issue of, at least in part, as you put it, pirating of Canadian talent by helping to fund Canadian talent and keeping it in Canada. One of the other things we are seeing, though, with the rise—if I can put it this way—of the Canadian brand, is that increasingly Canadians do not want to leave Canada. I have repeated this a number of times but I think it is important. Our cities are very liveable and dynamic, we have a great education system and they are safe. All of a sudden, we, as Canadians, are realising that Mecca is not always south of our border. We have a lot to offer. What is happening now is that a lot of Silicon Valley companies, rather than luring talent to Silicon Valley, are investing in Montreal, Toronto and Edmonton, largely. In the last year, let us say, we have seen Facebook, Google, Google DeepMind, IBM, Microsoft and Uber all make significant investments in those three centres. That is a welcome change for us.

Q178 **Lord Puttnam:** What restrictions does Canada have with respect to the use of personal data? How does this impact on the development and application of AI?

Dr Alan Bernstein: I am not an expert, by any means, on personal data. I know the Federal Government has legislation on the books and there are personal privacy commissioners, both at the federal level and in

almost all the provinces, if not all the provinces, and the Northwest Territories. There are many statutes and acts on the books on the use of data and privacy. Those are generic in the sense that they do not pertain specifically to any one technology but pertain to privacy and consent across the board, including artificial intelligence. Whether or not the Government or Governments are contemplating new statutes to deal specifically with AI, I would not know. I cannot go beyond that. I know that Governments are certainly concerned about privacy and consent around all these technologies. I would include artificial intelligence, genomics and genetic data and a combination of the two. These are important areas for us because they pertain to advances in not just the science but how that science is applied to improve health. It is clear that the public need to trust these technologies and trust the privacy and consent issues around them if they are going to be used in the maximum way.

Lord Puttnam: Can I follow up on that? Canada, clearly, certainly to my mind, is a lot closer to the UK in terms of its value base than to the US. Has the concept of data trusts, or similar arrangements for the safe and responsible sharing of datasets between public and private organisations, been considered in Canada? Let us say, there is a more prosocial approach than the one in the United States.

Dr Alan Bernstein: That is an excellent question. There is a very large open access movement in Canada. I will give a somewhat long answer, so I apologise in advance. There is a major consortium that was first funded by the Wellcome Trust in partnership with federal agencies in Canada, including an agency I used to work for, the Canadian Institutes of Health Research, on solving the three-dimensional structure of proteins. That initiative, which now includes Sweden among other countries, is an open-access, publicly available dataset that transcends public-private. GlaxoSmithKline is also a funder of that initiative. The Montreal Neurological Institute is starting a major open-access initiative. The AI group, to get to your point, especially in Montreal, believes very strongly in repositing all its algorithms and thinking about AI in an open setting. I think that is a value, as you pointed out, that is shared between Canada, the UK and, I must say, many American scientists as well.

Q179 **Baroness Grender:** You have already given us a sense of some international collaboration. Do you think there is a need for a more formalised way of having international collaboration? Do you believe Canada has a role in this? Who is best placed to facilitate this kind of co-operation? Can you talk to us about any other specific, global AI initiatives that Canada is participating in? Do you think there are initiatives that should exist that currently do not?

Dr Alan Bernstein: These questions are near and dear to my heart, as president of CIFAR. The essence of the CIFAR vision and model is that we bring together scientists from around the world to address questions of importance to the world. At CIFAR, we do not believe that any important question is unique to any one country. In AI specifically, our programme in AI, which we call Learning in Machines & Brains, involves scientists from at least half a dozen countries including Canada, the US, the UK,

Japan and France. That programme, which has been going under various names for the last 15 or 20 years, has always been international in its scope and always will be. That is at the core of the CIFAR vision. That is the first thing I would say in answer to your question.

The second thing I would say, and we have touched on this a bit, is that the challenges of AI for society—the economic, privacy and philosophical issues, et cetera—are not unique to any one country. There are many initiatives that are national in scope in the area we call AI in society. At CIFAR, we believe strongly that there needs to be a global approach to those issues. One of the things CIFAR would like to see happen, and we are taking some steps to make it happen, is to bring together, on an international scale, funders from around the world who have a major interest in the impacts of AI on society, so that we can develop a more coherent, integrated discussion around the implications of AI for society. That will be of greater use for policymakers in government and for the public, as opposed to the current situation where a lot of academics and organisations are now writing about various aspects of AI. Of course, a lot of this is about predicting the future. There are some data-driven predictions which should be paid attention to and some that are mere speculation. It is hard for all of us to sift through that. We would like to see a more coherent approach to this.

Thirdly, CIFAR is hoping to start a new programme. About a week ago, we launched our next global call for ideas. We are hoping that one of those new programmes will be in the broad area of AI and society or, more broadly, new technologies and society—again, bringing together academics from around the world. I do not know if I have answered your question. This has to be addressed on an international scale.

The Chairman: Thank you very much, Dr Bernstein. The final question comes from the Lord Bishop.

Q180 **The Lord Bishop of Oxford:** Thank you, Dr Bernstein. We have the responsibility of making a series of recommendations at the end of this inquiry. If there was one recommendation you think our Committee should make, what would it be?

Dr Alan Bernstein: It relates to the last question I was asked. I think the recommendation would be that the UK, perhaps with Canada and CIFAR, should develop an integrated, international approach to the societal, philosophical and ethical implications of AI. These are not issues that are unique to the UK or to Canada, or to the US, France or what have you. We have a lot to learn from each other and should be working together on this. Historically, both our countries have a track record of developing these international initiatives and ensuring that they are effective and accomplish the goals they are meant to accomplish. One of my VPs at CIFAR is going to the UK next week to a Ditchley conference to discuss some of these issues, so we already are establishing strong bonds and relations with colleagues in the UK around these issues. I would like to see them develop further. To me that would be a great recommendation, Lord Bishop, coming out of your Committee.

The Chairman: That is a great note to end on, Dr Bernstein. We all appreciate your evidence today. It has been very interesting. There is a

CIFAR – Oral evidence (QQ 172 – 180)

lot we have in common, in terms of our national approaches, clearly. Thank you very much indeed. We appreciate your time and liked the backdrop to your office.

Dr Alan Bernstein: Thank you very much for your time. I appreciate your interest.

Citizens Advice and Competition and Markets Authority – Oral evidence (QQ 85–94)

Evidence Session No. 10

Heard in Public

Questions 85–94

Tuesday 7 November 2017

[Watch the meeting](#)

Members present: Lord Clement-Jones (The Chairman); Baroness Bakewell; Lord Giddens; Baroness Greender; Lord Hollick; Lord Holmes of Richmond; Lord Levene of Portsoken; The Lord Bishop of Oxford; Viscount Ridley; Baroness Rock; Lord St John of Bletso; Lord Swinfen.

Examination of witnesses

Colin Griffiths and Will Hayter.

Q85 The Chairman: This is the tenth formal evidence session in our inquiry and it is intended to help the Committee consider the impact of artificial intelligence on consumers. We have just had a business-to-business session and this is, if you like, the other side of the coin. Many apologies for the delay in getting cracking, but, as you can see, this Committee easily gets overexcited.

I will go through the ordinary rubric. The session is open to the public. A webcast of the session goes out live and is subsequently accessible via the parliamentary website. A verbatim transcript will be taken of your evidence and this will be put on the parliamentary website. A few days after this evidence session, you will be sent a copy of the transcript to check for accuracy. We would be grateful if you could advise us of any corrections as quickly as possible. If, after this session, you wish to clarify or amplify any points made during your evidence, or have any additional points to make, you are welcome to submit supplementary written evidence to us.

We have before us Colin Griffiths, policy manager, Citizens Advice, and Will Hayter, policy director, the Competition and Markets Authority. Would you like to introduce yourselves for the record, and then we will begin with the questions?

Colin Griffiths: I am a policy manager at Citizens Advice and the policy lead on smart metering, smart homes, digitally derived services and, increasingly, the internet of things.

Will Hayter: I am a project director at the Competition and Markets Authority. Most recently I was responsible for our digital comparison tools

market study. Our mission is to make markets work well for consumers, business and the economy. I also have a strong interest in this subject.

Q86 The Chairman: I am going to start off with a big question. What are the biggest advantages and disadvantages that AI could bring to consumers over the next 10 years? You might want to put that in the context of the Government's recent AI review by Dame Wendy Hall, and whether it touches upon the issues you think it should have touched upon.

Colin Griffiths: You can probably break down the advantages and disadvantages and, to some extent, categorise them by things that AI enables consumers to do, and things that it will be used to do to consumers. The disadvantages tend to fall in the second camp and the advantages into the first. The advantages have already been discussed in earlier sessions. New services are emerging that will help consumers to navigate increasingly complex markets and help them make better, more informed decisions. There is improved accessibility. There is a lot of work on chatbots, in language translation and in similar systems.

AI allows increasingly tailored services, which consumers like. Because you have more data, you are able to tailor a service much more specifically to a consumer. In the energy world, energy efficiency advice is particularly good when you have a smart home. It can look at exactly how you live your life and rather than giving very generic, occasionally not useful information, it can give you something specific. Finally, one would hope that the benefits of the increase in efficiency and productivity would be passed down to consumers and not just held by the industries which get them. There are the broader potential benefits, which others have talked about: the impact on healthcare and various new ways that data can be used to improve the day-to-day lives of consumers.

The disadvantages—and this is where things are increasingly done to consumers—are if it disrupts markets in ways that can disadvantage consumers. Insurance is a good example. Insurers are always keen to have more data about consumers, because it helps them more carefully analyse and target their products and services. There are always going to be winners and losers in that situation. Insurance entails pooling risk, when data becomes more precise, for every consumer who gets a much better deal on their insurance, there is going to be someone who may get a worse deal or be excluded. The same will be true for financial services as they start to merge credit ratings into datasets about how you live your life. I was intrigued to note that some people I have spoken to already deliberately make their junk food purchases with cash and their gym membership on card, such that in the data world they look as if they are a better prospect for insurance. That is happening now and people are starting to make that decision at the most engaged end.

The Chairman: That is very savvy.

Colin Griffiths: There is also a particular risk for the digitally excluded, because it seems likely that a lot of these services will be offered online and through digital services. If you are not online, you will still be subject to all the AI analysis that happens with your insurance company or financial services provider, but you might not be able to take advantage

of those new services that are appearing elsewhere that could help you be a more empowered consumer.

Finally, it is a challenge for regulation, not an insurmountable one but it is a challenge. The regulators are quite well built for evaluating simple pricing strategies, but once you get into the world of a black box of AI, and no one quite knows how it reaches decisions, regulators have their work cut out for them in evaluating that. It is worth noting that the risks and benefits are not mutually exclusive; consumers will both benefit and be at risk.

The Chairman: That is a formidable balance sheet. Do you agree with that, Will?

Will Hayter: Yes, with large parts of it. First, what we can say with certainty is that AI and data and so on will underpin new services. I am probably not the first person to mention that it is 10 years since the iPhone came on to the market, and 10 years ago who would have predicted that we would have had a computer in our pockets—all of us in this room probably—that is capable of doing all the amazing things it is able to do? It would have been very hard to predict that back then.

Others have talked about bias, security and privacy issues. Focusing on competition and consumer issues, an overall point is that in markets that work well, services that make firms money are likely to be the ones that consumers want. That is how competition works when it works well. I am going to focus on a particular angle touched on by Colin earlier, which is navigating complex markets, and will try to think about a pessimistic and an optimistic scenario. The pessimistic scenario is that the technology makes things difficult to navigate and makes the market more opaque, and perhaps consumers lose trust and disengage from markets. The more optimistic scenario is that the technology is able to work for consumers.

The Chairman: Price comparison apps and things of that nature?

Will Hayter: As an example. To try to make that a bit more concrete, because it is a bit of an abstract idea, imagine a scenario where there is a company that will, with your authorisation, take on your bank statement and credit card statement. It might look at all the ways you spend your money, suggest a number of ways you might like to switch your household bills to alternative providers and perhaps monitor the market to keep track of who might give you a better deal. It might suggest that you get rid of your travel insurance because you already have it through your bank. It might suggest you unsubscribe from Amazon Prime because all the TV content you see there you already have via Netflix, or whatever the combinations might be. You might book flights, et cetera. It might do all those things at a lower cost than a human-powered version.

It is important to note, of course, that consumers vary a lot so you cannot talk about them as a monolithic block, and, as Colin has mentioned, there are those who are less engaged and those who are more engaged. Again, there is a pessimistic and an optimistic scenario. The pessimistic scenario is that the disengaged lose out. The optimistic scenario is that some of these tools are able to work for the more vulnerable and, for example, help us to tackle our own behavioural biases. We will always value the short term over the long term and we

tend to be averse to losses more than we value gains. Algorithms, AI and the tech might be able to help us to counteract our own tendencies.

The Chairman: I am going to bring in Lord Levene in a minute, but I want to ask a question. Obviously, the Government's AI review did not really cover the consumer area. Do you feel there is any stirring in the forest at all? Are the Government aware of the consumer implications at this stage and do they have any kind of strategy or thoughts about strategy?

Colin Griffiths: I believe there is but it is quite nascent. Exercises such as this inquiry are very useful because this is all still to come. There have certainly been areas—and it might be touched on later—that I have noticed. Recently, for example, the chief executive of Ofgem started talking about consumers owning their data and saying that that this is vital to competition. We are starting to enter a world where an energy regulator is talking about that, which you would not automatically think of. My impression is that it is beginning, and I think the Government are taking it seriously. I am very keen that they take it seriously, because, as I am sure has already been made clear, this is a very fast-moving area where there is a lot of change. It is particularly complex and opaque and crosses sectors very rapidly. I am very keen that government engages and is not as siloed as it has been.

Will Hayter: I would add that as an independent authority we do not tend to speak for the Government.

The Chairman: That is why you are here.

Will Hayter: Indeed. What I can say is that in talking to colleagues and other authorities in this area, this is very high up on everyone's list. We are all conscious of the need to build on our existing understanding of some of these issues and bring in additional expertise.

Lord Levene of Portsoken: Colin, would you expand a little on what you said about insurance? If I understood you correctly, the people who understood it get a better deal to the detriment of others. In other words, it is a sort of zero-sum game. Could you expand on how you see that?

Colin Griffiths: My specific example about gym memberships was a bit glib.

Lord Levene of Portsoken: I was talking about insurance, not about gym memberships.

Colin Griffiths: But that was relevant to financial datasets being used to analyse people for their insurance. You are gathering how someone lives their life from one dataset and applying that to a different service. The concern is that, at the moment, with a product such as insurance, you pool risk for a lot of consumers, and, effectively, there is a cross-subsidy, or however you want to phrase that, but the more precise the data is about each individual consumer, and the more tailored an insurance product becomes, the more discrepancy there is going to be in cost. You can argue that it is reflective, but it also significantly changes that market, and there are some consumers who are likely to be losers and some who are likely to be winners.

Baroness Bakewell: Do you think it is going to disadvantage the old, because this technology, as you say, is moving fast? It is very quick and becomes very glib. As you age—and we are all going to live to be 100—you are less able and less adaptable to new technology. We have a demographic in which the values of AI will decline as people get older.

Colin Griffiths: I would say that is definitely a risk. There are ways to mitigate it through making these products more accessible, but, yes, that is absolutely a risk, especially as these things start to moving into services such as analysing your Amazon versus Netflix consumption—an example Will gave earlier—and finding that you can cancel one of the accounts. That is a service that will almost always, I should imagine, be offered online or via an app. If you do not have access to those, that is not a service you will be making use of, so I think it is a risk.

Baroness Bakewell: AI might help with the care of older people, but how are they to be kept abreast of the technology? There is a problem to solve there, is there not?

Colin Griffiths: There is. There is always a role for customer education and making them more informed. I am also somewhat wary of that. I appreciate that I am saying this as a representative of an advice agency, and we advise millions of consumers all the time, but there is a tendency within industry more broadly to say, “The solution to this problem is to educate consumers and to make consumers think more correctly about this”. Especially when you get into really complex markets, and when so many of these products are sold on convenience and making life easier, they also come with the line, “You are going to have to understand some really byzantine business models”, and there is a limit to where consumer education can get you. A more top-down approach is needed to make sure that there are basic standards for these things.

The Chairman: We will explore that a little further later.

Q87 **Lord Levene of Portsoken:** How aware are consumers of the role of AI in products that they use? Is it important for consumers to be aware of AI and its implications?

Colin Griffiths: As an organisation that gets a lot of consumer evidence from millions of consumers coming to us, the short answer is that they are not particularly aware, but it is a very difficult question to poll, because most people have heard of the concept of AI, but the boundaries between a sci-fi conception and what is really happening are quite narrow. The more you dig into a consumer’s understanding, the shallower it will tend to be. Consumers are increasingly aware of the use of algorithms and that they’re in a world where, “I have to interact with my computer. Flight tickets seem to change price because I looked at them before”, and people are getting increasingly suspicious about that sort of thing.

The other thing that has been very striking to me over the last few years is a significant increase in consumers being aware that their data has a value. I was particularly struck in a contact we had recently about a smart meter, where someone wanted their smart meter removed and to be reimbursed for the value of the data it had collected while it was on

their wall. That is starting to creep in as people know that their data has a value to someone, somewhere along the way, even though, as I have mentioned already, a lot of these business models are particularly confusing to people.

There is an increasing amount of asymmetric information on this. When we talk to consumers and do research, we often find that people are comfortable with the idea of a shop and they understand that it buys in bulk, it increases the price slightly—that is its margin—and that is the shop's business model. When they are interacting with a large company, for example Google, they think, "I am getting free email and free maps and this service has pretty much replaced my satnav and that appears to be free", and they are somewhat aware that Google are selling adverts and seems very wealthy. However, there is a much bigger gap in people understanding how on earth this company is making money. They know it is making money from them—they are kind of aware of that—but the short answer to your question is that they not very aware of AI specifically, or at least do not fully understand it.

Lord Levene of Portsoken: Does it matter?

Colin Griffiths: This touches on the point I made earlier about consumer education and how that has a key role. Consumers need to be up to speed on this and aware of the implications, but I am also reticent about putting all the onus on consumers by saying, "It was your job to understand".

The Chairman: Quite. Should they be told that AI is integrated into products of whatever kind?

Colin Griffiths: Yes, some information that helps deal with that opacity of how you ended up where you are, and how you got the price you received, would be useful, yes.

Will Hayter: I would echo much of that. The important point is that consumers are generally informed in the markets where they are shopping for products. That is both for them individually to get the right answer and to ensure that the right kind of competitive pressure is kept up on the suppliers that are offering those services. In particular, it is critical that they are clear on what they are getting, and the terms on which they are getting it, to support overall trust and confidence in the market.

To give an example of a live issue for us at the moment, we have just opened a consumer investigation into online hotel booking sites, which is all about whether the rankings that are presented on those sites are clear.

The Chairman: Whether they are bumping up their own ratings and so on.

Will Hayter: Whether the basis for that is clear, exactly. It is also important to bear in mind that it is not necessary for consumers to know everything about what is going on behind the scenes. Using the analogy of a supermarket, most people probably have a general idea of how a supermarket works. That does not mean that they have to know every detail of the toing and froing between the supermarket and its suppliers,

for example, as long as they are generally confident about what they are getting and the terms on which they are getting it.

The Chairman: We potentially have a trust issue, do we not? Is there a role for a kite-marking scheme or something of that sort, on a voluntary basis perhaps?

Will Hayter: Kite-marking schemes can be powerful if they have really big brand recognition and consumers understand what they are there for. This is slightly straying outside our area of expertise, but there is an AdSense logo, I believe, which is supposed to identify tailored ads. It is a small blue triangle on ad results on websites. As a consumer, before I learned about this in my professional capacity, I had no idea what that was or that it was even there. You may get firms signing up to these schemes, but unless consumers understand what they are there for and what they are communicating, their value is more questionable.

Colin Griffiths: One point I probably should have made more strongly—Will picked it up—is that you can help quite significantly here by consumers knowing their rights and that there is still accountability here. That is the other concern that consumers often have about this. It starts to feel a lot more anonymous and harder to engage with and understand. We already have issues with consumers wanting to know, “How do I complain? How do I get redress? What is my route?” In a world where there is a black box of AI and a company can more plausibly say, “We don’t really know. The machine made the decision”, or indeed, “We contracted that out to a separate AI company”, it is particularly vital to ensure that there is always a very clear route for any consumer who has a problem to get redress and have complaints dealt with. That is more about being aware of your rights and what you are entitled to, for want of a better phrase.

The Chairman: Is there any obvious route currently on a complaints scheme? Is the ICO the obvious place now or would you say that is a role that needs to be developed?

Colin Griffiths: I certainly think the ICO will need to keep up with AI. You can make complaints through the ICO if there is a data privacy breach—and I am sure the ICO is doing this as well—but as more and more services start having this threaded through them, it is going to become more important that you factor it in. I apologise that I have not given you a really neat answer to that. I feel as if I should have the complete solution.

The Chairman: We will keep coming back at you.

Q88 **Viscount Ridley:** As background to a question about consumer concerns, one of our previous witnesses, Sage, said in its written evidence, “Rhetoric in the media is largely negative”, and, “The biggest threat we face right now is not an existential risk to humanity or widespread job loss”. What are the biggest concerns related to the use of AI that you hear about from consumers? Are consumers already too spooked by this technology?

Colin Griffiths: I will open by saying that although we hear from a lot of consumers, it is very rare that someone comes to us and says

specifically, “I am worried about AI”. I am extrapolating somewhat from consumers who are worried about data-driven services or what these companies are doing. Probably the most commonly expressed sentiment is the jobs question of who this is going to put out of work. Outside of employment concerns, accountability is one of the most common, because it feels like it is being taken away and is happening to you. It comes back to the distinction I made right at the start about things being done to you. The second a consumer feels that something is being done to them, they generally dislike it. Innately it makes them feel, “It must be something I wouldn’t have chosen to do”. There is a general unease.

One striking point on data-driven services is there is a tendency—and I hear this a lot—for people to say, “People don’t really care about this. They put everything on Facebook. They don’t really mind”, but when you talk to consumers, even those who are using all these services, there is quite a strong unease that comes through when you dig into it. “I am using this but it was a kind of all-or-nothing offer”. “I wanted to be on this service because everyone else was and it seemed useful”.

In some research that we did, one consumer summed it up quite nicely when we were going through the customer journey of signing up to digital products. He got to the terms and conditions, which were in .5 font and very complex, and he immediately said, “They know I am not going to read this. They know that if I do read it, I won’t understand it. Whatever this is about is protecting them, so I’ll click ‘yes’ because I want the thing”. Every time that happens, consumers become a little more nervous and wary. Media drives this to some extent, as does past experience of companies, which might replicate existing concerns about junk mail. Now the junk mail is digital, and there is the feeling, “They know what I am doing in my home”. To some extent, people intuit where it could lead without having a really solid grasp of what it is.

Will Hayter: We tend to hear concerns about specific products. It is interesting that you mention rhetoric versus what is going on in practice. In our recent project on digital comparison tools, the rhetoric in the media was very negative. There was unease about what these firms were doing. However, when we asked consumers, for the most part they had reasonably good feelings about these sites and thought they were getting a good result. Some 90%-plus of people thought they were well served by those firms. The biggest issue of mistrust was with the use of data, which we have already touched upon and we are very alive to.

Viscount Ridley: As a veteran of various genetic debates, I am aware of two possible ways this could go: first, the genetically modified food way, which went highly negative very quickly and stayed that way; and, secondly, genetics in humans, which, despite several people’s attempts, never bothered people that much as they could see that the benefits outweighed the risks. It is quite interesting to think of that analogy. Specifically, what do you think can be done about these concerns?

Will Hayter: The answer to that depends on the specific market, the specific context and the specific product. As Colin said, consumers would not necessarily be able to articulate what they feel about artificial intelligence. I would probably struggle as a consumer. It is important for

us to look very carefully at specific examples and get the evidence about them.

I mentioned our hotels case as a particular point. We are in the process of drawing in specific views about that sector. There are forms of artificial intelligence that support that business model in how the ranking is done and how the reviews are generated and checked. We are interested in whether consumers are being misled in the specific context of when they are booking hotels.

Viscount Ridley: It is case by case and context-specific.

Will Hayter: Exactly.

Colin Griffiths: There is also something to be said for early experiences being quite crucial in this world, because, increasingly, products are coming on to the market and people are making use of them, and you tend to see a spike. To use smart metering as an example, “Watchdog” ran a series of programmes on concerns about the safety of smart meters, and there was a significant spike in the number of people cancelling smart meter installations with their energy suppliers. People tend to react, and it often feels, especially in the area of data, as if it is not that big of an issue until it is, and suddenly there is a big rush to it. It makes a stronger case for making sure that at least some core principles exist at the start that stop anything particularly egregious happening that could eliminate all the excellent services coming through that could be really helpful in the future. That provides a case for making sure that those early experiences are better and that some of these questions have been thought through rather than everything being reactive, where you wait for a scandal to happen and then you work out how to handle it.

Viscount Ridley: That is slightly different from what Mr Hayter was saying about it being case by case.

Lord Levene of Portsoken: Can I go back to what you were saying about people wanting something so badly that they will just tick the box at the end of the day? What, if anything, is being done, or what can you do, to force companies to condense that into a very short set of sentences that have to be read? Some of them are pages and pages long. There is no way anybody could ever get through it, or even understand it. Is anything being done about that?

Will Hayter: That depends on the specific kinds of terms that you are talking about. If you are talking about data protection, clearly the ICO is in that territory. If it is more about general terms and conditions, we and other enforcers are there to make sure that the terms and conditions that are given to consumers are fair.

The Chairman: On the web, the Ts & Cs for applications in particular are notoriously lengthy, complicated and opaque, yet one is required to tick them to move on to the actual product, whatever it may be. Lord Levene has made a very important point: most consumers do not have a clue, do they?

Will Hayter: It is a genuinely difficult issue, because often the terms that lie behind these services can be really complicated. The digital comparison tools study illustrates that. Data exchanges that are going on

behind the scenes are quite complicated. If you try to get an insurance quote from a comparison site, it is getting and exchanging data with 40-plus insurance companies, and that is an intrinsic part of the service that is being offered. It is quite difficult for consumers to understand that. On that specific example, we have pushed firms to explain more clearly to try to improve the perceptions, and therefore the trust, of what is going on there.

If you go back to my optimistic scenario at the beginning, there is a very interesting paper by an academic called Michal S Gal, entitled *Algorithmic Consumers*. The optimistic version of this is that in future somebody's algorithms and the tech can be brought to bear to help consumers navigate those complicated terms and conditions, but, clearly, we are not there yet.

The Chairman: You called them algorithmic consumers.

Will Hayter: That is the name of the paper.

The Chairman: I wondered whether we had already identified a cohort of algorithmic consumers. I want to be one of the first.

Q89 **Lord Hollick:** The digital world is populated by some very large companies and a lot of very small companies. In the context of making sure that we develop AI and it becomes a national asset, the reliance of AI on those large companies, and the fact that they have bottomless pockets of unpaid tax dollars, means that they can outgun any other company and can outbid any other company for talent and for other businesses. How does the Competition and Markets Authority look at that? What remedies can you recommend to address what could be a showstopper for UK-owned AI development in the UK?

Will Hayter: The first thing to say is that a large company, or a company with a powerful position, is not necessarily a problem. The problems come if that position is abused, and that is where the Competition Act comes in.

Lord Hollick: I propose that it could easily be abused.

Baroness Bakewell: That is quite likely.

Will Hayter: If it is abused, we have flexible tools to tackle it. One part of the Competition Act is about the abuse of dominance. To use an offline example, we have just fined some pharma companies £90 million for excessive pricing where they had a dominant market position. We also used our markets regime system in our energy markets and banking market investigations. Our Open Banking remedy, which is allowing consumers to take control of their data—for example, of their bank account transactions—to support them in getting better, more tailored services, to take their information to a small start-up and allow them to provide a service, is exactly what that was directed at. It is supporting competition between start-ups and the big banks. The important thing is not necessarily the size of the company itself but that the position can be confronted by a small start-up. Interestingly, in the study that we have just done on comparison tools, a lot of the firms that are trying to offer the newer services that I described earlier—digital concierges or

automatic switching—are smaller start-ups, but it is important that the data that they are relying on to provide those services is available.

Lord Hollick: So you are not concerned that the dominant position of, say, a Google, which has 80% of the search market, or an Apple, will stifle the growth of the AI in the UK?

Will Hayter: Again, not in and of itself. The recent Google case, which is a European Commission case, of course, rather than ours, where Google was fined for leveraging its search position into the comparison shopping services market, is an example of how to use the flexible tool of the Competition Act—or in that case Article 102 of the Treaty—to take on that abuse of a dominant position.

Lord Hollick: Are you monitoring this space carefully to make sure that competition is not being stifled and that start-ups are not being crushed by having a single counterparty to go to?

Will Hayter: That is exactly the sort of thing that we do.

Lord Hollick: Are you doing that on a proactive basis, or are you waiting for companies to come and complain?

Will Hayter: It is a mixture. We prioritise our resources based on a whole range of inputs. To take a couple of examples, the banking market investigation came to us as a reference from the Financial Conduct Authority. We receive complaints as and when they come. We monitor carefully for cartels, for example. We have the whole leniency application process to encourage people to come forward and report cartels.

On the flipside, the digital comparison tools market study and the recent hotels case are both own-initiative pieces, where we have felt that was an important thing to look at, so we have gone out and done so, gathered views from stakeholders and acted as a result. It is a mixture to try to use our resources efficiently.

Lord Hollick: Would it be fair to say that, as things stand at the moment, you see no evidence that these very large companies are in any way distorting competition and start-ups, but you are keeping a close watch on it?

Will Hayter: We are very much keeping a close watch on it.

The Chairman: On the basis that you are doing that, clearly data is really important in this context. Whether it is a competition matter or a merger matter, do you consult the Information Commissioner when it concerns data? Is there a relationship between the two regulators in that sense?

Will Hayter: There certainly is. It is becoming increasingly important, and that relationship is key to making sure that we look at these things in the right way. We are in the process of improving that relationship and drawing us closer together. There was a strong data protection element in the digital comparison tools piece. We set out four principles for these services: clear, accurate, responsible and easy to use. The “responsible” part was focused on the data protection piece, given the importance of trust in the use of people’s data that we found in that sector. Throughout that piece of work we consulted with the ICO.

The Chairman: I suspect that when there are mergers between organisations that both hold significant datasets, that will increasingly be a matter that you will want talk to the Information Commissioner about.

Will Hayter: Absolutely. What is interesting to note is that most of the merger decisions that deal with the agglomeration of two big datasets have so far been taken at the European level, so we have not thus far dealt with that very specific example, but clearly that seems quite likely in the future.

The Chairman: That is coming down the track.

Will Hayter: Absolutely.

The Chairman: We got so bound up with the competition issue. Colin.

Colin Griffiths: A particular concern on the competition point is where a company can build a closed system. Looking at the smart home, for example, the system of analytics about your heating, your lighting and how you live your life improves gradually over months and years as it learns about the habits of you and your family and where you are. I had a meeting with a smart home company not that long ago where I asked, "What's the process for switching away from you to a new service provider? The product you are selling is the analytics, but can I take the raw-usage data that you have collected and move that to another supplier?" The answer was, "No, we'll delete it". Effectively, that means that it can start locking in a consumer. It might have a better or cheaper product available elsewhere, but a consumer does not want to start from scratch with years of analytics about what you do to learn your preferences again, even if they think the new product might be better.

That is where the GDPR is quite welcome in the data portability recommendations that it makes, because that helps consumers switch. One of our biggest concerns about the competition element is that if someone can build a closed system they can kick the ladder away, for want of a better phrase, and continually lock consumers into their ecosystem, or however they want to brand it.

The Chairman: That is really interesting. Is that ability to pass data between suppliers at the consumers' behest?

Colin Griffiths: Yes, and we have worked on that. I was involved in the MiData programme. The ability to easily switch has always been a key principle. Consumers do not come to us and talk in quite those technical terms. They say, "I want to be able to switch and I don't see why I can't move this". There is a desire for that, a desire that generally only emerges at the point of them switching and realising that they cannot. It is not something that people see ahead of time.

Q90 **Lord Giddens:** You may feel that the question I was delegated to ask has been answered, which was: are consumers concerned about how their personal data is used? The answer to that is "not very much", so, with the Chairman's agreement, I might amend it slightly. Children are consumers. What specific issues arise in respect of children and young people where you cannot simply say they are orthodox consumers? You might want to comment on that. Do you see a role for the protection of

data through blockchain technologies of some sort that could be bracketed on to consumer rights?

Colin Griffiths: The answer to the initial question is yes, and we get contacts about this. It is interesting that you mention children, because one thing I found most striking—and this comes back to the point I made earlier: that people are sometimes quite dismissive of consumers' concerns—is that industry often mistakes using a service for being happy with the service. The fact that you are on Facebook must mean that you do not care that you have accidentally or deliberately told everyone where you are and when you are going on holiday and things like that. There is a consent question there that people get quite muddled about: they are happily using it and therefore it is fine. There is this strong sense of unease that I mentioned earlier.

What is really striking is that if you talk to consumers about their children, they are suddenly much more concerned. "No, I wouldn't want my child agreeing to those terms". "I wouldn't want my child on Facebook or Instagram, or, if they were, I would want it really locked down". I do not think it is putting it too strongly to say that some of those people feel guilty about how much they put out there. They say, "I know I shouldn't have done that, but I don't really understand it, and they'll get my data somehow anyway". It is very striking what happens the second you bring children into it. There is an imperfect analogy with junk food. People will eat junk food but will not let their children eat it. It is a similar thing with a lot of digital services. People are much more concerned there. A useful way to gauge people's true views on this is to ask them whether they would let their children do that.

Lord Giddens: A surprising number of people around the world give young infants, almost newborns, smartphones to play with. You can buy a potty with a stand for an iPhone.

Colin Griffiths: I believe it is called an iPotty.

Lord Giddens: It is a serious issue. Children are not the same as other consumers.

The Chairman: It gives a new meaning to "potty training", does it not?

Lord Giddens: It leads to a life of sitting on the loo and looking at iPads.

Colin Griffiths: As a parent of a young child, there is also a valid point to make—it is similar to the terms and conditions conversation that we had—about the moment when, having decided that you want the thing, you are presented with a little box and you go, "Yes, that's fine". I can understand the idea, as someone who does not own an iPotty, that, "This makes my life easier in the now, and I will consider the implications later". It is that same decision-making. I was not expecting this to be quite such a potty-centric discussion, but there is a valid point to make about immediate convenience over the long-term implications.

The Chairman: You can always guarantee that Lord Giddens is going to get to the heart of the matter.

Lord Giddens: The root of the matter.

The Chairman: Thank you very much. And on blockchain?

Will Hayter: I do not profess to be a technology expert, but, again, blockchain is a technology that we are conscious of in a broad way. We hear plenty of things about its ability to do things securely and to handle data. It sounds as though it has the potential to be an enabler for some of the broader concepts of privacy and trust in markets and so on that we feel so passionate about. Quite what that means in practice we will have to wait and see.

The Chairman: Eventually, engaging in best practice or whatever towards consumers may become a standard that people need to adhere to.

Will Hayter: Potentially, yes.

The Chairman: Without putting words into your mouth, Will. Let us move on to Baroness Bakewell's question.

Q91 **Baroness Bakewell:** How can data best be managed in the interests of the public? I use the term "managed" because quite a lot of people in their written submissions have said the Government have to do more, the Government have to do plenty. How do you feel about that?

Will Hayter: You will have heard lots of angles to this debate—important points about privacy and so on. I will not focus on those.

From a competition point of view, which Colin has touched on already, there is a point to make about the opportunities arising from the management of data. Quite rightly, much of the debate tends to focus on the threats to privacy and data protection. If we only perceive a threat and we never share any data, we do not get the benefit from all the services that AI and data can support. It is worth thinking about how the opportunities for data and AI can be made to work for consumers.

Going back to my optimistic scenario about the algorithmic consumer, et cetera, a lot has to happen in practice for that to work well, and a lot of that is about the management of data. I have already touched on our Open Banking remedy, which was all about freeing up your data, with the right controls and systems in place to do so, so that companies can use that data to offer you better services and more choice. If you take that through the energy sector and smart meters, or the mobile sector and your usage history, your transaction or viewing history on various different platforms, the algorithms need to be able to work smoothly on your behalf and pull that data from one part of the cloud to another to offer those services.

The principle of portability is in the GDPR already. In practice, there is a lot behind that in the interoperability of different systems and common standards for data; how it will happen that you will issue a command to a firm to get hold of your data and use it in the right way.

Baroness Bakewell: What do you think about the Royal Society's proposals for a data stewardship body? What do you think about data trusts generally?

Colin Griffiths: I am thinking about consumer outcomes and how things should ultimately work for consumers. Whenever we talk to consumers and we do research—and we also get lots of contacts in with people

concerned about this—the two themes that always emerge across different demographics and opinions on all this, from people who are resigned to people who are very worried about it, are transparency and control, which tend to be the two key words that we use. Transparency means giving consumers the ability to see where their data is going and what it is being used for and who by. Even if this is not information that consumers will use immediately, even if they are just signing up, just knowing that they have that ability to check in the future is quite consistently of great value to them.

The second aspect is control, which is the ability to make decisions to change options. This touches again on the “you are in or you are out” model, which very much does not allow for that. Options might take the form of trade-offs. We have tended to advocate for a world in which, because so many of these services are beneficial, a provider should be able to come to you and say, “In exchange for your data, I will offer you this”. The consumer should be in a position to say yes or no and to tailor it. Some companies are increasingly experimenting with sliding bars, in a sense: “If you give us your location data, we can provide you with this”, or, “If you give us this we can provide that”, rather than, “This is a flashlight app for your phone and we need location data”.

Baroness Bakewell: Should all this be groomed by the Government and by legislation?

Colin Griffiths: Those principles should be, yes. The core principles that ensure transparency and control should be put into place so that people have those. Achieving them and framing it correctly such that you do not completely eliminate products that would have been really useful is where the delicate balance is made. At the heart of any decision-making or policy-making on this, consideration should always be given to consumers having transparency and control.

Baroness Bakewell: Do you agree?

Will Hayter: Yes. A lot of that in principle is in the GDPR and implementation of it.

Baroness Bakewell: Should it be pushed further? Is the GDPR adequate?

Will Hayter: It is more a question of principle versus practice. The principles sound good. We touched on the portability principle, which is great. As the consumer, I ought to be able to move my data from one place to another or give someone else access to it to support services for me. There are a lot of questions about what that means in practice. Going back to what I said, let us try to think about the opportunities there as well as the threats.

Q92 **Lord Swinfen:** When artificial intelligence systems or the algorithms that underpin them malfunction, or otherwise make erroneous decisions that cause individuals harm, do new mechanisms for legal liability in redress need to be considered, or is the status quo sufficient? If new mechanisms are needed, what form should they take? How should victims be compensated?

Colin Griffiths: To some extent, the answer to this hinges on a point I made earlier about regulators being considerably more joined up. The issue with AI is that it cuts across so many different sectors. I have talked already about several of the key principles that you need to adhere to. They include consumer ownership of data, so they have that transparency and control, clear routes for redress, and trying to avoid the one-step-removed issue with AI where there is often a black box. Another key element that will be needed—and I do not know whether you need new structures to do this or whether you can deal with it within current structures—is a commitment from providers to an expectation of how a service will work rather than just offering things as they are.

To use a specific example, and going back to children, a toy manufacturer recently made IoT-enabled toys with microphones, cameras and all sorts of things that were taking quite a lot of data of young children. Their devices were hacked very extensively and were completely open for anyone to access. With some baby monitors, people could even talk to children at the other end. It is a particularly troubling example. The response from one particular large company, which I made a note of to make sure I did not misquote it, was to add to their terms and conditions—there are the terms and conditions, again—“You acknowledge and agree that any information you send or receive may not be secure and may be intercepted or later acquired by unauthorised parties. You acknowledge and agree that your use of any software or firmware is at your own risk”. That was their approach to dealing with that.

That seems entirely unacceptable as a way of doing business. My personal inclination is that if you are doing that, you should have to put it on the toy box in big letters like a cigarette warning, and that might make people take that a bit more seriously as a risk of the product. I have not entirely answered your question about whether it needs something. To some extent, there is an element of a watching brief. If regulators get up to speed, if they work together, if they acknowledge that this cuts across all their areas, it could be done under current redress systems but, if not, there should be scope for more action.

The Chairman: What do you mean by the current redress systems?

Colin Griffiths: The current redress systems that work in this market. To use energy as an example, we are in a world where Ofgem, the data regulator, now has to sign off data plans from energy networks because they have extensive privacy impact documents and they go to their regulator, which is the energy regulator, to get those signed off. It is going to be a case of Ofgem, which was not built with detailed data analytics in mind, getting up to speed very quickly. That is going to be true across the board for regulators.

Will Hayter: It is definitely an important question to consider. Our answer tends to be a bit of a watching brief, because there are many ways in which consumers can be harmed, and many of those ways exist at the moment. The question is: are there new types of harm or is it more the case that the kinds of harm that we are talking about are brought about by AI? To illustrate what I am talking about, there is already legislation against discrimination. AI is one way of discriminating.

There is legislation in consumer protection regulations to prevent companies misleading consumers. AI is one way of facilitating a consumer being misled. The hotels case that we have just opened is about technology-based companies potentially misleading consumers. There are tools to deal with that at a principles-based level, whether it is supported by AI or not.

One example in the competition space, which is a big topic for debate, is the potential for AI and algorithms to be used to collude. In that case, we have found so far that, again, general principles apply. You might not necessarily have a smoke-filled room any more, but an email from one firm to another saying, “Let’s agree to fix prices”, even if an algorithm is being used to make that happen, is still susceptible to the same principle-based legislation, as we found in a case where two companies were fixing prices for posters on Amazon Marketplace, for example.

The Chairman: And they were using an algorithm?

Lord Swinfen: Some algorithms can mutate. Does this not change the legal position?

Will Hayter: You have hit exactly on the nub of that particular debate on algorithmic collusion. Again, the pessimists in this scenario might say that this is a completely new way of colluding, so there is no way for the current legislation to tackle that. We think the jury is still out. So far, we have not seen these kinds of algorithmic mutations leading to problems that cannot be tackled under the existing law. In the posters example that I mentioned, there was an email chain that we could follow, in which two firms agreed to collude, and they just enacted it through an algorithm. Depending on who you talk to, there is still quite a lot of scepticism that even the most sophisticated deep-learning algorithm can really learn to collude, because there are still many of the same incentives at play for someone to duck the agreement, or the tacit agreement, and reduce prices and compete for more market share.

The Chairman: That is fascinating.

Q93 **Lord St John of Bletso:** Colin, you mentioned that the consumers you had spoken to did not see AI as a threat. Do either of you believe that we need a specific watchdog or regulator to protect consumers with respect to AI-based products and services?

Colin Griffiths: I might have addressed this somewhat in my previous answer when I talked about regulators needing to be much more on the ball and joined up and not operating in the distinct silo of their specific industry. Data is making all those industries join up, and an energy network is now a data network. The first recommendation in a recent report, *ReShaping Regulation*, was—I will read it so that I do not get it wrong—“Regulate for how consumers consume not how businesses are organised, reflecting the blurring and bundling of products and services”. That is quite a solid recommendation on regulation. This does not mean completely ruling out an AI regulator, but because that AI regulator would have to be so cross-cutting and in every industry it would be quite a different-looking beast than an industry regulator, which is how we

currently structure it. Either way, it would look quite different and might not fit our current conception of what regulation looks like.

Will Hayter: You heard from the previous witnesses that AI is in every market in some way; data is in every market in some way. Any regulator in any sector, or a cross-sector competition authority such as us, has to be thinking about AI and data already. I mentioned the digital comparison tools piece from our recent turf, as it were. We did a piece of work on online reviews and endorsements using our consumer powers. That, again, is a market where AI and algorithms can be used to check reviews as they are put online, for example, or to generate false reviews for that matter. Again, the same principles apply. I mentioned our Open Banking initiative. There are bots in the event-ticketing and hotels markets. I mentioned the posters collusion case, and Colin talked about the connected toys issue. Even for us as one authority, lots of the things that we are already doing are informed by an understanding of AI. We have to think about how AI tech data is driving the markets that we are looking at.

It is worth making the general point that we should be cautious about jumping to new regulation and new regulators for fear of holding back the innovation that drives a lot of these new services. That is not to say, of course, that nothing needs to change. I have mentioned this already, and we are very conscious of the need to beef up our understanding of this, to build on the cases we have already looked at and to draw in new expertise, and, indeed, to keep talking to other authorities, as Colin mentioned. As an example, the comparison tools study involved the coming together of most of the economic regulators—the FCA, Ofgem, Ofcom, the CAA, Ofwat. We talked to all of them as well as to the ICO, the Advertising Standards Agency, et cetera, throughout that work. It is important that we all think together, particularly if and when people start buying products in a bundled way across services.

Q94 **Baroness Greender:** We talked earlier about how valuable an exercise such as the one this Committee is conducting is, but we only really add value if we come up with great recommendations at the end. This is your chance to tell us what you think. You have only one each; you have heard how strict the Lord Chairman is on this. What do you think we should recommend? What is the one key thing?

Colin Griffiths: My take home, would be to keep consumer needs at the heart of any of these products as they are developed. There is a tendency for excited technical people to build exciting technical things without always considering the consumer. Specifically, there is the question of distributional impact—the winners and losers, and making sure that, even if there is some really brilliant stuff happening here that is helping a certain bracket of consumers, that is not to the detriment of or is simply completely leaving behind another group. That would be my key message.

Will Hayter: My recommendation is this: let us try to make the most of the opportunity that all this data and new technology provides to get it all working for consumers as opposed to against them or 'to' them. We need to think about the building blocks for data and AI to work, again, on

behalf of consumers. We would highlight the start of that as the portability provision in GDPR. We need to make the most of that and make it work in practice. That is not the start; there is plenty more to come besides.

The Chairman: Thank you both very much indeed. That has been a really interesting session. It has been the first time we have really focused on the consumer, and we will take to heart your final two recommendations. Thank you very much, Colin and Will. We have had a very good session indeed.

Competition and Markets Authority and Citizens Advice – Oral evidence (QQ 85–94)

Competition and Markets Authority and Citizens Advice – Oral evidence (QQ 85–94)

[Transcript to be found under Citizens Advice](#)

Digital Catapult, NVIDIA and Dr David Barber – Oral evidence (QQ 38–45)

Digital Catapult, NVIDIA and Dr David Barber – Oral evidence (QQ 38–45)

[Transcript to be found under Dr David Barber](#)

Dyson, Fujitsu and Ocado – Oral evidence (QQ 105–115)

Evidence Session No. 12

Heard in Public

Questions 105–115

Tuesday 14 November 2017

[Watch the meeting](#)

Members present: Lord Clement-Jones (The Chairman); Lord Giddens; Baroness Grender; Lord Hollick; Lord Holmes of Richmond; Lord Levene of Portsoken; The Lord Bishop of Oxford; Viscount Ridley; Baroness Rock; Lord St John of Bletso; Lord Swinfen.

Examination of witnesses

Paul Clarke, Dr Joseph Reger and Dr Mark Taylor.

Q105 **The Chairman:** This is the 12th formal session for the inquiry. We have Paul Clarke, chief technology officer of Ocado; Dr Joseph Reger, chief technology officer EMEA, for Fujitsu; and Dr Mark Taylor, global strategy and research director of Dyson. Thank you very much indeed for coming today.

I will go through the usual rubric and then ask you to introduce yourselves. This session is open to the public. A webcast of the session goes out live and is subsequently accessible via the parliamentary website. A verbatim transcript will be taken of your evidence and this will be put on the parliamentary website. A few days after this evidence session, you will be sent a copy of the transcript to check for accuracy. We would be grateful if you could advise us of any corrections as quickly as possible. If, after this session, you wish to clarify or amplify any points made during your evidence, or have any additional points to make, you are very welcome to submit supplementary written evidence to us. Would each of you like to introduce yourself for the record, and then we will begin with questions?

Dr Mark Taylor: Good afternoon, everybody. My name is Mark Taylor. I am responsible for strategy and research at Dyson. Dyson is an engineering and technology company which develops domestic and professional products for consumers and businesses around the world.

Dr Joseph Reger: Joseph Reger. I am Fujitsu's chief technology officer for Europe, the Middle East, India and Africa. I am also a Fujitsu fellow.

Paul Clarke: Paul Clarke. I am chief technology officer at Ocado. We are the world's largest pure-play online grocery retailer and we build all the technology that we operate ourselves.

Q106 **The Chairman:** Thank you very much. I will start with a pretty broad

question. What are the biggest advantages and disadvantages that AI could bring to consumers over the next 10 years? In that context, it would be helpful to know why your companies have invested in AI in particular. It might be useful to give your impressions of the Government's AI review and say whether that has relevance to your businesses. Would you like to start, Dr Taylor?

Dr Mark Taylor: Some of this may already have been covered in previous evidence sessions, but I certainly view the emergence of artificial intelligence, particularly at this point in time, as offering a fourth industrial revolution, following the mechanisation of the textile industry, the automation of factories, and digitisation in the last century and this century. If you consider that it has the possibility of being a fourth industrial revolution, of course any industry or Government or citizen needs to take that incredibly seriously.

For Dyson we see massive opportunities in using artificial intelligence techniques. In the product development process it is helping us to design better products—using high-speed computing and machine-learning techniques to optimise air paths in cyclones in vacuum cleaners, for example. It also enables us to make products that are much more personalised and differentiated, that are much more adaptive and which can adapt to their environment over time. That creates an opportunity to develop a new technology in a new area that brings a product benefit and allows us to be more successful as a company.

To speak to the disadvantages, as with any change, it is difficult to foresee all the possible implications of that change when you are standing in front of it and it has yet to unfold. Obviously, there are many challenges, as have probably been voiced in previous sessions, around data privacy, algorithmic transparency, labour markets, et cetera. I am incredibly optimistic about the future. The human race has been incredibly resilient and has had a great ability to deal with all those revolutions that have gone before and adapt and change and become stronger as a result. I do not claim to be able to see what will happen through this, but I remain very optimistic that we will find ways of adapting and using the technology for all our benefits.

Dr Joseph Reger: You asked about the consumer aspect of it and I would like to respond to that, but I would like to emphasise that that is maybe not the most important point, and I will come to that.

From a consumer perspective this is very attractive: it is going to be a safer, healthier, more prosperous and certainly more comfortable world and, therefore, the consumer does not see the price tag on it. The price tag on it is large and it says that consumers need to give up a big part of their privacy and provide personal data so that it works well. It is a deal, but it is not clear whether it is a good deal or not for the consumer. We see that in Fujitsu and since we are a responsible business—two years ago we won the Prince of Wales award for responsible business and we are very proud of it—we entertain these ideas and investigate what this means to people in general.

For a country such as the United Kingdom, as for most larger European economies, that is an important aspect. I would not argue with that.

However, the ramifications and consequences of artificial intelligence for the industries could be even bigger and more significant because it is about competitiveness in the future world and the ability to participate in international value chains, and so on, so the impact on the industrial side could be much larger.

In reference to Mark, predicting it is difficult; however, the public in general across the whole of Europe is being a bit sedated by statements such as, “It’s not going to be so bad. The first industrial revolution with the steam engine brought about more jobs and the second did too and so did the third; therefore this one—the fourth one—will do the same”. That is a big mistake because the current one is unlike any of those before and the honest-to-God statement is that we do not know what the net effect is going to be.

The Chairman: To take up one point—I will come to Paul Clarke in a minute—you seem to imply that there is a competitive adoption dynamic of AI that will simply drive companies to have to adopt AI, in a sense, to be competitive, and that that is almost an inevitable process.

Dr Joseph Reger: Indeed, and that competition is going on now. The foundations of it are being laid and industries will be redefined. Incumbents are having difficulty staying in their own industry and newcomers are taking the businesses of existing companies. There is an incredible restructuring of the value chain space going on across the world. I travel a lot in Europe and I do not know of a single European country where no attempt would be made to establish AI centres and be competitive and leading in some way, because everybody is trying to get on top of this.

Paul Clarke: Focusing on this through the consumer lens, as we have been asked to do, I would perhaps set it in the context of our business. Our business runs on an intersection of five disruptive technologies: the internet of things, big data, robotics, AI and cloud. Of those we see AI, in the Tolkien sense, as the one to rule them all. It is the one that lets you do the really exciting things with the others. It pervades what we do across our business, and yet, at the same time, like many others, we are just getting started here.

I see amazing opportunities for consumers going forward and we are working on a lot of these the moment. There is faster lower-friction shopping, greater personalisation and adaptive user interfaces that can respond to different people’s shopping styles and agendas; better, smarter recommendations, predicting what customers want ultimately before they have a clue themselves. That is the journey we are on. New kinds of interfaces are emerging—voice is a huge one at the moment and augmented reality, and AI is at the heart of those. There are increased levels of customer service on the back of greater reliability and efficiency. It is helping customers, and indeed employees, make fewer mistakes, for example using AI for monitoring and oversight, to be a bit like the third gyro in the aircraft, sitting there on your shoulder trying to watch for when you are about to make a mistake, augmenting you as a human being. It is about making better use of scarce resources, whether those be time, energy, water, land or transport networks. That is very central

to what we do in the routing systems that we create: using all the data from today's routes to drive smarter ones tomorrow. We are trying to keep ourselves safer, whether it be with physical or cyber security, and helping to manage our privacy. I am sure we will get to that later on. Those are very central to what we do at Ocado as an online retailer.

The disadvantages are hard to describe from a consumer's perspective. Taking the last one, you are in an arms race there. As well as being able to use AI to try to keep people safer, clearly, other people are going to use AI as new kinds of attack vectors in cyber security and fraud. Along with other members of the panel, I am very optimistic about this future.

Lord Levene of Portsoken: Dr Reger, we have heard quite a lot of evidence, and I do not think much of it so far has said to us, "You don't realise how bad this might be. It really could get very serious. We have opened something up here which we could all suffer from". Accepting for a moment that that might be the case, does that not in itself open up the possibility that the effort being made up until now is not necessary because it will be handled in another way, and that that effort could be put into new areas which we do not even have the ability to pursue at the moment?

Dr Joseph Reger: In all likelihood that will happen. However, the basic issue is that AI is quite possibly the most powerful technology in the history of mankind, and all powerful technologies, without exception, have been used both as a force for good and a force for bad. AI is being used, essentially, for weaponisation of code these days and therefore that is a situation we need to look into. It is very unlikely that it will get out of control at this development stage. We are at the very beginning of AI, and the current phase of AI, which is commonly dubbed as artificial narrow intelligence, is a phase where that is not such a big danger.

What worries us at Fujitsu in developing these technologies and seeing their potential is that there is no societal debate about where we would like to take it to, and what kind of societal agreement needs to be in place so that we know what kind of control we need—if we need it. There is no agreement about how quickly we are going to use it and introduce it in various segments or aspects of life. Technology is developing at exponential speed. For that reason, as an accompanying process, we need these kinds of societal debates right now.

The Chairman: We will try to push that very debate along ourselves. Lord Swinfen.

Q107 **Lord Swinfen:** Gentlemen, how difficult is it to hire staff with the requisite skills to develop or make use of AI? What impact has AI had on productivity in your businesses? Have you automated any roles? I would be surprised if you have not. How informed are your existing staff about the opportunities and risks?

Paul Clarke: Data science and AI are at the most overheated end of the computing skills spectrum. There is definitely a massive shortage, as has been noted in the AI review, of graduates and postgraduates emerging with those skills. What is not talked about enough is the fact that those skills lie at the end of what is a pipeline of digital literacy that stretches

all the way back to primary school. There is a massive amount more that both government and business can do to help look after that pipeline across its entire length. Definitely, if we want a more diverse set and more graduates emerging at the end of the pipeline, we have to do much more at the start, whether it be, as the recent Royal Society report mentioned, making sure that we invest in more qualified teachers, or considering mandating schools to offer these subjects in digital literacy up to GCSE and A-level, and considering mandating these subjects in the curriculum up to GCSE like we do with maths and English.

That pipeline does not stop at university. It needs to continue on into industry, and we need to do much more to incentivise industry into that process of continual learning, particularly when it comes to subjects such as AI. There is no getting away from the fact that initiatives such as the apprenticeship levy have carved a hole in the available budget that companies, including ours, have to spend on that continual learning. We would argue very strongly that that might be better transformed into a training levy that can be used by companies in a much more holistic way to make sure that they continue that process. As was said in the previous session, we are going to have to fuzz this boundary between education and work in quite a considerable way, and we are going to have to see it as a much more hand-in-hand and continuous process of reskilling and reinvention across one's lifetime. There is much more we could do there.

To finish, one of the challenges is the fact that we are looking at what is a massively non-linear disruptive transformation in a very linear way in terms of the problems, the impact and our necessary response to it. I am sure we will get on to that later.

Dr Mark Taylor: I echo pretty much all those points. It is extremely difficult to hire AI talent in the UK. It is a very important, growing and sought-after area, and graduates and postgraduates with machine-learning skills command many job offers and extremely high salaries. Yes, it is very difficult. I completely agree that the problem facing us starts at school. It starts with teaching computer science and STEM subjects, and encouraging their study, and bringing machine learning into the curriculum, as I note they are about to do in China—training our children and the future workforce in what will be one of the core skills, in order to be successful as individuals and as a society in the future. It translates into university. We do not have enough STEM students at universities. We are not training enough people with these particular AI machine-learning skills. Therefore, we are not producing enough graduates or postgraduates in this. That is about looking at the curriculum and funding for schools and providing the right capability for schools to teach the curriculum and, obviously, expanding the number of places at university and postgraduate PhDs and MScs.

That will all take time. I mentioned at the beginning how important this was as a revolution, and it is important that the UK is successful at harnessing the use of AI. We also need to make sure that we can hang on to the graduates who are doing postgraduate training when they finish their master's and PhDs. Some 70 per cent of those in STEM are from overseas. We need to encourage them to remain in the UK and continue to pursue their discipline in industry or in academia, because we need to

retain that talent that we have helped to train and develop here in the UK for the future. A broad response is required through secondary, through university and then at postgraduate level.

Dr Joseph Reger: If I start with the second part of the question, yes, we have automated roles. We are using artificial intelligence technology in particular in the form of machine learning, which is 99 per cent of artificial intelligence today. Whenever anyone says AI, they mean by that machine learning. We do that in the area of service automation. We have chatbots for our customers to take trouble tickets in service situations, and so on. We use it for cyber security because these machine-learning routines are capable of identifying threats the first time they occur, which is when they are the most dangerous, through anomaly analysis and other methods. We do that and it is difficult to find the right people for it; the market is essentially empty. We reskill a large number of our engineers. We have an internal programme called Fujitsu Distinguished Engineers, which is a merit-based programme, and within that we have an artificial intelligence community. As with many modern movements, it is only very slightly organised. It is self-organising and people get together and teach each other about the latest developments, and that is the fastest way to do it.

I do not believe that universities can provide us with AI-ready people because it is a mixed skillset. I would have trouble identifying which faculty should produce them because there is maths, information theory, physics, complexity theory and all sorts of other things in there, and, therefore, it might be better if we just had capable people coming out of universities who get training on the job and become very useful in real projects very quickly.

One quick point I would like to make is that we have no time. The first Industrial Revolution with the steam engine was 200 years ago. The second with electricity was 100 years ago. The third with automation and numerical control of production machinery was 50 years ago. The next one is likely to be 25 years from now—in other words, within the lifetime of a person. These things have accelerated to the level that we do not have time for long preparations and we cannot wait until our universities or other institutions provide what the industry needs; we have to provide it ourselves.

Viscount Ridley: Mr Clarke, you talked about a pipeline of digital literacy going back to primary school being necessary. In its written evidence Ocado said that coding should be mandated just as English and maths are. That goes a bit far for me. I am a bit sceptical about this. Could you help dispel my scepticism? I learned FORTRAN and BASIC in the 1970s and that was of very little use to me in later life. Surely anything you teach now in school is going to be way out of date by the time these kids leave. Also, is teaching coding not like teaching kids how to build internal combustion engines in the 1920s? What we need to learn is how to use these technologies, not how to make them. Finally, are coding jobs not going to be the first to be automated by AI? Is reskilling of well educated, generally educated people—in STEM subjects particularly—more important than coding in schools, which is what a lot of people say we need? Sorry, I am not picking you out, but I am not convinced.

Paul Clarke: I completely share your view. That is why I use the phrase “digital literacy” as opposed to “coding”. I see digital literacy as being a much bigger portfolio. It includes things such as data literacy: how you harness data, how you visualise it, how you model it, how you understand bias. Data is core. It is one of the foods of AI. Teaching our children to have mastery of it, to be data whisperers rather than just at the mercy of it, is extremely important. Understanding things such as the ethics and the philosophical challenges around these technologies is very important. I completely agree with you that we need to rethink the whole curriculum in terms of future proofing it, because many of the things we are teaching our children now—and I say that as a father and as an employer—are going to be as disrupted as the encyclopaedia has been by the internet.

Therefore, as I am sure has been said before many times in your evidence sessions, we need to develop people’s interpersonal skills, creative skills, intersectional thinking, strategic thinking, agile thinking—the process of reinvention. How do we teach our children to view reinventing themselves not as a mid-life crisis but as a skillset, if you like, because they are going to do it many times? Preparing children in the widest possible way for this much smarter, more automated future is what we need to do, and we need to do it across the whole curriculum. Fundamentally, it is not just about teaching our kids to code. Viewing it like that is tinkering at the margins. I quite agree with you also that those skills will become out of date very quickly. As was said in the AI review, and by Dr Reger—and I would completely agree—mathematics and a whole portfolio of other skills are really important here; it is not just coding.

Q108 **Lord Hollick:** Dr Reger, you have urged us and the Government to hurry up and get on with it. The train has already left the station. What steps should the Government take and what measures should they introduce to accelerate the development of AI and its deployment across the UK, in particular in the business sector?

Dr Joseph Reger: Governments in general—the UK Government might be an exception, and I hope they are—like to regulate. AI technology does not need regulation because it is a competitive race and the faster the United Kingdom progresses in that race, the better it is for the country. As I said, it needs a parallel activity that discusses the possible issues and the challenges that come with it. We need a societal debate on it and to agree what we will do in parallel with all the research. If we want to have any sort of control, be that for ethical or moral reasons—we know of instances in which machine learning is already producing results that are not acceptable in many countries for ethical and moral reasons—we have to do it now because these things have to be at the very foundation of those future systems. Retrofitting them is rather cumbersome, very expensive and probably not doable. Therefore, this needs to happen now. My advice would be to have no regulation on the technology side. The country should try to instigate as much innovation and as many projects as possible, but start a debate where one can assess the possible downsides as well.

Q109 **Baroness Rock:** I would like to move on to consumer awareness. We have heard in previous sessions that awareness of the concept of machine learning is relatively low, whereas awareness of its applications is much higher. How aware are consumers of the role of AI in products, or indeed in services that they use? Do you think it is important that they are aware and have an understanding? Dr Taylor, could we start with you?

Dr Mark Taylor: Most of us are interacting with AI on a daily basis, whether we are doing a search, using a voice assistant or online shopping. The list is endless. This year there has been a widely published study which suggested roughly two-thirds of consumers did not know they were interacting with an AI entity when in fact they were. That was a survey of about 1,500 people across half a dozen different countries. I do not know how well the survey was done but it seems plausible. That is the point: the algorithm—the AI—is intended to be invisible and behind the scenes. To a very large extent, particularly as we are talking at the moment about narrow AI, as the gentleman to my left mentioned, that is not a major issue. When an AI algorithm is making very significant life-or-death decisions, or near to, over an individual, transparency becomes essential. When it is translating some text for you from one language to another for your personal use on a website, maybe it is less important to know as a consumer how it is done. The context as to how AI is being used, for what purpose and how the individual is impacted by that is the most important lens to look through when thinking about how aware people should be.

The other angle of AI is obviously around data and how people are using it. You have algorithms and data, and how data is used is incredibly important to individuals. I am sure we all care very deeply about our own data, as does everyone else in the rest of the country. Transparency around data use is important. The GDPR regulations that are coming in next year are very welcome and are a very good strengthening of data protection. There is room to go still further—it was spoken about in earlier evidence sessions—and personal ownership of data may be something that we want to move towards as a society.

Baroness Rock: We are going to come on to a data in a second so maybe Dr Reger could talk about consumer awareness within products and services.

Dr Joseph Reger: My view is that the consumer must know. Not only must the consumer be aware; the consumer needs to understand the details. It is not enough to know there is some AI or machine learning going on. They need to understand how the system works, at least on some level, and what kind of data it uses and where it got that data so that consumers can make an educated choice. Currently consumers are accepting contracts and it is akin to the problem with the smartphone, where lots of things appear to be free but they are paying for it with their data; it is just that they do not know about it or it has never been explicitly said that that is what is happening. We need to raise this to the level of a contract—opting in, opting out—with an educated consumer saying, “This is going too far, I am not going to do this; but this is acceptable and I am going to accept that part”. That is not done.

Everything is in a single bucket today and it is much easier to accept and just go with it. As I mentioned in my statement at the beginning, AI and machine learning is a very comfy world for consumers; therefore, they are enticed into accepting maybe more than they should.

The Chairman: Paul Clarke, do you have the same view about the fact that consumers should know?

Paul Clarke: Definitely. To answer it as an exam question, they are relatively unaware, but becoming more aware as devices and cars get smarter and we have voice hubs in the home. People are being exposed to it, but, at the same time, taking our business, they probably have no idea of some of the back-end uses of AI behind the scenes, controlling huge populations of robots collecting their groceries and the role that machine learning plays in being able to do that. There, you are talking about tasks in terms of complexity and data and real-time control that are beyond what you could do with humans. There is some awareness and it is growing.

I would agree on both of those points. On the transparency point, there will be the kinds of decisions that AIs make where it is more important, but, at the same time, there is no absolute here. We have to benchmark the decisions that are being made against a decision a human would have made. Humans are not totally transparent in the decisions they make and they are certainly not free of bias. It is not as if we are comparing what an AI might do with some utopian absolute. There is a huge amount of research going on at the moment on how to “white box” what is currently a black box, and I will not go into that. I know we are going to get on to data and there is a lot more to say there.

Q110 **Lord Levene of Portsoken:** I think we have really covered this, but just to spell it out a little more precisely: are consumers concerned about how their personal data is used in commercial products and services? Do they mind?

Dr Mark Taylor: Do consumers mind how their data is used; is that the question?

Lord Levene of Portsoken: Yes.

Dr Mark Taylor: Yes, I would anticipate that they do. It is very important to be very clear with consumers how their data is going to be used. At Dyson we do not share data with any other company. We do not sell data. We have a very clear end-user licence agreement between individuals who use our products and services, and the company. We try to be very transparent about what that means and use very simple language. It is vitally important that there is transparency, honesty and ethical and moral principles behind how a company uses the data from consumers, and it should not do anything with it that it has not explained to the consumer it is doing.

The Chairman: Can I follow up on one point? You talked about using simple language with data, but of course we did not get on to the point about how you practically explain the use of AI generally in the products and so on. Do you think the same sorts of rules apply and that, if you use simple language when you are applying it to particular areas such as the

ones you have mentioned, that is the right way forward? Is it possible to adopt the right kind of language in those circumstances?

Dr Mark Taylor: Yes, that seems perfectly possible to me.

The Chairman: Without alienating the consumer or making them think, “I am not going to buy that”. Are there not some issues there?

Dr Mark Taylor: I perhaps have not thought this through entirely, but with Dyson we try to explain how our products work. It is part of how we operate as a company. We try to explain the technology behind our products. With artificial intelligence and the use of machine learning we would want to explain how we are using a technology to improve the product and be very transparent and open about that. We want to be very clear about how we are using the data and how we do not sell it or use it for any other purpose than to improve the performance of that product, and how we dissociate data from products from consumer data and follow the law in doing that correctly.

Dr Joseph Reger: Could I add a one-sentence comment to whether consumers are concerned? It depends on the country. I am in an international role. I can say that in Japan it does not seem to be a big issue, in Germany it is a very big issue, and the United Kingdom is somewhere in between. It is not homogenous within a country.

The Chairman: They like robots in Japan as well, do they not?

Dr Joseph Reger: Yes, very much so. Because of that, there are interesting consequences of data usage as well. There is an issue here, and it has to do with the fact that most of the methods that are used to collect consumer data to good purpose work only if there are a lot of them. Everybody has to contribute and only then is it really useful. However, what happens if whole groups in society decide that they do not want to be part of that? The applicability of it, in particular in public administration, is diminished, and therefore it is a problem. We cannot discuss the matter by saying data is one thing and people are like that. We need to go into much deeper detail: what kind of data, for what kind of people, in what role and for what purpose? I alluded to the fact that the societal debate needs to detail that and give it a structure, because the average statements can be misleading.

Paul Clarke: Just to link this back to the previous question for a second, one of the reasons why it is so important that consumers understand the role of AI is that it is part of developing this trust between the providers of services and the consumer—both regarding the role that AI is playing and, as has been said already, the role that the consumer’s data plays in enabling AI to provide those services. What is particularly interesting here is that there is a slightly “damned if you do, damned if you don’t” situation. On the one hand, we know that our customers expect our systems to get to know them over time. They expect them to become knowledgeable about what they want. You could sum that up as someone saying, “I have been shopping with you for years. Surely you know I don’t like fish; why do you still show it to me?” On the other hand, consumers are equally capable of being offended if you make a conclusion about them, even if it is correct, such as, “How dare you assume I’m vegan”. In a sense, the poor old AI is caught somewhere in

the middle. That is why trust and understanding is so important in order for people to realise that it is a new kind of contract.

The Chairman: Presumably your brand is very important in that respect as a commercial entity.

Paul Clarke: Absolutely.

The Chairman: That is the trust creator, in a sense.

Paul Clarke: Totally. How data are used and for what purpose, and the security and privacy implications around that, are absolutely central to that; otherwise, I agree very much with what has been said.

Viscount Ridley: A point that has not really come up in our inquiry so far that is quite important, and which was raised by your fish story just now, is anthropomorphism. Part of the reason we are offended and worried about this is that we think of it as a person, and we need to teach ourselves not to. The classic one is the satnav—you hear the irritation in her voice.

Paul Clarke: I completely agree. It is absolutely extraordinary, as was said in the previous session, how almost first-generation devices when they come into our homes become members of the family. They certainly have in my home. When it is said that these technologies cannot be used in a caring way, I am sure there are going to be, for the foreseeable future, limits on that, but I would definitely never say never. I would also say that there are huge aspects of caring that are not just about empathy. This is one of the many big societal challenges that AI and associated technologies such as robotics can help us solve: how we care for an ageing population where otherwise we will lack the resources to do so.

Q111 **Lord Hollick:** When one of your autonomous delivery vehicles mounts the pavement and knocks over a pedestrian, are you liable or is the manufacturer of the car liable? For the other manufacturing companies, where does liability lie and are changes needed to clarify this?

Paul Clarke: AI definitely raises all sorts of new questions to do with accountability. Is it the person or people who provided the data who are accountable, the person who built the AI, the person who validated it, the company which operates it? I am sure much time will be taken up in courts deciding on a case-by-case basis until legal precedence is established. It is not clear. In this area this is definitely a new world, and we are going to have to come up with some new answers regarding accountability. To answer your specific question, I think it would depend on the details of the case.

Lord Hollick: To Dyson: when you are manufacturing products that go haywire, what happens?

Dr Mark Taylor: All our products are covered by legislation in the markets in which we sell them. There are safety regulations, electrical regulations and various other laws that we have to comply with. At the moment, I do not foresee a situation with our products where we would fall outside the existing legislation. Looking at the transcript of a previous evidence session that you ran with three far more capable legal minds

than mine, there was a view that quite a lot of existing legislation covered many of the foreseeable outcomes of malfunctioning products or services with AI within them.

The Chairman: We are coming on to another question on the ethical side, which has some bearing on this.

Dr Joseph Reger: Let us remember that as a company we encourage and serve our customers who would like to use machine learning and artificial intelligence to improve their competitive situation. Supposing this works well, it would be really bad if liability and legal issues took away what they had just gained by using these technologies in productivity and so on. The problem is very obvious now in the car industry with autonomous cars, where it is clear that the car manufacturers will have to take responsibility for these systems. If you take that as a model for the rest, it appears to some of us that it will have to be the company at the end of the value chain which sells the product to the customer that takes responsibility, because it would be an impossible expectation on the consumer or the end customer to be able to figure out what kind of value chain is behind it, who sold what, who has done what and who is legally responsible. That would be totally impossible for them, and, therefore, we need an end point in the value chain with that kind of responsibility. I know this is hard for industry. Our customers who are using this technology are facing these issues already. We need a legal system that keeps up. Currently that is not case, as mentioned, but it needs to keep up because these products are hitting the market already and therefore the questions of liability, responsibility and accountability need to have a new definition very soon.

Q112 **Lord Giddens:** Feel free to answer my question as you wish because it covers some of the issues that have been discussed. What role do ethical guidelines or principles play when your companies develop AI systems? I do not know how far you are integrated with the giant platforms, but when you look at the amount of data that Google or Amazon or any of the five massive companies have on you, it is awesome. They can know more or less where you are, what you are saying and track your habits. How does one deal with the implications of that? I do not know how that specifically interacts with your more narrow AI, but it is a very generic issue. There is some counter-reaction which I think companies here will have to take on board. If I can put it this way, I have my spies in Silicon Valley who keep me informed every day. There is a huge battle going on now in Silicon Valley. These companies which were once gods are now deeply embattled and have to be embattled in some way because they cannot any longer take for granted the power they have had over our lives and the invasion, which we have kind of acquiesced in but has all sorts of consequences, some of them deeply destructive. Therefore, it would be interesting to know how the ethical principles you apply are affected by the wider knowledge environment and the surveillance which the big companies conduct on our habits.

The Chairman: Dr Reger, you look poised to answer that question.

Dr Joseph Reger: The giant platforms exist, there are only a handful of them and for the most part they collect data about consumers. This is a

concern and we have to debate it. However, let me come back to my original statement that there is another AI machine-learning opportunity here and that is in the industry. Companies—car manufacturers and so on—collect their own data about their own products, and I am sure Dyson is collecting data about its own products that is not directly in Google, and that is clever and how it should be, because companies which use machine learning for industrial purposes like secure platforms and protected data. Fujitsu is very much in the business of providing these alternative platforms. We have our own cloud and AI technology that can be used for these purposes. To compete with Facebook and Google on data about consumers, particularly that available on social networks, is very hard—in other words, impossible—so that is not our objective. We focus very much on the value creation of our customers as a company, whatever they do, and we support them to build their own datasets. I do not believe that the future will involve a single dataset controlled by one company. It is not going to happen because it is not in the interests of any industry.

Lord Giddens: If a customer asked you, “Would you release all the data you have on me to me?” would you do that?

Dr Joseph Reger: We do not directly target the consumer markets. We develop technologies that would be applicable and we would be working with a company which does that for the end consumers.

The Chairman: Can I ask a question that underlies some of the interests of the Committee? Am I right in thinking—and I will come on to you, Paul Clarke, in a minute—that none of you has publicly available ethical guidelines on your adoption of AI or on the way that you incorporate AI into your products or services?

Paul Clarke: That is correct at the moment. This is a fast-moving area. Technology companies are starting to form frameworks or groups to try to work out what kind of ethical frameworks are necessary to deal with that. Clearly, other organisations such as the Royal Society will have an important role to play in guiding what would be an appropriate template for that. We have governance structures in place that obviously look after the data and risk, but, you are right, we do not have a formal published set of ethical guidelines on AI at this point.

The Chairman: What would you say as regards Lord Giddens’s questions on the Silicon Valley giants?

Paul Clarke: As was picked up in the AI review—

Lord Giddens: I only meant it as a broader example of the fact that ethical principles should be brought out into the daylight, specified and made known to those whose data is being used.

Paul Clarke: Being clear about what you will do with people’s data is obviously very important and that includes, potentially, how they will be used to deliver smarter, better services on the back of that data. It is more than that, because we need to create new ways for these data to flow to more easily, where appropriate. We have to create new kinds of data marts because open data is only a small part of what we will need. We need ways for companies to be able to exchange data, but to do it

with the appropriate kinds of passports and metadata that make it clear what that data can be used for. I am not talking here about personal private data but other kinds of datasets that companies may want to exchange with each other. Some will have value but they will be very important, once again, if the UK is to make the most of this opportunity. If everybody keeps everything to themselves, we will not create the richness of the intersection of these datasets where some of the most exciting things are going to happen.

Q113 **Viscount Ridley:** Dr Taylor has definitely answered the question on transparency, but perhaps we could have a little more from the other two witnesses. It is about transparency and the extent to which we should be able to look inside the black box and see how an AI reached the decision that it did. The evidence we have heard in this inquiry has been surprisingly contradictory. Some people have said that AI makes this easier, the coding is simpler, and you can more easily interrogate these systems and find out what is going on in them. Other people have said it is impossible and you are holding it to a ridiculously high standard because we do not even hold human beings to this standard, et cetera. Where do you fall on this? How does the GDPR's call for a right to an explanation affect your businesses?

Paul Clarke: I mentioned earlier that I thought transparency was not an absolute, and clearly we have to benchmark this against the transparency and bias that exists within human decision-making. Transparency from AI is desirable but, as I said earlier, at the moment it is beyond the state of the art, certainly in terms of convolutional neural networks, but a lot of work is going on there. Do I think it is desirable to be able to do that? Yes. Is it easy to do that at the moment? No, and in many cases it is not possible, so we have to judge the decisions that are being made based on the quality of those decisions, and we have to choose where to use AI and maybe where not to use it.

Dr Joseph Reger: If I start with the GDPR question, GDPR is very much needed, and I do not think anyone wants to know whether we like it or not. It is coming for definite and is needed. Our company has gone through enormous efforts to be prepared for GDPR, but it is a different question from AI because in GDPR you can provide the transparency. We need to be more precise. AI might have the mechanisms, and we hope that they will be developed, for transparency. Machine learning, which is by far the overwhelming majority of AI in current applications, does not. When you open up a box that has been trained, what you get is a couple of million floating-point numbers and that is about it. It does not give you any insight into why and how decisions have been made. Very often those decisions have nothing to do with the algorithms and the box but with the data that has been used for training it. These machine-learning engines are generic in nature and they will learn what kind of data is provided to them. The real question is what kinds of data are being used, not so much what kinds of systems we are using. If a credit-rating system develops a racial bias, which has happened in documented ways many times, the question is not what kind of machine-learning box it is but what kind of data has been fed in. Even if we find out what kind of data, and people could argue that is the data we have and it is apparently

evidence of certain biases, as a society we still might decide that it is not acceptable for us. Even if the data leads the systems to those conclusions, it is not acceptable for us and we want to have a mechanism against that bias. At that point, we have to build the systems that do not make these mistakes, which is why I was mentioning that we need the debate on this now.

Q114 **The Lord Bishop of Oxford:** You may want to answer this quite briefly. Do you believe we need a specific watchdog or regulator to protect consumers with respect to AI-based products and services?

Dr Mark Taylor: Data protection, algorithm bias, algorithm fairness, algorithm transparency, everything we have been talking about, are all vital for public trust and acceptance of the use of AI in products and services. To the extent that these are not already enshrined or covered by law, whether it is in the impending GDPR on the data side or laws that govern the products that those AI algorithms are expressed through, yes, having some framework to protect us from bad actors might be necessary. I go back to my previous point: you met three far more qualified people than me to talk about the law and regulation on that, and I think they drew some good conclusions.

Dr Joseph Reger: We will find that different countries respond to this challenge differently. In some countries, legal frameworks are being developed already and they will have the legal frameworks for it rather soon. Whether that has been exaggerated, we will know when we see them. In other instances, it will be more case by case, and case law, and then we shall see where we arrive. I come back to what I said earlier on the technology side—we do not need watchdogs. We need an understanding of the kind of situation we are heading into. There needs to be a debate, with the participation of the technology companies. This is one of those instances where the political sphere and the technology sphere need to work hand in hand.

Paul Clarke: I do not want to repeat what has been said, but we have to be mindful about the interplay between innovation and legislation. We have to be very careful that we do not legislate ourselves into a self-fulfilling future that may be just as uncomfortable to us in other ways as the one that we fear, and that drives us towards the legislation in the first place, because of unintended consequences. Clearly, although we are an island, we are not an island in the sense of how these technologies will play out, and we cannot afford to play King Canute. Therefore, we have to be mindful that what we do not do, others may. We have to be careful about how we balance legislation and innovation going forward. Clearly, we have some very important institutions already such as the ICO. That has a very important role to play in the data. I am sure some of those may get widened in the future.

Q115 **Baroness Greender:** You get one recommendation each to wish upon us as an early Christmas gift, but only one because we are very strict. What would you like us to recommend when we conclude this report?

Paul Clarke: Only one?

The Chairman: It is not Christmas yet.

Paul Clarke: It would be a kind of meta wish. I would wish you to look at both the opportunities and challenges through a much more disruptive and non-linear lens. The challenge going forward is how we can make the scale of adjustment and response that is required within the five-year term democracy that we are in, because I think some of the changes that are needed are beyond that.

The Chairman: The two timetables do not match.

Dr Joseph Reger: An artificial intelligence council for the country would be really good, if it is put together so that it represents all the different streams and interests. It needs to have industry representatives who would like to use AI as soon as possible and as much as possible for competitiveness, productivity, and so on. It needs to have political representatives who understand AI as an opportunity to serve better the citizens of the country. It needs to have some people who are more sceptical. Professor Stephen Hawking is one of those people.

The Chairman: Or yourself, indeed.

Dr Joseph Reger: Thank you for that, sir. There needs to be this mix because there need to be debates and a good fight in that AI council.

The Chairman: Thank you. Dr Taylor.

Dr Mark Taylor: AI offers an enormous opportunity to the UK economy and the most important wish that I would have is for a comprehensive strategy around how the country is going to benefit from its exploitation. Maybe it is a bit of a 'cheat' wish, but it is one that spans education at secondary, university and postgraduate level and spans into industry and how industry can be involved in that strategy for funding, for retraining and for awareness for many segments. It would be a holistic strategy.

The Chairman: Thank you very much indeed. That brings to an end this evidence session. I am afraid it is always a bit of a gallop, but thank you very much indeed for the insights you have provided. They are extremely valuable and we look forward to looking back at the transcripts, because we always reflect on what we have heard by reading the transcripts of the evidence. Thank you very much indeed.

Professor David Edgerton, Professor Peter McOwan and Professor Sir David Spiegelhalter – Oral evidence (QQ213–223)

Professor David Edgerton, Professor Peter McOwan and Professor Sir David Spiegelhalter – Oral evidence (QQ213–223)

Evidence Session No. 22

Heard in Public

Questions 213–223

Tuesday 19 December 2017

[Watch the meeting](#)

Members present: Lord Clement-Jones (The Chairman); Baroness Bakewell; Lord Giddens; Baroness Greender; Lord Holmes of Richmond; Lord Hollick; The Lord Bishop of Oxford; Viscount Ridley; Lord St John of Bletso; Lord Puttnam; Baroness Rock; Lord Swinfen.

Examination of witnesses

Professor David Edgerton, Professor Peter McOwan and Professor Sir David Spiegelhalter.

Q213 **The Chairman:** May I extend a very warm welcome to you? We have before us: Professor David Edgerton, Hans Rausing professor of the history of science and technology, and professor of modern British history at King's College London; Professor Peter McOwan, vice-principal for public engagement and student enterprise, Queen Mary University of London; and Professor Sir David Spiegelhalter, president of the Royal Statistical Society, Winton professor for the public understanding of risk, University of Cambridge, and chair, Winton Centre for Risk and Evidence Communication.

This is the 22nd formal evidence session for the inquiry. The session is intended to help the Committee to discuss the public narratives on artificial intelligence. In view of the presence of Professor Peter McOwan, I declare an interest as chair of the council of Queen Mary University of London.

I have a little rubric I need to go through before I ask you to introduce yourselves. The session is open to the public. A webcast of the session goes out live and is subsequently accessible via the parliamentary website. A verbatim transcript will be taken of your evidence and will be put on the parliamentary website. A few days after this evidence session, you will be sent a copy of the transcript to check for accuracy. We would be grateful if you could advise us of any corrections as quickly as possible. If, after this session, you wish to clarify or amplify any points made during your evidence, or have any additional points to make, you are very welcome to submit supplementary written evidence to us. Perhaps you would like to introduce yourselves for the record and then

we will begin with questions.

Professor Sir David Spiegelhalter: I am president of the Royal Statistical Society, which means that I am statistician, but my current job is working on public engagement with statistics and risk. I also have a history in methodology work in Bayesian methods for artificial intelligence, which I used to do.

Professor Peter McOwan: As the Chairman said, I am vice-principal for public engagement and student enterprise at Queen Mary University of London. My research area is in artificial intelligence and robotics, particularly social robotics. I am the institutional lead for Queen Mary's Engage Watermark Gold Award from the National Co-ordinating Centre for Public Engagement, and I was awarded the Mountbatten medal for public engagement by the Institution of Engineering and Technology.

Professor David Edgerton: I am Hans Rausing professor of the history of science and technology in the Centre for the History of Science, Technology & Medicine at King's College London's Department of History. I am a historian of modern Britain and of technology and science, focusing on Britain and more globally. I am the author of a book called *The Shock of the Old*, a global history of technology in the 20th century.

Q214 **The Chairman:** Thank you very much indeed. You may want to answer only certain questions, so if by telepathic communication between the three of you you decide with nods who would like to lead on a particular answer, we would be very happy with that, because some of the questions may be right within your province and others may be slightly outside. We will play this by ear as we go through, if you are happy with that. There is no expectation that you will answer every single question.

I will start with a very general but quite topical question in many ways. What, in your view, are the present dominant narratives concerning AI at the moment? Are these broadly accurate? Are they helpful or harmful? We have in mind that artificial intelligence is often characterised in fairly Utopian or Dystopian ways, so we have that contrast. How do we get beyond that into a slightly more realistic perspective on the possible opportunities and the risks of AI? That may well be a question that each of you wants to answer.

Professor David Edgerton: You might be better informed on this question than I am, but my general impression is that something called the fourth industrial revolution has framed many accounts of AI; indeed, AI seems to be the central novelty in these stories of the fourth industrial revolution.

It seems to be a fairly recent term, perhaps popularised at the World Economic Forum at Davos two or three years ago, but in fact it is a term that has a history. Harry Elmer Barnes, an American historical sociologist, who talked about a "technological conception of history", described a fourth industrial revolution that was about to happen in 1948. This fourth industrial revolution would arise through atomic energy and supersonic transport. In the 1960s, he changed his mind and said the fourth industrial revolution started in 1935 and that the world was going through a fifth industrial revolution. Around 1960, there are any number

of fourth industrial revolutions going on. Indeed, some people think that AI represents, by implication, the second industrial revolution—the most important change since the first industrial revolution. We have this very particular sense of AI as a world-transforming innovation, which I think is a very unhelpful way of thinking of any particular technical advance.

Professor Peter McOwan: When I saw that this was one of the questions to be posed I did some research on this, and the most up-to-date paper I could find was by Fast Company and Horowitz, which was published this year, in 2017. They looked at the longitudinal study of the reporting of artificial intelligence in newspapers and used sentiment analysis and crowdsourcing to have those particular stories labelled as either positive or negative. They discovered that up to about 2009 there had been fairly low-level coverage of artificial intelligence and what there was was as much positive as negative.

After 2009, there was a massive increase both in the coverage of artificial intelligence and, in particular, in the coverage of more negative aspects of AI, such as fears of loss of control to artificial intelligences and the ethics of that. That included things like employability and the general negative impact on the employment market, although there were very positive aspects such as its use in healthcare. The caveat was that the data was taken from the *New York Times*, so there may have been a cultural bias because it was a US-based paper, and because, being a non-tabloid, there might be a different kind of sentiment analysis in there.

That reflects my personal views on this, which are there are some very positive stories about artificial intelligence—the recent discovery of the planetary system, for example—but those are very much being outweighed at the moment by the negative stories associated with AI. Very often those negative stories are somewhat sensational, not surprisingly because they are picked up by the newspapers and very often are not based on the credible technicalities of what is available to us at the moment and include a very large pinch of future gazing.

Professor Sir David Spiegelhalter: People are not being given a good or accurate impression of what is actually going on. Part of the problem is this phrase “artificial intelligence”, which I am deeply suspicious of. I worked in this area 30 years ago, and people were deeply suspicious of it then. Massive claims were being made by what were then called expert systems that they would discover mines and would do all this stuff. It was all utter puff, and there was so much puff. This area has been full of puff for decade after decade.

Now, there is some fantastic work going on in machine learning, algorithms, self-driving cars and all this technology, which is quite extraordinary. Whether any of it could be considered intelligent is another matter. Okay, it can play chess and do that kind of stuff, but I believe that AI should be reserved for generalised artificial intelligence, something that really shows intelligence. When that is displayed to me, I will be truly impressed. Until that happens, I will be equally impressed by the amazing technological developments in machine learning, algorithms, pattern recognition, voice recognition, speech processing and automatic translation. These are extraordinary feats that have been done.

Professor Peter McOwan: I would entirely agree with that. One issue with the term “artificial intelligence”, as David rightly points out, is that it is bandied about fairly liberally. All that today’s artificial intelligence does is find patterns in data. We are constantly finding patterns in data ourselves. When we read something, for example, we are segmenting the letters from the page and looking for those sorts of patterns. That is all artificial intelligence does. Depending on the data that we put into it and what we do with the results coming out of it, there are a whole load of interesting things that you can do with it. But, ultimately, it is simply about finding these kinds of patterns in datasets.

The Chairman: You all confirm that there is an issue here. Whether or not it is a dramatic change in technology—and you are all almost suggesting that it is not—there is a narrative issue here. How do you suggest that we get over that? For instance, are the creative arts or science fiction a way of getting through this? Are there ways in which we can overcome this that are not necessarily purely linear, in a way?

Professor David Edgerton: Can I suggest a way, which is to look back to the past to understand how unoriginal these arguments are? I think that will prevent a certain amount of unnecessary repetition.

The Chairman: May I say that that sounds like quite an intellectual approach to the problem? I am not decrying that, but an intellectual approach is not necessarily going to crack this problem, is it.

Professor David Edgerton: It could also be a very practical approach. If I read a few lines, perhaps the method will become clear. This is from some time ago: “The essence of modern automation is that it replaces the hitherto unique human functions of memory and of judgment ... Computers have reached the point where they command facilities of memory and of judgment far beyond the capacity of any human being or group of human beings who have ever lived ... In America technological change is beginning to move now even more rapidly in the white collar professions than in engineering because it is so much easier to programme operations of costings, of wage sheets ... than it is to programme an engineering job ... We can now set a programme-controlled machine tool line so that, without the intervention of any human agency, it can produce a new set of machine tools in its own image. And when machine tools have acquired, as they now have, the faculty of unassisted reproduction, you have reached a point of no return where if man is not going to assert his control over machines, the machines are going to assert their control over man”. That was a speech by Harold Wilson, 1963, “the white heat of the scientific revolution”, the scientific revolution being, effectively, the second industrial revolution.

The Chairman: We keep telling people that we have been here before, do we?

Professor David Edgerton: Yes. Unless there is evidence to the contrary, one should assume that most of this rhetoric is just reheated nonsense from a hundred years ago.

The Chairman: Are there any other approaches?

Professor Peter McOwan: I would be delighted to pick up on the science-fiction side of things. Science fiction is a double-edged sword on this particular stage. I am sitting here in front of you because I was inspired by robots on the television in "Star Trek" and "Doctor Who" and so on, and very many of my colleagues doing artificial intelligence research have been pulled in through the kind of excitement that comes from that media representation of AI. However, we also go to the movies and say, "Oh no, not 'another set of robots trying to take over the world' kind of story".

Why is that? It is interesting to look at the broader cultural issues. Japan, for example, has a very positive view of robotics. There are a number of different reasons for that. One is the fact that the predominant religion in Japan is Shintoism, which allows non-human objects to have a soul, whereas in the Judeo-Christian view of the world the only thing that can have a soul is a human being, so anything that does not and pretends to be a human being is evil.

Robots came from the Golem legend. Golem looked after the Jews in Prague and was a large clay man who was brought to life by putting special incantations and words into his mouth. In the Japanese culture, robots are the heroes. They are the ones who are saving the world and making the world a better place. In western culture, they tend to be the kinds of things that go up against Tom Cruise and in the end lose, which is good because there is a moral structure there, but there is intrinsic evil, and, of course, both of those are neither true nor false. The technology is morally neutral; it depends on what we wish to do with it that makes it good or bad.

The Chairman: It shows the power of culture though, does it not?

Professor Peter McOwan: It shows the absolute power of culture.

Professor Sir David Spiegelhalter: Media representations of any scientific technology are incredibly important. The people working in this area, the practitioners, need to take more responsibility for the media representation of their subject. If there are puff stories, they need to be called out. There needs to be more engagement with journalists and the people telling the stories so the stories that are told are gripping but accurate. We will come on to historical analogues later. When we think about things like GM or vaccines, the only way to get good stories into the media is for the active engagement of the scientists with the media to take very seriously the way in which their work is being presented.

Among the people working in this area, there needs to be a certain number of what you might call ambassadors, spokesmen, performers—people who get out there who will engage with the media, either to present themselves or to work with journalists to get accurate stories for this to be taken really seriously.

As an example, we have been trying to improve the image of statisticians. I do not think we are seen as dangerous, just dull. We have been working hard and have a whole programme of training young RSS ambassadors to get out there. One has just been made head of statistics for the BBC. I was on Radio 5 Live this morning spouting on about "the statistic of the year", and that sort of stuff. We are actively trying to

change the way in which our profession and our work is portrayed in the media.

The Chairman: So clearly it is about academics seeing themselves as performers as well.

Professor Sir David Spiegelhalter: Yes, I am a performing statistician.

Professor Peter McOwan: I am a performing roboticist.

The Chairman: I am sure you are all role models in that respect.

Professor Peter McOwan: That is a particular issue, because we have bought into this absolutely. The question is how you motivate and reward others to do it. There are a lot of people researching this in universities up and down this country and many of them are unaware of the concerns of wider society, because they have never bothered speaking to people. We need to think about how we incentivise that, possibly through the knowledge transfer network, and the kinds of things that are coming up, to make sure that people are going out there and talking, deflating the puff and telling people the exciting stories about the technology and the advantages of it, as well as becoming better researchers because they become aware of the context in the country around them and the sensitivities. Breaking down that wall is incredibly important.

The Chairman: It is about a proactive approach.

Professor Peter McOwan: Correct.

Baroness Greener: I guess the dilemma that you have provided the Committee with is where we get the balance right between the “utter puff” pieces versus the “embrace the change now” pieces, because we have a responsibility as parliamentarians to make sure that a future generation is ready to embrace change and to ensure that it has sufficient skills to do that. Where do you think the balance should lie for us as Committee in that?

Professor David Edgerton: I do not think that Parliament or the Government have the responsibility to ensure that people embrace changes that are dictated from above. It is the Government’s and Parliament’s responsibility to give people choices over which they can exercise their collective judgment. We should not assume that the stories that we tell about AI will reflect what will come to pass. I disagree quite strongly with the view that in essence we have to educate the public about a future that we already know is going to exist.

The Chairman: We like disagreement on our panels.

Professor David Edgerton: With your permission, I will continue to disagree. I do not think that the problem is really the puff generated by the media. There is a problem of elite understanding, and the notion of the fourth industrial revolution is a perfect illustration of that. Why do people who are supposed experts in this area talk in this extremely crude, ahistorical, unanalytical evidence-free way, not just about AI but about a whole number of other novelties? It is far too easy to blame the media for this way of thinking.

Baroness Greener: Where should we get the balance as a Committee?

Professor David Edgerton: There are many sorts of people who are experts in the way the economy and society are developing. The danger is that if one identifies a particular technical advance as a driver of wide social change, the authorities on that become the experts in that technique. That in itself is an unhelpful approach, with apologies to my colleagues.

Professor Peter McOwan: Actually, I completely agree with you. There was very much that aspect in what I was saying: that as researchers in technology we use jargon and are elitist. We do that to protect ourselves. Most jargons are created in that way. Also, we are not aware of the wider contributions of the social sciences and the arts. Throughout my life I have ensured that I produce that mix. It is not about going out there and proselytising to people and telling them, "You must learn about this"; it is about saying, "Isn't this interesting?" and explaining it to them in such a way that they can pick up on it. It is not about forcing people to change their mind or to agree, because that would never work.

We need to be open about it and let people be aware of all the things that that technology is capable of doing that is already making their lives far better and happier than they were before but which vanishes unseen underneath. It is like mathematics and statistics; they are all there, but nobody sees them.

Lord Giddens: I fully accept the force of most of what was said. Let me say that *The Shock of the Old* is a terrific book and deserves the enormous esteem in which it is held, but everybody struggles with the definition of AI. From what you are saying, you would not define an autonomous vehicle as displaying any form of intelligent behaviour. Only a few years ago, that was thought to be completely impossible. An autonomous vehicle has to respond to all sorts of previously unknown circumstances; otherwise, it is not autonomous. Would you include that in the category of AI or not?

Professor Peter McOwan: One would have to put it in the category of AI in the way that AI is generally understood. I would say that it was reactive and autonomous.

There is a really nice analogy in the Braitenberg vehicles, an idea created by Valentino Braitenberg. These are very simple; you connect something that detects light with something that has a motor through a positive connector between them. If you switch that on, it will move towards the light. If you change the connection, it will move away from the light. If you anthropomorphise that, you will look at that and say, "That little robot likes light", and, "That little robot gets frightened by light". You can take these very simple reactions and build them into much more complicated machines, each of the individual component parts of which is a simple reactive circuit, yet it exhibits very complex behaviour. At some stage you would say that it is doing something intelligent, because that is what you would expect a human to think about doing at the time.

Lord Giddens: I do not want to take up the Committee's time, but autonomous vehicles have to respond to many new situations that they have never encountered before, so they have to learn, surely, otherwise we would never trust them on the roads. There is something quite

significant going on there, I think.

Professor Peter McOwan: And they can have accidents, as we have seen. That is partly because they are trained on very large datasets of the sorts of things that happen, but there are situations in which the pattern of stimuli and information that is coming into them they have not specifically seen before.

Lord Giddens: It would not have general intelligence, but it would still be a machine that has to be adaptable and has to learn. You could never trust it on the roads if it were not in some sense amazingly complex.

Professor Peter McOwan: In the same way, a human being driving would come across a certain set of circumstances that they had not come across before.

Lord Giddens: I think we should draw the line there.

The Chairman: We will continue to speculate on that.

Q215 **Viscount Ridley:** I am supposed to ask you about historical comparisons and whether they are useful in understanding what is happening today, but we have already started on historical comparisons quite well, so could I broaden the question to bring in other technologies and what happened in other revolutions? I am very aware that sometimes technologies come along—reproductive technologies, mobile phones—and the public say, “Fine, that’s great”, and at other times, such as when shale gas or biotechnology come along, the public say, “We don’t like that”. I have been in contact with scientists before this freight train of controversy has hit them and they have not seen it coming. Which is this, and how do we know which this is? Sorry, Chairman, it is an off-piste question. I apologise.

The Chairman: No, it is absolutely logical.

Lord Giddens: It is a completely fair question.

Professor David Edgerton: I would start with the observation that we have any number of techniques that we use in the present and that we have used in the past. Those that have become the subject of controversy are a very tiny minority. That is where we need to start. Secondly, we invent many more techniques than we can ever use, which means that we must reject most of the techniques that are on offer.

One problem that we have here is that a lot of the standard discussion suggests that it is bad to reject new techniques, but we have to do it all the time. We cannot have 100 different kinds of telephone or 400 methods of skinning cats. We have to reject them. Scientists themselves have played a very important role in rejecting, choosing or not proceeding with the take-up of the majority of things. The public have a perfect right, indeed a duty, to reject most things that are on offer. We have a very complicated world stuffed full of changing techniques where we necessarily have to exercise choice. I would frame the question in quite a different way.

Professor Sir David Spiegelhalter: The thing that distinguishes the example that you gave is that people will warm to a technology if they

feel personal benefit from it. In the 1970s, when surveys were undertaken of what technologies people were scared of, microwave ovens were up there with nuclear power stations, but people got to like microwave ovens. Still nobody understands how they work, but they find them very useful. The question of mobile phone masts and the use of mobile phones causing brain cancer is not a big issue, because people like their mobile phones and do not want to get rid of them. People will warm to technology all the time. It is called the affect heuristic in psychology. Once you decide something is nice, you discount any possible criticisms of it. It is very popular.

You can see the opposite happening with things such as fracking where people do not feel a personal benefit from it and therefore object to it. They feel that someone else, a company that is doing it, is getting the benefit. You can see currently with Monsanto and the argument over whether glyphosate is carcinogenic or not. The farmers, who see a great benefit from using Roundup, or whatever, are very pro glyphosate; the campaigners, who do not feel there is a benefit to them of using it, are very anti.

The crucial thing in this area is whether people feel that the technology is useful to them. Of course, it is already massively useful to them. Every time they use Google Maps to do their driving, every time they take a picture and it identifies the faces and eyes and focuses, people use this technology. My feeling is that this technology will not have the same sort of intrinsic fear associated with it as some of the other technologies have from which people do not feel they are getting any benefit. It need not be like that.

The Chairman: Can you repeat what you said at the very beginning about this being an example of something? I am sure Viscount Ridley is fully aware of what it is.

Viscount Ridley: The affect heuristic.

The Chairman: The affect heuristic. Thank you very much indeed.

Professor Sir David Spiegelhalter: It is why people like nuclear power. It depends on your association with it; is it warm and fluffy or not? People who live next to a nuclear power station are often very pro nuclear power.

Baroness Bakewell: You speak of all these various techniques, most of which get rejected. Where does the whole range of choices which the public make or do not make interplay with the development, marketing, branding and advertising of something, which might be the wrong choice but succeeds because it wins its way into the hearts of the public by commercial interests?

Professor David Edgerton: It is rather concerning that we talk about technology only in relation to the final consumer. Technical choices are being made by all sorts of agents who are not the final consumer. The positing of a certain kind of consumer who is inherently distrustful of certain kinds of techniques does not seem to capture the problem at all. Lots of different bodies are taking decisions—some openly—and some cause controversy but most do not.

If your question is directed at whether the market system, the competition between private enterprises and particular research agendas of Governments produce the optimal technical development, the answer is almost certainly no, it does not. The implied question about how might we generate better choices and perhaps more optimal ones is really important and we absolutely need to ask it. We cannot ask it if we assume that novel techniques come out of the ether and we merely have to apply them. Again, it seems to me that we need to shift the discussion to ask what kinds of things we would like as a society and how we can ensure that they come about. There might be a very free market or a more statist answer, but that is where there is room for serious discussion, it seems to me.

Q216 Lord St John of Bletso: In a way, you have answered my question, because I wanted to touch on public trust. You have spoken about media representations and possibly media misrepresentations. Many of the AI community that we have spoken to have felt that the media have focused more on the threats rather than the opportunities. To that end, my two questions are: how can debate and discussion on AI be conducted in a way that will engender public trust, and do you know of useful examples from other countries that could be instructive?

Professor David Edgerton: Surely we want to generate knowledge and understanding rather than trust. That is my first point.

The Chairman: Not even trustworthiness.

Professor David Edgerton: Trustworthiness in the agencies involved in the discussions, yes, absolutely, but not trust in an abstract thing like AI.

Professor Peter McOwan: I agree.

Professor Sir David Spiegelhalter: I will channel one of your fellow peers, Baroness Onora O'Neill.

The Chairman: That is exactly why I mentioned that.

Professor Sir David Spiegelhalter: I thought so. I spout on about this all the time. We should not be trying to be trusted. As she said, you have to demonstrate trustworthiness, and that means having a degree of transparency but not just what is called "fishbowl transparency" where you tell everybody everything. As she says, the information you are using has to be accessible; they have to get it, understand it to some extent and be able to critique it.

This also goes to the idea of explanation, which we can get on to and which I am very happy to talk about. Again, you say that AI people have been saying that the media are giving misrepresentations, but, again, I am not blaming the media. When I talked about the media representations, I was not blaming the media; I was blaming these AI people. Why are they not working with the media and ensuring that the right sorts of stories appear? It is difficult, and you cannot control the media, but you can work with them. It is completely wrong to blame the media for the representations that are going on.

Lord St John of Bletso: Many would argue that artificial intelligence is not necessarily artificial; nor is it intelligent. It is a form of smart tech.

We noticed in a recent Eurobarometer survey that 74 per cent of those surveyed in the UK would “make more use of digital technologies if there were more widespread tools to improve reputation and trust”. That is really what we are trying to get at.

Professor Sir David Spiegelhalter: The idea of trust is important. The latest Ipsos MORI poll put scientists third behind doctors and nurses in terms of trust in this country which is extraordinarily high, so we are starting from a very sound basis in this country. If we could get on to this idea of explanation, it is important because it is to do with Onora O’Neill’s big point that people need to be able to check what is going on, critique it and assess why a decision is being made. We could talk a lot about to what extent the GDPR is going to make it a legal necessity to explain why a decision has been made, but, even if it does not, the idea of interpretability and explanation, or whatever you want to call it, of AI or algorithms, is incredibly important.

It is taken very seriously at the moment within the AI community. The latest AI conference is full up with how you can make DeepMind interpretable. It is incredibly difficult with some of these black-box algorithms compared with the simpler statistical things I grew up with, and still produce, where you can see the weight being given to each item of evidence. If you are using deep learning, there is no way you can work it out. People are desperately trying to work out ways to produce an explanation for why things are happening. In fact, there is quite a strong debate going on that it might be better to give up a little bit of predictive accuracy to have something that you can better explain to people. There are ways of making an algorithm more trustworthy, even if it is a black box, and that is to allow people to do “what ifs” and to start changing it. What was the crucial thing that tipped the balance between me getting a loan or not? I do not know if anyone else does this, but when I am doing my insurance online I always change my address a few times to work out what is driving the premium that is being quoted. I lie and find out what it is doing.

The Chairman: We are picking up some good tips.

Professor Sir David Spiegelhalter: It is a black-box algorithm, probably not a very sophisticated one and probably fairly crude, that is producing the premium, but I can reverse-engineer it by playing with it. That should be an option open to everybody for whom a decision is being made using some sort of algorithm.

Professor Peter McOwan: I would add benchmarking to what David said. If you have datasets that you sort of know what human experts would think was the right answer to, if your AI decision-making system is going completely off that, there is something slightly worrying in there. It is not just about looking at the inside of the black box. It is incredibly important, because part of the reason why these deep-learning systems are not used in safety-critical systems is because you cannot prove time and time again that they will work. This comes back to the point about smart cars; if you have an exception to it, it sits outside something called the frame problem. AI tends to work within a known set of parameters, and if you go outside that there is a problem. Human beings have that,

too. If you are put in a new situation, you have to make mistakes to learn from it. It is about understanding what is in the black box and what that black box is doing in a set of what would be considered sensible benchmarks for sanity.

Q217 **Lord Holmes of Richmond:** Should the public be informed when they are using a product that features substantive AI elements? If so, how? Would some form of kite mark or notification be a good idea?

Professor David Edgerton: I do not have a view on this.

Professor Peter McOwan: I think there is something in this, to be honest. One of the issues that I have also been involved in through my career is the underpinning importance of mathematics and the fact that it is quite difficult to get young people to study it, although things have improved and there have been a whole range of interventions and a whole set of reasons why that has or has not happened. I spend my time constantly reminding people that mathematics and statistics underpin so much of what they want—how their MP3 player works, and so on—because they are built into simple artificial intelligence systems. Raising awareness of the fact that it is there—in the same way in which the ingredients in foods that we eat and other things have to be labelled on the back—and saying “may contain traces of AI” sounds quite a sensible approach.

Professor Sir David Spiegelhalter: The problem with that is that you would have to define what the AI was to say, “This has 80 per cent AI” or, “This has 30 per cent”. It is not like food labelling, which is a very good example of a traffic light system and an extremely good example of public communication that people may or may not want to use, and it would be very difficult in this area. But you could put a kite mark in place, indicating that people can critique or check what is being done to them.

One of the big areas, of course, is recommender systems, such as Facebook sending newsfeeds to you. In Facebook now you can at least ask, “Why am I seeing this?” You do not get a very good explanation, but it is a start of being able to ask that extra question: “What algorithm is being applied to me at the moment?” One could expect—I am making this up as I go along—that you are told the extent to which the decision being made is automated. It is a bit like being told that you are being targeted with this advertisement: “This is not going out to everybody. This is aimed specifically at you”. You should be able to know that.

Professor Peter McOwan: It is a little like when they tell you, “This call is being recorded for training purposes”. There could be a requirement to say, “Humans are not in the loop”, or, “Automated decision ahead”.

The Chairman: That will be an interesting piece of communication. Thank you.

Q218 **Lord Giddens:** How can the implications of AI for personal data and privacy be communicated to the public? Could you comment on the implications for children in particular, because they are the first truly digital natives who have never known another world? You can see that

they cannot let go of their devices or be separated from them, and they have become, in a way, part of the machine. They are piling up data without knowing what they are doing. How can society address this issue?

Professor Peter McOwan: It is an incredibly complicated problem. There are two points that I would make. The first is about data protection legislation. We could go much further and look at the big issue of recombinant data, where a piece of data is taken from here and a piece of data taken from there and combined. In certain instances, you would give over particular data that you would not give over if you were going on to a different website. It is the same as when you tell your friends something that you would not necessarily tell your parents. If one could join those two datasets together because one knew your name, one would know a lot more about that information than you have been willing to reveal in either of those interactions. It is worth looking at how recombinant data is used.

The second example—this is one of my favourites and I use it in my lectures a lot—occurred back in 2010. I do not know if any of you have come across pleaserobme.com, but it was a website that went live for about 24 to 48 hours. It had a bit of intelligent decision-making built into it that looked at people's posts on various social media and could pick up such things as the geotagging that you have in photographs that you label, "I'm having a lovely time on holiday", or, "Here I am outside my house". The geotag for that tells you where the house is, but you know that the time stamp says that the person is over in Spain on holiday and that his geotags are ticketed over there. You can combine them to say that the house is empty. That came out of an incredibly simple set of rules that were needed to combine those. It scratched the definition of artificial intelligence, but it made a real impact, because people suddenly realised that could happen if your personal data was there and you were leaking all this information right, left and centre. Some more of those short, sharp shocks might be interesting.

Lord Giddens: One thing that preoccupies me, but not so much the rest of the Committee, is the implications of face recognition. Could you build that into your answer?

Professor Peter McOwan: Face recognition, absolutely.

Lord Giddens: I was thinking of the next stage.

Professor Peter McOwan: The technology for that is improving tremendously, but it is still very brittle at the moment. If you wear a hat, a shadow passes over your face and things like that make it very difficult to do. It can recognise a face, but recognising that it is your face is quite different. The Americans have been using it to scan for criminals in crowds at football games as people are going in and coming out. Once you have data that can look at a video feed and begin to extract other data from it that is robust, it opens a whole series of other ways of being able to build on artificial intelligence.

Professor Sir David Spiegelhalter: The particular thing about these technologies is that people do not know that they are giving data away and that things are being done with it. It seems to me that it will be

impossible to tell everybody in every situation when they are providing data. We will come on later to data governance mechanisms. Rather than each person being personally responsible for everything that is being extracted from them, it is an area where regulation and governance are appropriate. I do not think this is an object of personal responsibility.

At the same time, we think of this in relation to children in particular—I will introduce the term “data literacy”, which I will come back to later—as being enormously important. There are some new skills, particularly connected with social media, that are absolutely essential for modern citizens. They are starting to be taught in schools. There is fake-news detection by school kids. There is a lot of interest in this, and a lot of programmes and curriculums are being developed for this. It is enormously important, and there are some basic skills that need to be part of every future citizen’s armoury.

Baroness Grender: Do you have any thoughts about whether there is any potential for data to be a commodity for the young generation? Secondly, as the mum of a 12 year-old, I have to say that there is the whole concept of a smartphone—and all 12 year-olds have smartphones, however much you try to ban it—and they do not care about giving away all this data. Even though they are highly educated about how to use a smartphone, they really do not care. There is a generational divide. We have picked that up from other witnesses as well. If I say, “Don’t put on Instagram that we’re on holiday”, there is absolutely zero understanding of that from my 12 year-old. Help us with this generational divide. How can we overcome that? It is new, because a whole generation has smartphones and just does not care about how much of its information is out there.

Professor Sir David Spiegelhalter: Maybe it will have to be like our generation was taught not to talk to strangers. There needs to be a basic change in education so that people are warned from very early on. Twelve years old is too late.

Professor Peter McOwan: I agree. Some of the work that I have been involved in has been taking up artificial intelligence in discussions in primary schools. In the very simplified versions that we have, you can build a neural network with bits of string and toilet rolls, and young kids understand that and can play a very rudimentary game of snap. They are excited and interested by it and you are educating them about what is going on with data and how data can be used and so on. If that is there and it starts off in primary school and carries through, by the time they get to 12 years old, hopefully they are much more savvy about it.

We need early intervention and to embed discussions on artificial intelligence into the curriculum, because it is across the curriculum and it can really enhance the learning that young people are doing. Building that in across the curriculum and supporting teachers so that they feel more capable of being able to discuss these sorts of things is important in the same way that maths, physics, chemistry, biology, history and geography are important. These days, computer science, and particularly artificial intelligence within computer science, has become a very

important science for the fifth technical revolution, or the sixth, or whatever.

The Chairman: Or the second.

Professor Peter McOwan: Or the second.

Baroness Greender: Will a future generation ever be able to commoditise its personal data?

Professor Peter McOwan: By that you mean that an individual says, “I will tell you this if you pay me”?

Baroness Greender: “I own all this personal data and you can have it but you have to pay me”.

Professor Peter McOwan: There are companies that are already doing that. They will pay you to access your film roll so they can find out a bit about you and start to characterise you. It might not necessarily be money, but they will let you download an app that does wonderful things, and at the same time they will take some form of payment. There is no such thing as “free” on the web. You give your data and that is really what you are paying over.

The Chairman: Thank you. We must move on to Baroness Bakewell.

Q219 **Baroness Bakewell:** My question follows on very much from the discussion that we are having, which is fascinating. What do people need to know in order to make informed choices about their own lives when it comes to jobs, skills and preparing their children for the future—we have already touched on that, but it is an inexhaustible subject—given the uncertain impact of AI in the coming decades? It is interesting that the survey by YouGov shows that people are not very concerned. They do not worry about their jobs. Some 62 per cent were not worried at all. Should they be worried? Should we help them to be worried so that they can help to find a solution? I would like to know what you teach when you give presentations at schools.

Professor Sir David Spiegelhalter: It comes back to data literacy. Having a deep scepticism about everything you see in the media is the first thing, and teaching people the tools to deconstruct arguments that claim to be based on evidence—“The numbers show ...”, and all that sort of stuff. There is a set of principles that you can go through to decide how reliable something you are hearing is. This comes back to Baroness O’Neill’s point about trustworthiness. I do not like to think of educating people, but you just need to empower people to check on someone’s trustworthiness.

Baroness Bakewell: A lot of adults need to know exactly those things too and they do not exercise them, do they?

Professor Sir David Spiegelhalter: No. I would hate to think that it is too late.

Professor Peter McOwan: If you teach them in schools, by the time they are adults hopefully they will remember them.

Baroness Bakewell: Should people be worrying about jobs or the

changing landscape of employment generally?

Professor Peter McOwan: The one thing that we can be certain of is uncertainty in the future in the jobs market, and not necessarily just because of artificial intelligence technology but because of a whole range of other things. When I go into schools I talk to young people about flexibility and transferable skills—the ability to reason, to be entrepreneurial and so on. Look at some of the work that is coming out of think tanks. My favourite piece of data was presented at the previous session on AI. It was the percentage of jobs to be lost in constituencies depending on which Member of Parliament it was, which suddenly got people very interested; it was potentially up to 30 per cent to 40 per cent.

My feeling is that it will probably not be as bad as that, because there will be new jobs opening up. They will require you to be flexible in order to be able to move into them and to retrain. That shows the importance of education, but then I would say that, wouldn't I, as a university professor? I also think that the threat of artificial intelligence is likely to go in a direction that is not quite the way we imagine it. In the future, we could see large areas of beneficial blending, bringing together what AI does well and what human intelligence does well to work together.

A very brief example is an artificial intelligence system that I have been working on, which can analyse the way people look at images, for example, rather than having to pay lots of people to sit with iTrackers on. When I tried to commercialise that I was told that it would not go anywhere because if I undermined it and said to all these designers out there, "I have something that can replace you", it would get no traction. It becomes part of their computer-aided design, so they do all the clever stuff and it tells you this and it becomes a non-human in the loop. That is a likely direction that things will go.

Professor David Edgerton: We would do a disservice to the public if we told them that AI was going to be the major determinant of the employability of their children. In effect, we did that with IT, with space rockets and with aeroplanes before that. Each individual new technique will have a small effect in relation to the total economy. It is to misunderstand the nature of our society to assume that one particular technique will have a transformative effect that is out of proportion to all the other techniques.

Baroness Bakewell: But if you have automatic cars, it will put all taxi and lorry drivers out of work.

Professor David Edgerton: If we have them: that is an argument that people make. I can see lots of arguments for not having automated cars. I can see lots of arguments for buses and trains and for walking and cycling. That is my point, in a way. If we start with the assumption that we will have nothing but driverless cars, we will do a disservice to the public's understanding of what the future is likely to be.

Could I add another point, about politics and the economy? Brexit will probably have a much larger effect than AI on the employment possibilities that our children will have. Do we warn them about this in

schools? The price of oil or whatever might be more important in any particular situation.

Q220 Lord Puttnam: I rather agree with you. I am far less concerned about the impact of the productisation of AI than I am about the notion of the productisation of people. This troubles me greatly and it touches on the next question, I am sure, so I am getting ahead of Stephen. One thing that makes me unique here is the fact that I spent the first dozen years of my working life in advertising, and by the time I left I had no illusions whatever that if you offer advantages to people in advertising to find out more about their customers, they will take them. It is the misuse of data that troubles me very greatly.

Talking about people being informed, I like watching soccer, so I have a Sky set. It is quite specific, but the fact that by watching soccer the data on all my family's habits is being collected centrally and can be turned around and used in a variety of ways leads me to the whole issue of the misuse of data.

I jotted down an example. Why do I trust Which? but mistrust TripAdvisor? It is not just because I have grown up with Which? and I know that Which? is controlled by the National Consumer Council and that TripAdvisor does its best to manipulate me and push me by the number of sites it has to make decisions. This is about government. One of you touched on this. What is the role of government in ensuring that we are not just aware but are protected against the misuse of data that is being collected without our knowledge?

The Chairman: There is a competition to answer this, clearly.

Professor Sir David Spiegelhalter: You can see that we are racing to answer this. This comes into the whole area of data governance and organisations, which is in a state of flux at the moment in this country and elsewhere as to the role of regulation in data. As I said, I do not think it is up to every individual to have to protect themselves against this; it needs a higher regulatory framework. There are various bodies at the moment. The Information Commissioner's Office is doing very well but could be strengthened even further. There is the Nuffield Convention on Data Ethics and Artificial Intelligence, which has just been set up with the help of the Royal Statistical Society, and there is also the proposed Centre for Data Ethics and Innovation. There are various bodies, and I am not going to say what the right balance should be exactly or how it should be done. There is a massive need to grasp this and to develop appropriate governance and regulation. It has to be debated.

Lord Puttnam: I have a silly example from my own background. I started in advertising in 1957. The big book that year was Vance Packard's *The Hidden Persuaders*. You will remember that very well indeed. The issue that cropped up was the whole business of subliminal messaging. Within a year the IBA moved to be very clear about what you could and could not do on television in terms of speed of images. Whether they overreacted or underreacted is something that you will probably have a view on, but the Government moved like lightning. It was real. Do we still have the same mechanism and the will to use our agencies to move against the misuse of data?

Professor Sir David Spiegelhalter: Where it is happening at the moment and has happened for ages is in the setting of insurance premiums. There is very strong regulation, and the only genetic condition that you are allowed to ask about is Huntington's disease, so far as my understanding goes. You cannot ask about ethnicity. You can no longer ask about gender for car insurance. All sorts of rules have been put in where the companies would love to have more data because they could do finer stratification for the premiums, but it is not allowed.

Professor Peter McOwan: Lord Puttnam's point is important, because collecting that data allows them to produce services that could be very positive to you. It is a question of whether what is being done to you is something that you enjoy having done to you or something that you do not. That makes subliminal messaging very difficult. You chose an interesting example, and its efficacy is another thing that can be very strongly argued, but it has moved on very rapidly. It was a very contained, solid block of things, and it said that subliminal images meant that people would work like zombies and just buy whatever it was. You can ban that, because it is clear that that is what it does.

Recombinant data and data fusion from multiple different sources in and of themselves are not necessarily evil things, but they can be used for evil things. That can make it very difficult. As Professor Spiegelhalter said, data governance is important and you cannot block particular things at particular points and say, "You are not allowed to do this". You open up a broader situation when it comes to trans-border transfers, because there are cases of data havens where things that cannot be done under our data protection legislation can be transferred to other places and processed there and potentially brought back again.

Professor David Edgerton: On this point about the will or lack of will, there is the general phenomenon of a lack of will to control private corporations and interests in the public interest. We have seen, for example, a decrease in the quality of television programming over the last 20 or 30 years that is consequent upon that.

Lord Puttnam: You believe that there is a public interest element to this, which is certainly what I believe, that is not being exercised.

Professor David Edgerton: We know how to establish mechanisms to ensure that the public get quality of information. We have decided not to use those powers for all sorts of reasons with which we are very familiar, but we could return to a situation where we do.

Lord Puttnam: Thank you very much.

Viscount Ridley: Can I come back to the point that Professor Edgerton was making a few moments ago about the risk that we get carried away and prepare children for a future that might never happen? Can you give a specific example, maybe going back to that Harold Wilson speech you read from, of where we got a bit overexcited about a technology and built something into the curriculum or into the way we reskilled people that turned out to be a bit of a bust? Did we train far too many nuclear engineers in the 1960s, or did we miss the integrated circuit?

Professor David Edgerton: It could be argued that too many scientists and engineers in general were trained, given the number of jobs there were for them. Scientists and engineers assume that there is a desperate shortage, but the fact that so many go to the City suggests rather an oversupply. Certainly in the realm of public policy, all sorts of assumptions drove a particular technical development into what turned out to be a dead end, at least for the United Kingdom if not for the world as a whole. There was massive overinvestment in British atomic energy, for example, and overinvestment in supersonic transport. I think the United Kingdom would have been better off if Concorde and the advanced gas-cooled reactor had never come along.

These enthusiasms for very particular techniques have their negative side. One should not assume that it is all about exciting people to get into science. It leads to certain public policies that can reduce national and human welfare more generally. It is not a cost-free exercise to hype one technique over another.

Q221 **The Lord Bishop of Oxford:** The Government have recently announced a centre for data ethics and innovation, which will aim to lead in that area. What would be a good ethical framework for the development and deployment of AI, not only to build public trust, which we have covered quite well, but for a flourishing society? I observe in what you have said so far that you have been very strong on teaching people deconstructive and critical skills and building a strong hermeneutic of suspicion into children in particular, but you have not been particularly strong about appreciative inquiry, about identifying what is good and what makes for a flourishing society so that we not only avoid the dangers but build good qualities into our common life.

Professor David Edgerton: I thought that I addressed that point earlier. It is crucial. We need to get back into a position—we were there partially in the years after 1945—where we believe that we can take a collective view as to how to improve society and act on it collectively. There were very good reasons why many people began to reject that sort of approach, and I recognise that, but there was value in it and we ought to empower ourselves as a collectivity to think through the kind of society, the kinds of machines and the kinds of techniques that we want, and to decide, indeed, who should have control over those techniques.

Professor Peter McOwan: I agree. One has to take a utilitarian approach to this, because there are many positive and negative aspects, and it requires a decision to be taken on what would be allowed and what should not necessarily be allowed to create an area between those. That can only come through discussion. In the same way, if as a scientist I wish to do a particular experiment, I need to go through an ethics committee, and if that ethics committee has lay members as well as expert members, putting structures in place that give some confidence about the data, that would give us a slight idea at least of the red buttons that we should not be pressing.

Professor Sir David Spiegelhalter: There will be the centre for data ethics and innovation, and you are right that we have slightly emphasised the critical part, the appraisal part, which is where a lot of the ethical

debate comes in. The other side of what I would call data literacy is the innovation side, which has enormous excitement and benefits. Again, kids should be doing this stuff in primary schools. Building neural networks out of toilet rolls is so exciting and such fun.

Professor Peter McOwan: They get excited and come up with new ideas.

Professor Sir David Spiegelhalter: They are collecting the data. They are not just passive recipients of stuff on their phone. They are actively going out and innovating and constructing. It is never too early to introduce that. That should just happen anyway. I quite like the fact that it is ethics and innovation, because they are two sides of the same coin. One is to do with production and the other is to do with the control or criticism, and you need both.

Professor David Edgerton: Talking about things ethically can be a way of not talking about them economically and politically, and I think that would be a serious mistake.

Q222 **Lord Swinfen:** How should we be thinking about the international narratives about AI? There has been much discussion about an “arms race” emerging between nations developing and utilising AI. Is this a useful way of thinking about AI?

Professor David Edgerton: Yes and no. When we think about technical advances, we switch between a very global vision of a new technique affecting all of humankind to a very nationalist vision, without warning. We need to be very careful about that. A lot of talk about AI focuses, for example, on what the United Kingdom can get out of AI economically, in competition with other countries that are also developing AI. Sometimes those discussions about what the United Kingdom can get out of it are the same discussions that are happening in the world as a whole. The reality is that there is an element of competition, and the United Kingdom might be able to use AI without being an innovator in that area and, indeed, without exploiting British innovation in that area.

There are many separate discussions to be had about the relationship between the nation state and AI, and they should not be conflated, as they often are. At the moment, for example, it seems that the Government are looking to diverge from EU regulation precisely in these areas, I assume in the hope of getting an economic advantage. There are some very hard questions to be asked about what actual advantage the United Kingdom has vis-à-vis other parts of the world in AI that are probably worthy of serious examination.

At one point, I looked at biosciences. I found the claim that the United Kingdom had a unique strength in biosciences that had to be exploited to be not very strongly supported by the evidence. The reality is that there are five or six rich countries that are on roughly the same level when it comes to biosciences, and I would guess the same is true of AI, although I do not know in that particular case. One must be wary of the stories that we are increasingly told about these very particular British strengths waiting to be exploited.

Professor Peter McOwan: I agree. The whole issue is not simply about artificial intelligence but about materials research and bioscience and a whole series of things that can be considered an arms race with other countries. AI is one of a mix of technologies that, if they make their way into the marketplace, could change the way we buy and consume products. It is important to see that there is nothing special about AI; it is just part of the spectrum.

Professor Sir David Spiegelhalter: I do not think I can say much about this, except that we have to be aware of the enormous dominance of Facebook, Apple, Google and Amazon in this area. We cannot ignore it. That has to be taken into account in any thinking about our role in our national identity.

Lord Hollick: Professor Edgerton, you talked about certain areas of national strength, some of which may be exaggerated and some of which may not be fully appreciated. I would be interested to know what particular strengths in AI you think the UK has.

Professor David Edgerton: I could not answer that question.

Lord Hollick: Can any of the other panellists?

Professor Peter McOwan: We have a very strong machine-learning community—so much so in fact that, to pick up Professor Spiegelhalter’s point about Facebook and Twitter and so on, they are headquartering here in the UK because they can absorb elements from the UK university system. A colleague of mine who replaced me when I moved to become vice-principal got kidnapped by DeepMind in a very large bag with a very large pay cheque and we never saw him again.

The Chairman: So is it a threat or an opportunity?

Professor Peter McOwan: It is a combination of both. I am not going to come down on one side or the other. It is a threat to the UK from the point of view that it is very difficult to recruit lecturers in those particular areas, because they are all being picked up by these large companies if they are worth their salt, which means that it is difficult for us as a university sector to keep developing these areas. It is also an advantage in that we now have links with these larger companies, which is useful from the point of view of exploiting their knowledge base plus the spread.

Lord Hollick: Does it concern you that these companies are being bought by Facebook, Apple and Google?

Professor Peter McOwan: They tend to do them very quietly.

The Chairman: It is a real question.

Lord Hollick: Are we in danger of becoming a vessel state in the digital world?

The Chairman: I do not know where we have heard that phrase before.

Professor Peter McOwan: There is a possibility, because these are such large multinational organisations. One thing that continues to impress me is the fact that there are still a lot of very creative young academics coming into the system who are coming up with fantastic new ways of doing things. The point is whether at the end of that you can keep them

on. Sometimes you will, because the academic lifestyle is different from working with one of those large companies. We have fewer hammocks and we do not pay them as much, but there is a certain level of freedom within it, and a lot of people enjoy teaching the next generation of young minds. There are ways of keeping people in the university sector on the AI side of things, but universities these days are more interested than they have been in the past in co-working with these larger companies as part of the Government's industrial strategy and so on, and that is going to be incentivised.

Lord Hollick: Do you think that as we separate from the European Union that will make the UK less attractive, or will it not change the UK's attractiveness for people from Europe and the rest of the world to come to work here?

Professor Peter McOwan: I promised not to say anything about Brexit when I came here because it is such a contentious issue. However, if I must, what concerns me somewhat is that there are a lot of discussions about the pros and cons of Brexit, and in those discussions you very occasionally get a snippet of "Blah, blah, blah ... and science". It is all about trade and industry. That is incredibly important, but we as the UK were very well represented on things like Horizon 2020. I led a seven-year project on social robotics with people from across Europe, who needed to be there because we also looked at cultural differences. There are cultural differences in different countries in the acceptability of robots and AI. We might lose some of that, particularly the freedom of movement of academics, although hopefully, as the Government have mentioned, they will have visas that allow them to come in. Whether they will want to do that, only time will tell.

The Chairman: Thank you very much. A final question from Baroness Grender.

Q223 **Baroness Grender:** We have asked every witness this question at the end of every session. We want one specific recommendation from each of you that you think we should make in our report. Given that this is most likely the final time we do this, we are sure there will be an absolute cracker from each one of you before Christmas.

Professor Peter McOwan: So no pressure.

Baroness Grender: Absolutely none. Can you give a single recommendation about what we should recommend?

Professor Sir David Spiegelhalter: The public are not the problem; elite understanding of technical change and its impact is.

Professor Peter McOwan: I would say the same thing. The public are not the problem. The last research excellence framework showed that we have over 10,000 three-star and four-star researchers in this country. If we can reward the people who are researching artificial intelligence and motivate them to go out and not just disseminate but co-create with people to put it into the curriculum in schools and to work hard there, that will make a significant difference. To give one specific example, where people apply to Research Councils, now the overarching UK Research and Innovation, there are six words on the application form that

Professor David Edgerton, Professor Peter McOwan and Professor Sir David Spiegelhalter – Oral evidence (QQ213–223)

can change the country, and they are “must contain significant evaluative public engagement”.

Professor Sir David Spiegelhalter: My recommendation is what I said before. It should be empowering a new generation with data literacy, and that includes the ability to both innovate and critique.

The Chairman: Fantastic. That was very succinct. You can see that we did not want to let you go because we have had such an interesting session. Thank you very much. I have learned two new phrases: fishbowl transparency and recombinant data. Thank you very much indeed. Every little bit helps.

Fujitsu, Ocado and Dyson – Oral evidence (QQ 105–115)

Fujitsu, Ocado and Dyson – Oral evidence (QQ 105–115)

[Transcript to be found under Dyson](#)

Future Advocacy, Professor Dame Henrietta Moore and Professor Richard Susskind OBE – Oral evidence (QQ 95–104)

Evidence Session No. 11

Heard in Public

Questions 95–104

Tuesday 14 November 2017

[Watch the meeting](#)

Members present: Lord Clement-Jones (The Chairman); Lord Giddens; Baroness Grender; Lord Hollick; Lord Holmes of Richmond; Lord Levene of Portsoken; The Lord Bishop of Oxford; Viscount Ridley; Baroness Rock; Lord St John of Bletso; Lord Swinfen.

Examination of witnesses

Olly Buston, Professor Dame Henrietta Moore, Professor Richard Susskind OBE.

Q95 **The Chairman:** Welcome to our 11th evidence session. A very warm welcome to Olly Buston of Future Advocacy, Professor Dame Henrietta Moore of the Institute for Global Prosperity, and Professor Richard Susskind OBE. I have a little rubric I need to go through before we plunge in, which will no doubt bore my fellow Committee Members senseless, but I will carry on. This session is open to the public. A webcast of the session goes out live as is and is subsequently accessible via the parliamentary website. A verbatim transcript will be taken of your evidence and will be put on the parliamentary website. A few days after this evidence session you will be sent a copy of the transcript to check for accuracy, and we would be grateful if you could advise us of any corrections as quickly as possible. If, after this session, you wish to clarify or amplify any points made during your evidence, or have any additional points to make, you are very welcome to submit supplementary written evidence to us. First, perhaps each of you would like to introduce yourselves for the record and then we will begin with the questions. Olly, would you like to start?

Olly Buston: Thank you very much. I am Olly Buston. I am the CEO of the think tank Future Advocacy. We focus on maximising the opportunities and minimising the risks of artificial intelligence in the UK and around the world.

Professor Dame Henrietta Moore: I am Professor Henrietta Moore. I am the director of the Institute for Global Prosperity at University College London. We focus on envisaging a future flourishing of human societies.

Professor Richard Susskind: I am Richard Susskind. I wrote my doctorate in the 1980s in AI and law. I am president of the Society for Computers and Law. I am strategy and technology adviser to the Lord Chief Justice and I chair the Advisory Board of the Oxford Internet Institute.

Q96 **The Chairman:** Thank you very much indeed. I am going to start with a fairly broad, general question. How is AI impacting work today and how will this change over the next 10 to 15 years? We will include within the broad scope of that the kinds of jobs that will be lost as a result of AI adoption, the kinds of jobs that will be created and the sorts of skills that you think will be needed in the future and, if you think it is relevant, reference to the Government's recent AI review. Would you like to start, Olly?

Olly Buston: There has been a lot of research recently about the kinds of jobs and, more recently, tasks that are susceptible to automation. The assumption is that repetitive tasks which take place in a predictable environment are most susceptible to automation. The Industrial Revolution was all about the automation of physical tasks, and what we call the intelligence revolution is about the automation of intellectual tasks. A lot of research about automation and AI is based on the classic Frey and Osborne report. Their methodology has been improved by various people—the OECD, various accountancy firms, such as PwC, and others. Part of that has been about a greater focus on tasks, so becoming a bit more granular about what is being automated.

The McKinsey report suggests that if there are no more developments in AI capacity, even with existing technology a very large proportion of current tasks could be subject to automation, and 60 per cent of all occupations contain more than 30 per cent of tasks that are technically automatable with current technologies. If you add on the assumption in Moore's law, with the vast investment going into AI by companies, and indeed militaries, you would expect an acceleration of this process. The Google CEO, Sundar Pichai, says that the tech giant is "thoughtfully" applying AI across all its products, and IBM CEO, Ginni Rometty, said that her organisation is "betting the company" on AI. There is huge investment going in.

There is quite a lot of agreement that a lot of jobs will be lost. There is much less agreement whether more jobs will be created and what the net impact will be. There is also a lot of agreement that there will be extraordinary disruption and that impact will be felt unevenly across different groups. The piece of work we did, published last month, looked at automation in individual parliamentary constituencies in the UK. We found that the range was between 20 per cent and 40 per cent per parliamentary constituency of the proportion of jobs that were at high risk of automation by the early 2030s. There is a huge amount of change coming. The Government's recent report was excellent in thinking about how to maximise the opportunities. However, I think too little has been done by government—or any political party, in fact—in trying to think about the downsides and how we might ameliorate those and support the

people, families and the regions most impacted. Perhaps I will stop there because that was quite a long contribution.

The Chairman: Thank you. Obviously, you have been pretty granular in looking at the risks, and so on. Dame Henrietta, do you accept that that methodology is valuable, or do you have a different way of approaching it?

Professor Dame Henrietta Moore: No, I think that methodology is very valuable, but there is considerable uncertainty. Part of the problem is that in many areas we do not have good evidence on which to base predictions for the future. Extrapolating from the past in this area is likely to be unwise, let us say. We face great disruption and increasing social inequality. The changes we have already seen in the United Kingdom as a result of information technology—GPS and smartphone technologies—providing new kinds of platforms mean we have at least 10 million people in very insecure work. One of the problems we are facing at the moment, which is likely to get worse in the future, is that the quality of work rather than the number of jobs is something we need to attend to. That is a very important point. We are looking at a situation where quality of life and the ability of people to lead a larger life will involve some kind of new social contract, and that needs to come into place at the same time as these transformations are pushing through and not after we pick up the pieces later.

The Chairman: Professor Susskind, obviously you have focused quite heavily on looking at the impact on white collar, professional jobs. Do you recognise the analysis put forward by Future Advocacy, or do you see a different overlay?

Professor Richard Susskind: I agree with much that has been said, particularly that the past is not much guide to the future. William Gibson, the science fiction writer, once said that the future has arrived; it is not evenly distributed yet. I think that is right. If you look right across industry, for example at self-driving cars, in one way or another—I do not know whether it is five, 10 or 15 years—we know this is going to have a fundamental effect on 10 per cent to 15 per cent of the community. The interesting question about timing is that the way you frame it is significant. I like the quotation from a Silicon Valley chap who once said, “The best way to predict the future is to invent it”.

You ask the question as though, having experts along, we can give you a good idea of how it is going to unfold. We are all in the foothills and have some sense of likely development, but the pace of change is nothing to do with us. In fact, it might be, in large part, to do with the Government of this country and our industrial strategy. I would want to say, as a key point, that the future is not in some sense out there and some of us see more clearly how it is going to unfold; the future is far more malleable than that. Our challenge—those of us who feel passionate about the strengths and the threats here—is to be shaping rather than be shaped by this future.

I would say as a generality that many of the short-term claims made of AI are greatly exaggerated. On the other hand, I think the long-term predictions about AI are understated. That is to say, I do not think, as

some people do, that over the next few years we will see major unemployment across the professions, and so forth. But I think, as we get into the 2030s and 2040s, we will be living in a world where it is hard to imagine there will be the same level of employment available.

The term “AI” is problematic. My son, Daniel, and I have written a book called *The Future of the Professions*, and we use the phrase “our systems and machines are becoming increasingly capable”. Every day one hears of a new development, whether it be an app or a breakthrough in some AI lab. The point is this: historically we used to regard machines as good at undertaking routine tasks, so we looked at jobs in the workplace where it looked routine—it might be in the factory or the office—and said, “Yes, intuitively we can imagine that systems and machines will take those over, but surely when a job requires creativity, judgment or empathy”—the argument runs—“this is the distinctive capability of human beings”.

The first observation, if you will give me a couple of minutes to expand on this, is that we should not be thinking in an 1980s way about AI, which was very much my history, where an AI system was, essentially, a rule-based system and we wrote very complex decision trees to somehow replicate the jobs of human beings. We are way beyond that now. We now have machines, whether it be in game-playing, medical diagnostics, architecture, law, audit or tax, that are outperforming human beings. I am not saying this is pervasive but I say the future has arrived; it is not evenly distributed yet. We are seeing illustrations now which give us a glimpse and my guess is that in the 2020s many of our current ways of working will be fundamentally displaced.

My take on this is that the challenge we face is whether or not we compete with these systems or build these systems. By 'compete' I mean you might say, “I hear everything all the AI people are talking about for the future but I still think there’s a wide range of distinctive areas of work for human beings and that is where our industry and education system should focus on”. I take a different view; I think we are at a time of immense opportunity and the great opportunity for individual businesses and for careers—frankly, for the country—is to build a range of systems that will replace conventional working. Frankly, if we do not, others will.

The Chairman: Thank you. That was a very complementary set of responses. I have Lord Swinfen, then Viscount Ridley and Lord St John.

Lord Swinfen: Do you think schools are doing enough to prepare children for the modern, computerised workplace?

Professor Dame Henrietta Moore: No. I think that is very evident. There are a number of different areas here. One is that we seem to be doing remarkably little to teach young people about issues to do with ethics, data and privacy, which will be crucial in the future. We need a different way of approaching that set of questions, and we have done little about that. The second is that we have not thought about how we are going to teach children in schools in a way that allows them to move through a whole process of lifelong learning. We are still thinking in silos. We are still thinking about school, then university or college, then work and then something else, instead of thinking about how, in the future, we will have to make the boundary between learner and worker very porous.

You will not be going through school or university once and going into the workplace, learning on the job and going forward in the way we have thought of in the past.

The Chairman: I am going to stop you there because we are going to come on to that question later in a more specific way. I will bring in the others on that particular question later.

Viscount Ridley: Two of you have said that the past is no guide to the future. While of course that is true, none the less we should pay attention to the fact that there have been bloodcurdling warnings about the impact of automation on work ever since the threshing machine. Norbert Wiener, in 1949, said that putting computers in factories would usher in “an industrial revolution of unmitigated cruelty”. In 1964, the commission to President Johnson said much the same thing. In fact, as we have seen, automation has actually increased employment throughout this period. Why should it be different this time? As Herbert Simon put it, “The bogeyman of automation consumes worrying capacity that should be saved for real problems”.

Professor Richard Susskind: It is a question everyone in the AI field is debating. Is it different or are we just jumping up and down, seeing a repetition of a fundamental change, and new jobs will be created and we will all be safe? We looked at this in depth across eight different professions, from neurosurgery through to architecture and even the clergy—some areas one would think are not amenable to automation. Why is it different? The reality is that machines are becoming increasingly capable. They are taking on more and more tasks that, historically, we thought could be undertaken only by human beings. But new jobs will arise. The difference here is that machines will take on these new jobs as well. We have seen nothing in our research to suggest that the new jobs—and there are many new roles emerging—are roles for which human beings have comparative advantage over machines in the very long run.

Let me be clear: the 2020s will be the decade of redeployment rather than unemployment. There are huge employment opportunities to be involved in the development of the systems that will replace our old ways of working. Once we go beyond the 2020s, to the 2030s or the 2040s, it seems to me that as our machines become increasingly capable they will match human beings in cognitive capacity—I do not mean they think but in terms of problem-solving and reasoning—and in psychomotor or manual capacity, that is to say physical capacity; even in the area of affective computing, with machines that can both detect and express human emotions. If you look at the wide range of talents human beings have, in one way or another, machines—and it tends to be in different ways—can deliver the outcomes of human capacity, usually to a higher level. We are already seeing that now.

The difference in the earlier eras was we moved from a stage, for example, where people who were involved with physical tasks were perhaps replaced by machines and a whole bundle of cerebral activities emerged. Again, the point is that there is no evidence in any of our research to suggest that many of the new jobs emerging are, in the long

run, jobs for which human beings have an advantage over machines. It is to take too narrow a view of the nature of this change to think that, somehow, there will be a whole set of jobs that are the exclusive preserve of human beings. There is a moral question. We might say we want to draw limits on what machines ought to be able to do, but in terms of their capabilities, even if there are no further transformations—and I am of the view there will be many more, if we extrapolate from the technology we have today—it seems to me it is very hard to avoid the conclusion there will be far less for human beings to do.

The Chairman: Dame Henrietta, do you agree?

Professor Dame Henrietta Moore: This is a contentious point that everyone is debating at the moment. I certainly think, absolutely, that in the next 20 years lots of new jobs will be created. They have already been created in the last 20 years with the degree of automation we have already had in AI, and so on. The more worrying question is what will we do? These machines should improve productivity. At the moment, productivity is dismal, so these new machines coming in should improve productivity. That will create more wealth. What will we do with that wealth? How will we invest it wisely for the future? When we keep talking about jobs, and jobs in the quantum, we are missing a big area we need to think about around the job issue: what will we do with the wealth we create? Will we carry on doing the kinds of things we have done with it in the last 30 years, or will we do something different with it?

A job is only one part of the way of thinking about how you might conduct yourself in this period. For example, with the new forms of artificial intelligence we have, the whole question of the relationship between work and leisure comes under a degree of erasure. The notion we have is that you are either working or you are not working, or you have a job or you are idle, or you have a job or you are unemployed. This is a very 19th century view of the job. Modern jobs of the future will not be like that; they will be much more distributed. In a way, we can talk about jobs but we would be better talking about the character and nature of work and its capacity for creating value for society.

The Chairman: Thank you. I am going to give Olly the chance to answer this as well, because this is a really important question.

Olly Buston: The first point is that I would not be quite so dismissive of the profound human, social and political impact of previous rounds of automation. Arguably, fascism, World War II and communism could be seen, to some extent, as the playing out of the great inequalities created by the Industrial Revolution. It took a while for our societies to catch up in terms of distributing that wealth and creating a fairer society. It is the uneven impact of this wave of—

Viscount Ridley: I am sorry. What was fair about the 18th century? It was a very unfair and unequal society before the Industrial Revolution.

Olly Buston: I am not a historian of the 18th century, but I would suggest that perhaps there was an equilibrium that was massively disrupted. Even though it was not fair, it was disrupted and it took a while to create.

The Chairman: I will leave you to debate offline, if I may.

Olly Buston: That may be quite a weak answer.

The Chairman: Perhaps with the assistance of Lord Giddens.

Olly Buston: Again, the uneven impact of this wave of automation could have profound human, social and political consequences. One thing we have found is many of the constituencies with the highest proportion of jobs at high risk of automation are in the north and the Midlands—areas that have suffered previous rounds of industrial decline. There are social and political implications of that. You could suggest that the Trump election and Brexit were, to some extent, a reaction against automation as much as against globalisation. There are profound social impacts here and it is partly because of the unequal impact, even if the net effect is greater economic prosperity.

On the question of whether there will be more work in the future, potentially there is a profoundly optimistic future in which machines do more work, humans do less work and we are able to focus on the things that make us more human and we flourish. I feel that the line of least resistance and the way our societies are currently structured is probably not that. The path of least resistance is probably greater inequality, greater unemployment, and people's lives being undermined by unemployment, the loss of sense of purpose and communities damaged. We need a huge human, social and political effort, I think, to get on to the optimistic scenario which is achievable if our great politicians take us there.

The Chairman: You can certainly include us in that.

Lord St John of Bletso: I am sure my Lord Chairman wants to move on to other questions, but I will make one brief observation and ask a question. Clearly, machine-learning robotics has and will continue to have a major impact on job losses for blue-collar workers, and greater efficiencies. To what degree will it impact on white-collar jobs? Professor Susskind, you mentioned it earlier on but I am not that clear as to what impact it will have on the white-collar market. I understand the blue-collar market impact.

Professor Richard Susskind: My suggestion is, and the research supports this, that it will have as pervasive an impact on white collar as on blue collar. The technologies that challenged and replaced much of blue-collar work did not need to be as sophisticated as the technologies that will replace white collar work. I was speaking recently to 2,000 neurosurgeons and we were talking about this. You would think that neurosurgeons are the very peak of the professions. You would think they would not be susceptible, but they were entirely open to this. The point I made to them was that in the end people do not want surgeons, they want health. In the end, people do not want lawyers, they want justice. In the end, people do not want tax advisers, they want to minimise their exposure to tax liability, and so forth. The point we are seeing emerging in white-collar work is not simply that the technology is used to turbocharge or change the way we work; it allows us to think of entirely new ways of solving real problems.

We tend to think of the future of white-collar workers in terms of how we can preserve and enhance the way we have always worked. The deeper question is: to what problems do white-collar workers provide the answers? It is health, education and so forth. Technology is enabling us to answer old questions in entirely new ways. If you think, "I can't imagine a lawyer or a doctor being replaced", that is wrong. That is what Daniel and I call the "AI fallacy", which supposes that the only way we can match the performance of white-collar workers and professionals is by somehow replicating or copying them. The profound change we are going to see is the outcomes we require will be delivered in entirely new ways; for example, surgery that is non-invasive. In 50 years' time we will look back and say, "It's unbelievable that we used to cut bodies open". It is not, as most people would want to discuss, robotic surgery in the strategic long term; it is non-invasive procedures.

It is the same in the work we are doing in law; it is not about solving problems but about avoiding problems. It is the same in architecture, and so forth. We have to have a different mindset in thinking about the future, think less about how the technologies can copy our brilliant people today and think about whether there are new technologies that can deliver the outcomes we require in new ways. The exciting thing is the jobs that will be created in developing these new technologies. The tone of most of the questions—and I entirely understand it—in all these debates tends to be negative. I am hoping for British industry, for our educational establishments and the health of our nations, that we see this as a remarkable set of opportunities. We have some of the finest researchers, thinkers and workers—and, I hope, fledgling companies—in this area. We should be wondering how we turn this to our advantage rather than wondering what barriers exist.

The Chairman: Thank you. I am going to make apologies to the Lord Bishop, Lord Giddens and Baroness Greener otherwise we will spend all our time on the first question. I will move to Lord Levene.

Q97 **Lord Levene of Portsoken:** Perhaps I can start with Professor Susskind. In your opinion, how robust are predictions about job losses and job creation as a result of AI and automation? What methodologies are normally used?

Professor Richard Susskind: As I always say, there is no evidence from the future. This has been covered to some extent by Olly, who pointed out, whether it be the McKinsey study, the Oxford study or a whole bundle of other studies, that generally people look at the work that is undertaken today, the tasks involved, and think, "To what extent do these tasks require creativity, or are they routine processes?", and they break them down and say, "That percentage of tasks could be undertaken by machines we have today, but it is hard to imagine how these machines could undertake these tasks". I think this is fundamentally flawed, for the reasons I just mentioned.

What we are thinking about when we analyse work in this way is: how can we use technology to replace the way we currently work? The more interesting question is: how will technology allow us to work entirely differently? I have to say I find many of the predictions about job losses

to be entirely unreliable. They are based on a primitive conception of how we work today and how we might work in the future. Job creation is even harder because we can identify, in broad terms, that we will need people to develop systems and to analyse processes and we will need knowledge engineers, and so forth. As I keep on saying, we are very much in the foothills just now, so a lot of the quasi-science that many of the major consulting firms and so forth are putting out is not rooted in deep theoretical foundations. We are all feeling our way here. I do not want to overstate. My son, who is an economist, is always holding me back from predicting in terms of years. We do not know if it is going to be five, 10 or 15 years. We are feeling our way. I go back to the earlier point: the point is not to try to predict; it is to try to shape it.

The Chairman: Dame Henrietta, do you agree?

Professor Dame Henrietta Moore: Yes, I do agree substantially with that. We do not know what the jobs of the future will look like. The more worrying thing is that although the work we have in front of us at the moment and the predictions coming out are what I would call “evidence light”, I still think we need to work with them and do them because that is what we have in front of us. We need to take some steps here, these steps have to be incremental and we have to move forward with what we have. It is more important for the United Kingdom as a whole that we link this kind of discussion to what we imagine the new industrial strategy will mean. It has to mean that we will have to address what happens in specific places. This is where I think Olly’s work has been so important. It is not that everybody is going to do everything, but we have to have something for everybody somewhere. We have to recognise that we have huge parts of the country which are very underproductive at the moment. One of the things we ought to be thinking about, as these new technologies come on stream, is how can we use them to bring those parts of the country into that more productive space. This requires critical thinking, which has not yet been done.

Olly Buston: I think Richard and Henrietta are both absolutely right that we are feeling our way, but it is very important, as you said, to feel our way, it should continue and we should try and feel our way better. Professor Susskind is right that the more interesting question is how we can build a totally different society. I think it is still a very important question to ask what is going to happen to our current society. If you are a long-distance driver, 55 years old and in Texas, that exciting conversation about creating a new society may not feel very relevant but the more immediate, although not very immediate, issue of what is going to happen to your job is still very important.

The research we did takes this feeling on its way and looks at a more granular level, which introduces further uncertainty. We are talking about smaller sample sizes. A parliamentary constituency may have 50,000 or so jobs. The evidence is less accurate there but the questions are important to ask. We found, for example, that John McDonnell’s constituency, Hayes and Harlington, has the highest proportion of jobs at risk of automation. However, it is meaningless to stop there. The point is to ask a set of questions about why that might be and to robustly challenge that assumption. Why that was is because Heathrow is in that

constituency, so there is a belief that there are a lot of jobs that coalesce around an airport that may be automatable—be that warehousing, driving and so on. It is important to get more granular but important to be humble about the level of certainty. It is important to ask these questions. We want to bring together companies, politicians and others at a local, community level and have a proper dialogue with the people making the hiring and firing decisions and making the decisions about whether they are deploying AI or not. Having those local-level discussions in individual constituencies is the next step to get a more accurate sense of what the future might look like.

Baroness Greener: We have had witnesses who have said to us that empathy and caring will not be able to be replaced by a machine, and therefore you could see a future where the status of those kinds of jobs is raised as other jobs fall away. What do you think of that?

Professor Richard Susskind: I have written a lot about this. It is absolutely clear, certainly in the way the professions, for example, are currently constituted, that empathy is absolutely vital. There is an area of computing which you, I hope, have been exposed to known as affective computing—machines which can both detect and express human emotions. A machine now more accurately than any human being can look at a human smile and tell whether that smile is fake or genuine. A machine more accurately than any human being can listen to two female voices and tell whether or not they belong to a mother or a daughter. We have to acknowledge, it seems to me, on the basis of data rather than any feelings these machines have, that in terms of recognising emotions machines will probably outperform us as well. It seems entirely conceivable, in terms of responding, that they will—again, through machine learning, massive bodies of data and past experience drawn from human beings—know appropriate ways to communicate. Will that machine feel the same way as the psychotherapist or the GP? Almost certainly not.

I would like you to think for a second about Japan, where they have a severe shortage of nurses. In one hospital, at least, they have a robotic nurse in every room connected to the patients' gowns wirelessly. The system can detect, through agitation, sweating or noises emitted by the patient, when some kind of response is appropriate. It might be a comforting form of music. Often the question here is not: is that better than the wonderful nurse that, say, my wife is? The question is: is that better than nothing at all? When we are thinking, for example, of the long-term care of the elderly, I am not here to say that machines will be better than some of the many wonderful carers we have, but the company that can be afforded by a machine will surprise all of us. Anyone who has used an Amazon Echo will know, even with a little system like that, one establishes some kind of relationship.

There is a great headline by John Searle. John Searle is a philosopher at Berkeley. The day after Watson, IBM's computer system, won in a TV quiz show called "Jeopardy" in the United States—the computer system beat the two best ever human "Jeopardy" champions—his headline was, "Watson doesn't know that it won on 'Jeopardy'." I think that is perfect. Machines do not know that they have won a game, or would not know

they are being empathetic. In so far as you think empathy is the other person sitting in your shoes, I do not think machines will have that feeling.

The Chairman: I am afraid I must stop you there, Professor Susskind. We must move on to the next question. Thank you.

Q98 **The Lord Bishop of Oxford:** You have explored quite a lot the inequality in the distribution of change and I have to say, having listened to a lot of evidence around AI, this is the scenario which I find most chilling and disturbing. I admire your optimism but I am not yet at the point where I can be optimistic on the basis of the evidence you have given. I have spent most of my life in Yorkshire in the north of England, in communities recovering from the decline of the mines and the steel industry, and those are clearly the communities which are going to be hardest hit by the disruption which is coming. What do we need to do now to do the thinking and the preparation for disruption which, from what you and others say, is inevitable because we cannot avoid the competitive gains in productivity which will drive this cycle?

The Chairman: Would you like to start, Dame Henrietta?

Professor Dame Henrietta Moore: We need to start taking charge of the situation and begin to address it directly. Often when we talk about infrastructure investment we see this as a form of investment; when we have to do anything with social care it is a cost. The first thing to do is to think about investing in these systems of social care and how we manage this process of transformation. We have an awful lot of catching up to do because we already have extreme immiseration—and that is not too fine a point—in many parts of the United Kingdom. We have neglected to invest in these places and people over a very long period. We have a legacy that we will have to address as we address this big change. The first thing we need to do is start thinking about what would be a new social contract for the future and how it will work. We will create wealth in the future, we will create value, and there will be productivity gains from machines. What will we do with that wealth? If we look at platforms at the moment and the kinds of integrated technologies we have, we can see that they tend towards two things: one is a natural monopoly by certain individuals, and the other is creating lots of jobs around them which are very insecure, which are talked about in terminology always as being very disruptive. In fact, they are not disruptive at all; they are a very old-fashioned form of making money out of people, which is quite an easy thing to do.

In this situation of moving towards a natural monopoly, we have all these robots. Let us assume that they are distributed across the United Kingdom performing various kinds of wealth or value-creating tasks—whatever those may be. Who owns the robots? Is there some way in which, at a local level or indeed at a national level—some kind of sovereign wealth fund, or whatever you can imagine—those who are not directly designing the robot or benefiting directly necessarily from the robot or are not capable of engaging with “What kind of algorithm do we need?”, or whatever, will get a dividend out of that wealth creation? How are we thinking about that, for example? There are many ways in which

we can think about the ways we create wealth and how we distribute it. That is before we get on to what we do about education, and so on, which I imagine is another question, so I am going to stop there.

Olly Buston: One of the fundamental things that feels as if it is happening is the economic modelling based on capital and labour seems to be breaking down a bit. In the past, capital provided the infrastructure and labour did the work. Going forward, capital will do the work in the form of robots and machines. That poses great challenges in terms of inequality, our taxation system and welfare system. Bill Gates has called for a robot tax. I think he meant it slightly flippantly to try and spark a conversation; I do not think he meant we should go and count all the robots and tax them. We are mainly talking about software anyway, so it would be a bit silly. There is a need to think about taxation. There is a problem in that the second-biggest contributor to our tax take is income tax, and that disincentivises the hiring of labour. That is a problem we will face in the future. We need to think about that and about welfare. Obviously, there is a lot of excitement around a universal basic income. It is great that trials of that are going on around the world, including in Scotland, in Fife, and elsewhere, which is really positive. That definitely has potential as an idea. The part that may be missing is the sense of purpose that most people have. Currently, people get a lot of their sense of purpose from their work, so it is not just about income. That probably goes back to the idea of needing a new social contract. Part of that is how you give people purpose in this new society.

Professor Richard Susskind: I do not share your pessimism on one dimension. The good news is that the work we have done suggests we have at our fingertips an opportunity to make health services, legal services, educational services and all sorts of services immediately available. The liberation of knowledge, the removal of the gatekeepers, it seems to me, is one of the most wonderful and exciting prospects. The most profound inequality is the one that has been identified: that we will see the concentration of capital, which will be the main source of value, in the hands of a very few. We are seeing it already. How we redistribute the value from that capital seems to me to be one of the great economic and political questions of our time. We are nowhere near tackling that but, be in no doubt, that is going to be a fundamental issue for us. A universal basic income and a little bit of tax on a robot does not begin to address this fundamental issue that the overwhelming value in our world will be focused on the data, the software and the hardware that will be generating the value. All the indicators now are that that value is going to reside in a very small number of individuals and organisations. We need to think deeply about that. That raises fundamental, traditional questions of political theory.

Q99 **Viscount Ridley:** Will there be challenges for companies and organisations looking to integrate AI into the way they work in a significant way? What might these be? Do you think this could have an impact on how quickly jobs might change or be automated?

Professor Richard Susskind: The fundamental challenge here—people call it disruption—as I often say, is you cannot change the wheel on a

moving car. You have many very successful, conventional businesses and they will be challenged by this technology. Some technologies are, in the jargon, sustaining; they support and enhance the traditional way of working. Increasingly, AIs are disruptive, they replace the old ways of working. It is very hard for traditional businesses, while they keep going, to replace their old ways of working. The most successful way, it seems, of doing this is by starting new, parallel businesses. Rather than trying to graft AI on to the old way of working, the challenge for most businesses is starting up parallel AI businesses into which more and more of their businesses will flow.

Professor Dame Henrietta Moore: It is partly about starting with that question of what problems humans want technologies to solve. Very often, companies are not doing that but thinking about how they can jump on the bandwagon or how they can avoid being disrupted by the disruptors. All those kinds of things are very blinkered ways of thinking. We see that because we see the levels of investment in AI, of course, have gone up exponentially, but they are not what they ought to be if we are going to face this challenge going forward in British business. British companies overall are investing quite slowly. That might be because they do not really understand where these technologies are going, but I think it is also because they are, exactly as Richard says, thinking, "What can we do, given where we are, to incorporate these things?" We will very likely see that the pace of automation will not be as fast as it could be, in some areas, because it will be slowed down by this uncertainty of investment. Again, we have to understand that this question about AI and its impact on jobs has to be seen in the broader question of what is happening to the economy at large. There is a great deal of uncertainty in business now anyway. That is impeding people's ability to innovate.

Olly Buston: Access to good-quality data may be slowing take-up. Computer scientists say, "Rubbish in, rubbish out" and that data is the fuel of technology. That may change if we develop ways of structuring and labelling rubbish data and making it a bit less rubbish. Data is a big factor here, I think. Low labour costs, at the moment, may be a drag on take-up. I am an absolute optimist about all this as well. I think there is a huge opportunity for the UK in the delivery of public services in health and education and in massive economic opportunity. The Government should invest heavily, by its industrial strategy, in AI and pouring fuel on the fire of our amazing computer scientists and tech companies. There is a Brexit point here as well, I guess, in that those brilliant universities and tech companies need access to the highest quality of talent. The UK's ability to capitalise on the incredible assets it has in this area will depend a little on the continuing ability to attract the very best talent to our universities and companies.

Viscount Ridley: I am glad you raised Brexit because I wanted to ask a supplementary on that. Professor Susskind said something rather different, which is that Brexit discontinuity potentially gives us the chance of taking advantage of any new flexibility outside the EU; for example, an immigration policy that makes it easier for the most talented technologists from beyond the EU to work here and by cutting away the regulatory barbed wire that currently hinders so many companies. Is

Brexit going to help or hinder speeding up the slowing down that Professor Moore talked about?

Professor Richard Susskind: That was after my earlier and, frankly, overriding observation, which was—and this is what I genuinely feel—that at a time of greater technological progress than humanity has ever witnessed, when I would want our entire government focus to be on leading the way in exploiting these emerging technologies, most of our Ministers are preoccupied with disengaging from Europe. Without commenting on the merits of disengaging with Europe, the observation I made there was to try to offer some kind of rationalisation. My very strong feeling is that I wish I lived in a country where it was the case that our Ministers were excited by these technological developments and were thinking profoundly about how they affected our high-level strategy, international competitive positioning and the standard of living of our human beings, and so forth. It is the worst possible time to be considering a demerger.

Lord Giddens: My questions have already been answered about five times, I would say, at least.

The Chairman: Just be creative, as ever.

Q100 **Lord Giddens:** Given the possible impact of AI on the labour market, should the Government be considering how to mitigate this, or is it still too early? If it should start planning now, what options should it be considering? One way of focusing it might be, perhaps, to say a few words about the future of universities and labour markets in this context. Students are investing £50,000 for jobs that might not exist down the line. A huge crunch could come with student debt. You might want to focus in some way on that or give it a different slant, or you could say what you think of *The Patient Will See You Now* in light of the intervening years since Eric Topol wrote that book.

Professor Richard Susskind: It is a great title for a book. Can I comment?

Professor Dame Henrietta Moore: You go ahead. We are doing well, going left to right, or your right to left.

Professor Richard Susskind: My worry is that universities are generating 20th-century rather than 21st-century graduates.

Lord Giddens: You should feel free to say more on the general question, if you wish to.

Professor Richard Susskind: Can I focus on universities? Would that be okay?

Lord Giddens: Indeed.

Professor Richard Susskind: My worry is we are generating 20th-century rather than 21st-century graduates. I go back to my fundamental distinction: in the future, we either compete with the systems or build the systems. We need to think, “Are we educating our graduates to be people who are going to try and compete with these systems or build the systems?” My vote goes for educating a new generation of people who

will be developing the systems that will replace our old ways of working. If you look at law schools and medical schools, and so forth, what we teach and, to a large extent, how we teach has not really changed since I was there in the 1970s or 1980s. It worries me deeply. In fact, in many ways, a lot of what we do in universities—and I say this as a university professor—is teaching our students to do things that already machines seem to be outperforming us at. If you are right, for example, about empathy—I was not sure if you were taking a position there—and believe that empathy and interpersonal skills are distinctive human capabilities, they should be very much an emphasis in our educational establishments, and they are not just now. We still focus on the acquisition of content, if I can put it in technological terms. We need to think deeply about what we anticipate our workforce doing in the 21st century. As I say, the slice of the pie devoted to activities where humans can outperform machines seems to me to be diminishing in size.

Professor Dame Henrietta Moore: However, knowledge will be the key commodity, so we have to do something about the universities. I think the universities are responding. Many more universities are thinking about how you construct what we might call earn-as-you-learn or learn-as-you-earn degrees. In other words, we have to start thinking about how the boundary between worker and student is going to be much more porous, as I said earlier. You are going to go back several times for training. That means that the boundaries between companies and universities will shift again. We have to think very broadly about these issues. As I have already said many times today, it is not helpful to focus just on jobs and the quantum of them when we talk about AI and automation. We have to think about the broader range of things that this will engage with. One of them is how will we constantly reskill our workforce? What will happen with that workforce? For that, we can begin to see a glimmer. The question is: will that workforce be a very small elite? Will it be the elite out of a group of people who have been able to engage with a huge amount of debt and then pay it off? What happens to everybody else? This takes us back to some of the earlier answers we gave this afternoon. These things are hugely important, but I think universities are already thinking about how that will work.

Olly Buston: It is definitely not too early. Looking at the broader question, one area where more focus is needed is greater research on the uneven distribution of the impact and responding to that. If we can get a better picture of regional impact then we should be developing regional strategies in terms of investment, financial support and psychological support for people who may lose their jobs. We need more granularity in particular sectors—for example, the call centre sector. It depends on how you cut it, but there are about 1 million people employed in call centres. There is a great argument about when this moment will happen, but once the software reaches a certain threshold of quality it is possible that quite a lot of those jobs will go because it is a highly competitive market. Those jobs are concentrated in former industrial heartlands. We need to think through as to those sectors which will be most hit and the response. Driving is obviously another one, although there is huge debate about when it will happen.

The Chairman: Government should be doing that now, do you think?

Olly Buston: Definitely the research and developing a more detailed understanding, I think, and piloting some of those themes around UBI and thinking about different tax models.

The Chairman: Thank you. I am going to move on to Lord Swinfen. I am very conscious that this question has been partly covered, so I think we should start with you, Olly.

Q101 **Lord Swinfen:** Will the current education system adequately prepare young people for the challenges they will face in a world dominated by artificial intelligence?

The Chairman: Therefore, you may want to concentrate on non-university education, so to speak, at this stage.

Olly Buston: The answer is definitely no, for several reasons. The immediate, knee-jerk response, which at one level is correct, is that we need way more STEM skills—science, technology, engineering and maths—and that, absolutely, girls need to be encouraged to focus and develop those skills because the technology sector is very male-skewed. In the long run, as we have heard, it looks as if jobs that involve creativity, empathy, interpersonal skills and high levels of dexterity will be the ones that remain human for longest. You need to develop education systems that deliver those. As Dame Henrietta said, there is the nature of education, so how do we have lifelong learning? How do we have learning that fosters resilience, so people can adapt to incredibly volatile environments? There needs to be a White Paper on adapting the education system for the future. The children going into preschool and primary school now will be impacted by this. When we talk about creativity, it is very hard to teach a 44 year-old to be creative and it is pretty hard to teach a secondary school pupil to think creatively, so resilience and creativity are skills you develop very early in life. We have to think quite early in the education system.

The Chairman: Would you like to add, Dame Henrietta, Professor Susskind, anything beyond the higher education points you made earlier?

Professor Dame Henrietta Moore: Just to say that already many people in the British economy are not doing one job; they are doing multiple jobs. This multiplicity of tasks, where you come together to do something and then take up other tasks to do something else, is likely to be a key facet of the jobs of the future. You are not going to be doing a single job. You will not be just a programmer, or whatever it is. Thinking about how we train young people, both at school and university, to do those multiple jobs is something we have not done yet. At the moment, we see multiple jobs as a very low-end problem in the job market but it is going to be a problem for everybody in the labour market as we go forward.

Professor Richard Susskind: I agree with what has been said but have one slightly different angle. It is not so much focusing on what we teach but how we teach. We should be using AI in teaching. The advances being made in online, personalised and simulated learning environments

are remarkable. It is not just about teaching about AI, it is exploiting these technologies to improve the learning environment radically.

Q102 **Baroness Grender:** What do the public need to know right now in order to make informed choices about their skills, jobs and careers, given the possible impact of AI? Evidence suggests that people are not worried about this. Who should be responsible for making them worried, or knowledgeable?

Professor Dame Henrietta Moore: I think when people are asked if they think robots will have a big impact, they tend to say no, but a lot of people deny that climate change exists. We have a number of difficulties in this area about the future. The process we should be getting involved in, as Richard said, is shaping the future. We should be getting people who are working already involved in co-designing the technologies of the future to improve their work systems, improve their tasks and, also, so that they can imaginatively engage with a future that is going to be further technologized. One of the things we have not done is go to anybody who we believe is going to be affected by any of these jobs and say, "Okay, would you like to join a conversation about how you can shape your future work?" This is a form of democratic deficit and explains, in part, why people feel, as they say, "left behind" by processes such as globalisation, technology, and so on. It is partly because they have been left behind that they are not being engaged with. Yes, we need to do lots on the engagement of the public and workers.

Baroness Grender: Who should be responsible for that?

Professor Dame Henrietta Moore: It should be across the spectrum. It should start in schools, as Richard said. That is incredibly important. It should happen in the workplace and in the universities. We are beginning to do that already, and there are lots of things we can do there.

Professor Richard Susskind: I would like to see government-led public debate about this issue as well. I go back to the early 1980s when Mary Warnock was invited, in relation to in vitro fertilisation, to raise public awareness and discuss the numerous issues it raised, not just scientifically but, clearly, ethical as well. It seems to me the message is we are at this time of increasingly capable machines and when we are seeing greater and more rapid technological progress than the world has ever known. People should know this and expect a period of disruption and discontinuity. We need an informed, sober and stable, evidence-based public conversation, I would hope led, clearly, by employers, schools and so forth. It seems to me this is a perfect role for government to put some shape around this discussion and to identify the implications of these changes for all corners of society.

The Chairman: Olly, you have started a debate with your report, have you not?

Olly Buston: Hopefully. We do an annual opinion poll of public attitudes towards AI with lots of questions. One was that only 28 per cent of people are worried about job losses in their area due to automation. It does seem as if the public are either less worried or less aware than they might be, given what might be coming down the track. It seems there is

a job for government, employers and others. An area where a much bigger, broader public debate is required is around data. Artificial intelligence is supremely good at interrogating data, and the way we handle data as a society at the moment—

The Chairman: I am going to stop you there because Lord St John is going to ask the next question, which deals precisely with data.

Q103 **Lord St John of Bletso:** Indeed, it is on data. If labour becomes relatively less important in an increasingly automated world, should we be considering ways to remunerate individuals for the data they provide? How do we keep our personal data safe? By what criteria can value be ascribed to it?

The Chairman: You can carry on your train of thought, at this point.

Olly Buston: We need a new deal on data between governments, businesses and citizens. We have that slightly farcical social dance where you do not read the form, you tick the box, the company knows you do not read the form, you know the company knows you do not read the form but, somehow, quite a lot is built up on this. Consent is a fragile thing. For AI-based businesses to thrive this needs to be fixed. Currently, the default is, "You can have my data for eternity and here's my friend's data". We need to move to a situation where people are giving data through consent for specific periods of time to perform a specific function. That could all get horrifically time-consuming for everyone. We may need technological solutions to this. There are some interesting developments around data assistants that would help you manage your data.

The Chairman: Blockchain, or whatever may be helpful?

Olly Buston: Definitely, yes.

Professor Dame Henrietta Moore: There will be a further explosion of data, so we need to attend to this urgently. Teaching about it in schools is a starting point, as I said earlier. It is enormously important. For example, the new changes in medical technologies mean that it will be possible quite soon to do a whole MRI scan using something attached to your mobile telephone. That is very, very close already. That data will be uploaded to a cloud system and, therefore, it will be potentially held or accessed by many individuals. It is not just what I do on Facebook or on Twitter, which is the very thin end of the wedge, in regard to the data issue. Data will be one of the new ways of creating value, and we have to capture that value in the economy for the benefit of all. If we do not find ways of capturing it, some people we do not want to will capture it.

Professor Richard Susskind: I worry about the costs of trying to link the value of one's data and the mechanics of recovering some kind of payment for use of your individual data. It is part of the far bigger problem we discussed earlier, which is who owns and controls not just the data but the software, the technology, and so forth. I say again the biggest political and, perhaps, economic challenge of our day is to think about how we redistribute the value generally, not the value of the individual pieces of data but the productivity that these machines are creating for society.

Lord St John of Bletso: If I can ask a supplementary: some would argue that blockchain will become the repository of personal data. Do you see blockchain as being the ultimate repository of personal data?

Professor Dame Henrietta Moore: I do not think I can answer that question in any sensible way.

Olly Buston: This is not a very good answer, but where investment is towards companies with giant swathes of data, with greater data control you might get a shift to the value of companies being about the creativity and development of algorithms rather than how much data they have. That may be a positive development. Where blockchain may provide interesting solutions for how individuals can hand over parts of their data for particular purposes, it may be part of the solution.

The Chairman: Thank you. I am going to move on to Lord Holmes for the final question.

Q104 **Lord Holmes of Richmond:** Thank you for a fascinating session, in light of which this question seems a touch reductive but it is what it is. If there is one recommendation you would like to see the Committee make at the end of this inquiry, what would it be?

Professor Richard Susskind: I have been reflecting on this. One of the difficulties is that AI is so pervasive. It affects our work, it affects the way we govern, it affects the way we live and work, and so forth. We have not raised these at all, but profound legal issues arise in privacy, competition law, intellectual property law, crime and so forth. It occurs to me that, certainly within government, there is no single focal point; we are seeing an industrial strategy here and we are seeing a development in the justice system over there. My thought was I wonder if it would help, perhaps, in the Cabinet Office having a Minister responsible for AI who is looking sideways across AI, encouraging individual departments to exploit AI, thinking of the profound legal implications of AI and speaking to and encouraging uptake in industry. I know we have a Minister for digital strategy but I am thinking of something that cuts across government. My suggestion is that we have a Minister for artificial intelligence. This would also, as another point, raise awareness of AI, recognising it as a fundamental issue of our time.

Professor Dame Henrietta Moore: I would like to see a very strong recommendation about attention to the distribution of the benefits and of the wealth that will be created out of it. That could focus very profitably on a new kind of social contract, as we said several times today.

Olly Buston: I support both my colleagues' suggestions. The Minister is an interesting idea. Artificial intelligence was mentioned 19 times in the House of Commons last year, and that was more than it had been mentioned ever before. There is a need to elevate the political debate. I get the sense that within the Civil Service, Government and the Opposition there is not enough grappling with this. To add an additional third, I would say transformation of the education and training system; a new White Paper for education focused on STEM in the short term and, absolutely, encouraging girls into STEM; thinking how you foster creativity and interpersonal skills to prepare people for the jobs of the

Future Advocacy, Professor Dame Henrietta Moore and Professor Richard Susskind OBE – Oral evidence (QQ 95–104)

future, and changing how people learn to help them become more resilient and more able to direct their own learning going forward so they are prepared for a very volatile future.

The Chairman: I think you have about three things in there. That was very cunning. I am afraid it has been a bit of a gallop and my colleagues have not been able to ask all the supplementary questions they would have liked to have asked. We could, obviously, have extended this session considerably. Thank you very much indeed. It has been a very informative session.

German Research Centre for Artificial Intelligence (DFKI) – Oral evidence (QQ 163 – 171)

Evidence Session No. 17

Heard in Public

Questions 163–171

Tuesday 5 December 2017

[Watch the meeting](#)

Members present: Lord Clement-Jones (The Chairman); Baroness Bakewell; Lord Giddens; Baroness Greender; Lord Hollick; The Lord Bishop of Oxford; Lord Puttnam; Viscount Ridley; Baroness Rock; Lord Swinfen.

Examination of witness

Professor Wolfgang Wahlster [video link].

Q163 **The Chairman:** Professor Wahlster, it is very good to see you. I am delighted that the system and the technology works. A very warm welcome to our Select Committee. This is the 17th formal evidence session for our inquiry. This session is to help the Committee to discuss the international aspects of artificial intelligence and the policies that comparable countries to the UK are pursuing with regard to artificial intelligence.

This is a public evidence session that is being broadcast. People can log in and see the stream of the Committee's evidence session. A transcript will be taken of your evidence, which you will be invited to check for accuracy. If, after the session, you wish to clarify any points made during your evidence, we would very much welcome supplementary written evidence.

We have your curriculum vitae and your biography, so we know that you are the CEO and scientific director of the German Research Centre for AI, the DFKI. Thank you very much for your presentation, which we have digested and will look at further in the following week.

We would like to start, if we may, with our questions to you, particularly as time is short. Would you permit me to start the questioning?

Professor Wolfgang Wahlster: Yes, please. Thank you very much for the invitation. I am looking forward to the questions. I have received eight questions.

Q164 **The Chairman:** That is fantastic, so you know what is going to come. Let me start with my question. What are the biggest advantages and disadvantages that AI could bring to Germany over the next 10 years?

What is distinctive about Germany's approach to AI? We are trying to tease out whether you think that Germany has particular advantages or disadvantages when it comes to capitalising on or developing AI.

Professor Wolfgang Wahlster: We have a very special approach that is based on Germany's industrial strengths. In our AI programmes, both at the DFKI but at the national level in general, we work with our strongest export industries to inject artificial intelligence: first, into our premium cars made by the big car manufacturers; secondly, into our manufacturing machine industry, which is very strong in exports; thirdly, in the medical equipment domain, where Siemens is the world leader in high-end scanning machines; fourthly, into agricultural machines; and, finally, into appliances—Germany has very good companies in Miele and Bosch, with high-end equipment for households, including the kitchen.

We have discussed with our Government a strategy to transform these successful export products in these five major fields into smart products. That means that on top of these platforms we have AI applications in service, maintenance, and so on, and in the future revenues will be produced from the services provided by these exported machines. This is quite different from the US approach, which is based more on internet services. We are more in the physical domain; you know that Germany is well known for its engineering and manufacturing. Every second job in Germany is still dependent on manufacturing.

The Chairman: Is that a government-led policy of an industrial strategy that involves those very important German companies?

Professor Wolfgang Wahlster: Yes. In DFKI, we have strong industrial shareholders from all these companies. As you know, DFKI is a public-private partnership that is half government and half industry. In the research union, which is where we have meetings with our government officials, Ministers, even the Chancellor, and people from industry and academia, we decided that the strategy should go from Industrie 4.0, which is our flagship programme, for the next wave of manufacturing over so-called smart services, which are based heavily on AI. In September, we launched a very new thing: the academic industrial platform for AI learning systems. The topic of autonomy is very important here—Germany has most of the worldwide patents for autonomous cars, for example; more than the United States and China—and in other areas such as aircraft, ships and even trains AI is very important to make them autonomous.

The Chairman: Do you believe that AI is likely to have a positive effect on employment in Germany?

Professor Wolfgang Wahlster: Definitely. We see it already, but with AI and Industrie 4.0 we can have onshoring, which means that we bring back companies that moved their factories to low-wage countries. From Asia especially they are now coming back to Germany. One example is Adidas, which has opened new manufacturing facilities in Germany for its sports shoes. We are in the age of mass customisation, so the customer can decide what kind of sports shoe to wear, and the tailored sports shoe has to be delivered very quickly. This cannot be done remotely in Asia; it must be done very close by.

Q165 **Lord Swinfen:** Good afternoon, Professor. What issues does Germany have in attracting and retaining skilled artificial intelligence researchers and developers? How is Germany tackling this problem?

Professor Wolfgang Wahlster: On the one hand, we try to have attractive programmes for bringing back academic staff who have spent some time especially in the United States. This is a programme financed by the Alexander von Humboldt Foundation stipend where an associate professor or full professor can receive up to €5 million by coming back, so there is an incentive when you return to Germany from other countries. That is one thing. Over past years, the Government have invested heavily in the research infrastructure and clusters of excellence. We have an excellence strategy, which has attracted a lot of people. Believe it or not, we have been able to hire people from Google, where the salary is higher, but the cost of living in Berlin, for instance, is so much lower that the standard of living is higher. That is our argument.

Q166 **Viscount Ridley:** Professor, good afternoon. Could I ask you about the role of the German Government? You have touched on this, but specifically—we know about the Cyber Valley initiative—what problems is the German digital strategy 2025 aimed at solving? Is the German Government’s funding sufficient to enable Germany to capitalise on AI? Do the German Government believe that there needs to be AI-specific regulation?

Professor Wolfgang Wahlster: Although I am a civil servant I cannot speak for the German Government. From my consulting role on many of these committees I can tell you that there is a strategy of having big programmes starting with Industrie 4.0. More than €1 billion was attracted in funding from both the government and industry. The second smart service was €300 million from different ministries and there is the learning system from the Ministry of Research and Education. There is a lot of funding, which is one thing. There are clusters. DFKI is now the largest AI laboratory worldwide with industrial and governmental sponsorship. New initiatives are emerging. Cyber Valley is a relatively small one run by the Max Planck Institute, but there are many other hubs in Munich, Berlin and so on. There is a very good spin-off scheme at DFKI. We have 80 spin-offs in AI. Every week in Berlin new spin-offs in AI are being created.

Viscount Ridley: Could you comment specifically on whether the German Government are looking at AI-specific regulation?

Professor Wolfgang Wahlster: As far as I am aware, they are but only in very specific areas. It does not make sense to have general regulation. On the question of autonomous driving, we had a special committee on the ethics and legal problems on fully automatic driving. That is the only government initiative that has come to a clear conclusion. There are diffuse discussions, but not what I would call initiatives.

Q167 **Lord Giddens:** Good afternoon. Thank you very much for talking to us. How do you see DFKI’s role in developing and utilising AI in Germany? How does DFKI’s work relate to that being done by industry in the start-up sector? You might like to say a little more about the extraordinary

success of German manufacturing, with 23 per cent of the labour force in manufacturing compared with about 9 per cent in this country. They have already pioneered quite a lot of automated production. Do you think that AI will introduce a new level in manufacture?

Professor Wolfgang Wahlster: Definitely. That is very clear to me. I created the term “Industrie 4.0” in 2010, together with Professor Kagermann, who was the CEO of SAP and is now the president of Acatech—the National Academy of Science and Engineering—and Professor Lukas from our ministry of education and research. Our goal was to boost manufacturing. As you said, it is 23 per cent, but indirectly it is much more, because every manufacturing workplace triggers other jobs in the ecosystem. Every second job in Germany depends indirectly on manufacturing.

With AI we are going to what we call the fourth industrial revolution, which is a revolution for factory workers with improved working conditions but also a completely new paradigm of manufacturing, because we can manufacture tailored products that are exactly what the customer wants. The kitchen and the car are unique products, so we have batch size 1. We have shown this in our manufacturing. Last year, for instance, no BMW produced worldwide was identical to another car. There are so many variations with this product that it is possible to handle this complexity only with AI. The specific applications are automatic online quality control, predictive maintenance and worker assistance systems. When workers have to work on many different products every day they need assistance, so there is the high quality that one expects from German products.

Lord Giddens: Would you like to say a bit about the role of 3D printing?

Professor Wolfgang Wahlster: 3D printing is important, but I think it is hype. There are many niche applications, and we use it already in aerospace and the car industry, but it is not the solution for everything. Semiconductors and cyberphysical systems are very important for AI, but currently we cannot do 3D printing for electronics on the nanometer scale.

Q168 **Baroness Greener:** How do you see Germany’s AI sector relating to the UK currently? How do you think this will change in the future? Do you think that relationship will be one of collaboration or competition?

Professor Wolfgang Wahlster: Historically, the UK was the leading country in Europe in AI. It started in Edinburgh; as you may know, Professor Michie was one of the pioneers. So there is the University of Edinburgh, but other universities such as Leeds, Oxford and Manchester are very strong and we have good connections with them. German professors teach there. I was the president for many years of the European AI organisation, in which the UK played a major role. We have many friends there and great respect. The academic quality has been very good over the last 10 years, but unfortunately what is missing is the industrial application and collaboration with industry. Funding has been reduced too much in the UK. There is still a very good basis and an excellent tradition, but the potential of the academic talent is not used completely. That is my opinion.

Baroness Grender: What is your prediction for the future of that relationship?

Professor Wolfgang Wahlster: I hope that we can continue good collaboration. For instance, we always have two or three professors from the UK on our scientific advisory board at DFKI. They come a couple of times a year, look at our research and get inspiration. This will continue, but my only concern is Brexit and how this will work in relation to bigger European projects. We should try to connect on the industrial side, and Industrie 4.0 may be a good start, because I am convinced that every nation has to have a strong basis in physical products. You cannot make a living only with financial serviceengineering.

Q169 **The Lord Bishop of Oxford:** Professor, can you comment on data protection law? What implications does German data protection law have for the development and application of AI in Germany? We know that German data protection law is rigorous and strict. Could you comment on the impact of the GDPR? In the UK, we are exploring the potential of data trusts, a safe and responsible way of sharing datasets between public and private organisations. Is that being explored in Germany?

Professor Wolfgang Wahlster: Yes. This is a very important question. It is essential for AI that we have access to public data. I will give an example. In the traffic domain we require access to traffic light data. We have very good systems in our car industry that make it possible to save energy, fuel or battery by checking in advance by wireless services and AI when the signal will go red. Unfortunately, access to this data is not fully available at the moment. On Friday, we will have the first demonstration in Darmstadt, our Digital City, of a traffic light change being transmitted to a car, so you will know that it is not a good idea to expend gas on accelerating because when you get to the traffic light it will be red. That example illustrates that public data should be available for societal applications.

On the other hand, there is public concern about the privacy of personal data, which is very important given our history of the Nazis and, in eastern Germany, the Stasi. People are very sensitive about this data, which is good. Google does most of the research into privacy issues in Germany, and it is a tough test because most of the German population is very critical. For AI, I think this is okay, because, as I said earlier, we are more interested in data from machines, which is the domain in which we apply AI. I am not completely convinced about using AI in social media—Facebook, WhatsApp—which is the American domain. We should protect ourselves to ensure that this is not overdone. This is not so much the business of German companies; we are more in the technical domain. I have not heard that there should be any privacy protection of machine-learning data generated by manufacturing equipment.

The Lord Bishop of Oxford: What about the data trusts? Are you exploring anything similar to that idea?

Professor Wolfgang Wahlster: Yes. This is important. On the industrial side we have a big project called the Industrial Data Space initiative, through which we have formed a trust network where companies work together. Airbus is a very good example of European co-operation. These

are different sides and different ecosystems. Companies that are vendors can exchange their data in a trusted way without being compromised. In the area of AI, compromised data is a big problem. It is not the standard attacks and viruses, but if you get garbage data into your system the machine learning may lead to very serious problems because the data is contaminated.

Q170 **Lord Hollick:** You have painted a picture of a very close focus on the great strengths of German manufacturing and engineering. In that context, is there a need for international collaboration? If so, what role do you believe Germany might play? I have two points. What specific global AI initiatives is Germany participating in at the moment? What is the German Government's position on an international ban on autonomous weaponry?

Professor Wolfgang Wahlster: On the first question, I cannot speak officially for the German Government, but because I am involved in these activities I can talk about one big activity, together with the European Commission. We want to set up AI on demand whereby all the software tools that are developed for AI become more accessible as components that can be used as bricks, as when you build a house. That is a big initiative, which is being prepared and will be launched next year by the European Commission. The initial budget is €20 million.

We are discussing collaborations in Industrie 4.0 and Smart Service World with China and Japan. We have very good collaboration with the Czech Republic: the collaborative robots mission. We work on team robotics, where a couple of people work together with a couple of robots in the manufacturing context. A lot of this goes to rolling out the concept of Industrie 4.0. Our main mission from the government and research point of view is that Industrie 4.0 is a success. Last year, there were 60,000 scientific publications. The term is known worldwide. We have an advantage of two or three years, which we use, but we want to teach other countries to use this technology. Ultimately, it is very important that we can sell the technology but also that we can collaborate with other countries for the fourth industrial revolution, in which AI is one of the key ingredients.

Lord Hollick: Is the collaboration that you have spelled out with other countries at government level or a partnership between government and industry in Germany partnering with their opposite numbers in these countries?

Professor Wolfgang Wahlster: That is a good question. It depends. In the Czech Republic, for instance, it is based on both. It is a contract between the government and the two key research institutes, but it is also, as you pointed out, a collaboration with matching industry. In Germany, for instance, we have Volkswagen, and in the Czech Republic there is Skoda. In the specific project on team robotics, both companies are involved. We do the same in China and Japan, where we have very large subsidiaries of Siemens, Bosch and other German companies, but we work also with the research centres of Hitachi and Ricoh in Japan.

Q171 **Lord Puttnam:** If there is one recommendation that our Committee

should make at the end of this inquiry, what should it be? As a civil servant, can you say how closely the Federal Government are monitoring both the protection and abuse of intellectual property? Is this something the Government are engaged with?

Professor Wolfgang Wahlster: Yes, is my answer to your last question. We have some foreign investors in DFKI, for instance, and in other institutions. Our Government monitor this closely. Only companies investing in Germany in research, development and production are allowed to join. A foreign company, be it from the US or China, cannot work with us on projects if they only have sales organisations in Germany. This is the very strong criterion that has to be fulfilled.

On IP, at DFKI—I can only speak for DFKI—we have the principle that if a company pays fully for the research, they get the IP; we do not withhold any IP in this case. Most of our projects are in consortia or partially funded by the Government, and according to German regulations other companies with headquarters in Germany can use this IP.

Lord Puttnam: My core question is whether there is one recommendation that we should make at the end of this inquiry.

Professor Wolfgang Wahlster: My recommendation is that you go back to the roots. AI was invented in Europe, more or less. There were some people in the US, but the UK played a very dominant role in launching AI. My main recommendation is that you play to your old strengths and copy the German model a little to have stronger connections between industry and academia. That link is missing. There were some good efforts like the Catapult centres in the UK. I have discussed this a couple of times with the British Embassy in Berlin, and the conclusion was that Germany has an advantage in that academia-industry collaboration is in the genes of the industry and professors whereas in the UK this is still lacking.

The Chairman: Professor, we have had a splendid evidence session. Thank you very much indeed. You have given us a great deal to think about. We are extremely grateful to you, as I know you are between various meetings. We are all very grateful and have very much appreciated your evidence session. Thank you very much indeed.

Professor Dame Wendy Hall, Professor Michael Wooldridge and Professor Nick Bostrom – Oral evidence (QQ1–8)

Professor Dame Wendy Hall, Professor Michael Wooldridge and Professor Nick Bostrom – Oral evidence (QQ1–8)

[Transcript to be found under Professor Nick Bostrom](#)

Professor Chris Hankin and the Alan Turing Institute – Oral evidence (QQ 143–152)

Professor Chris Hankin and the Alan Turing Institute – Oral evidence (QQ 143–152)

[Transcript to be found under Alan Turing Institute](#)

Dr Hugh Harvey, National Data Guardian for Health and Care Dame Fiona Caldicott and NHS Digital – Oral evidence (QQ 128 – 142)

Evidence Session No. 14

Heard in Public

Questions 128–142

Tuesday 21 November 2017

[Watch the meeting](#)

Members present: Lord Clement-Jones (The Chairman); Lord Giddens, Baroness Greener; Lord Hollick; Lord Holmes of Richmond; Lord Levene of Portsoken; The Lord Bishop of Oxford; Lord Puttnam; Viscount Ridley; Lord Swinfen.

Examination of witnesses

Dame Fiona Caldicott, Dr Hugh Harvey and Professor Martin Severs.

Q128 **The Chairman:** A very warm welcome to our witnesses today. This is the 14th formal evidence session for the inquiry and it is intended to help the Committee discuss the application of artificial intelligence to healthcare in the UK. We are very pleased to have Dame Fiona Caldicott, the National Data Guardian for Health and Care, Office of the National Data Guardian; Dr Hugh Harvey, clinical intelligence researcher and consultant radiologist at Guy's and St Thomas' NHS Foundation Trust; and Professor Martin Severs, medical director, NHS Digital and Caldicott Guardian. The word "Caldicott" seems to appear rather today. We will no doubt discover more.

I have a little rubric I need to recite. The session is open to the public, a webcast of the session goes out live, as is, and is subsequently accessible via the parliamentary website. A verbatim transcript will be taken of your evidence. This will be put on the parliamentary website. A few days after this evidence session, you will be sent a copy of the transcript to check it for accuracy and we would be grateful if you could advise us of any corrections as quickly as possible. If, after this session, you wish to clarify or amplify any points made during your evidence, or have any additional points to make, you are very welcome to submit supplementary written evidence to us. Perhaps I could start by asking each of you to introduce yourselves.

Dame Fiona Caldicott: Thank you. I am Fiona Caldicott, National Data Guardian since 2014, but my work in this area goes back to 1996, hence the appearance of the name Caldicott in your bio-note. I was asked at that time to chair a National Health Service committee on whether patients' confidential information was being used when it did not need to be. The medical profession was very worried about identifiable

information being exchanged; for instance, when pharmacists ask for prescriptions to be funded where they had dispensed a drug with a lot of detail about the patient on the prescription. We found 84 flows of data where we did not think that that sort of use of identifiable data was necessary. That led to the establishment of Caldicott Guardians in each large health service organisation in the country. They are the people who safeguard the patients' information in the trust, in the hospital, so that it should be used only for the purposes for which it is intended. If difficult situations arise, they are the individual to be consulted by staff about the right thing to do. Hence the name tends to appear rather a lot.

Professor Martin Severs: I am a consultant geriatrician by trade, as it were, and I am a professor of healthcare for older people. I have been involved in information standards for over 30 years and set up the advisory structures to the Royal Colleges originally. I was responsible for then taking terminology out of private ownership in America into public ownership across the world and into a product called SNOMED CT. I then had the honour of serving under Dame Fiona on her second review. I have been the medical director of NHS Digital for the last three and a half years, and full time for the last year and a half.

Dr Hugh Harvey: I am a consultant radiologist across the river at Guy's and St Thomas' Trust. I have also had the opportunity to do research at the Institute of Cancer Research looking at some machine learning applications within the field of imaging. I have also worked within tech companies in London using clinical artificial intelligence and predictive analytics, and I now advise a couple of start-ups in the technical field in healthcare within London. I also sit on the Royal College of Radiologists' informatics committee and am part of an artificial intelligence working group with the Royal College.

Q129 **The Chairman:** I was going to ask what you did in your spare time, but I will not do that.

I will ask a very general question to start with. To what extent is AI already used in healthcare? Where in health do you see the biggest potential for the use of AI?

Dame Fiona Caldicott: It is being researched in universities and applied to health but I do not think it is widely used within the provision of services as yet. My colleagues can expand on this but I can give you one example, from our trust in Oxford. We have worked with the department of engineering on a software product which builds on algorithms that relate to the patient's vital signs; their blood pressure, temperature—the various things you need to know about a patient to care for them adequately. The algorithm alerts nurses when the patient is likely to deteriorate. If you are looking at a ward of 30 patients, you will be told through the electronic record of the patient that Mr B is likely to deteriorate because of what is happening to his vital signs. It means that the nurses do not have to go round routinely measuring those, which saves clinical staff a great deal of time. That is an application I am familiar with in our trust, but colleagues will know others.

The Chairman: I will come on to other members of the panel. You would

agree with some of our previous witnesses when they said that the adoption of AI was complementary rather than a replacement of professionals.

Dame Fiona Caldicott: Absolutely.

Professor Martin Severs: I agree entirely with Dame Fiona. I have a slightly wider definition of AI that includes statistical algorithms and high-powered computing technology applied to large datasets. There are a number of statistical algorithms already in place, particularly in general practice; not the machine learning types we have already articulated, but certainly algorithms are used quite commonly in general practice to help them manage patients. High-powered computing technology is particularly being used in commissioning in identification of high-risk patients—that is for either population management or individual management—and to help the commissioners look for patterns and mine the data such that they can work out how better to commission services.

Dr Hugh Harvey: Taking artificial intelligence as a broad umbrella for the entire field of computing which can mimic human intelligent thought, I use it every day. I voice-dictate my radiology reports—it recognises what I say and transcribes it. That has almost completely replaced radiology secretaries within hospitals over the past 10 to 15 years. That is quite a rudimentary form of artificial intelligence. If we are talking today about deep learning and finding complex patterns which humans are unaware of within large datasets, I can think of no examples in current clinical use.

The Chairman: Do you see any dramatic changes or is this all a smooth progress that takes place over the next few years? Will there be no great dramatic leaps forward, in that case?

Dr Hugh Harvey: It will be slow to start, but it will build and potentially become exponential as we figure out how to implement these in practice.

The Chairman: What about timescales?

Dr Hugh Harvey: That is the key question. I would be foolish to put a number on it but I would say it is going to be at least a decade before we see a deep-learning algorithm in proper clinical practice used routinely.

The Chairman: Is that what some people would describe as general AI or are you talking about some intermediate situation?

Dr Hugh Harvey: An intermediate phase, yes.

The Chairman: You are not making a forecast for general AI.

Dr Hugh Harvey: No.

The Chairman: We will not hold you to that, in that case.

Dame Fiona Caldicott: The pattern we will see is that trials will take the place of application. For instance, the one I have described to you is now being used in the community to measure the vital signs of patients with chronic pulmonary disease. That means there is a way of measuring that which can then be looked at from the hospital, saving a number of visits. The patient can manage their own health and the carers coming in will be able to look at the record, and so on. That will be researched in order to

show that the results are valid. There will be a gradual extension of the use, starting with some pilot sites. That would be the way the application would begin to take off.

The Chairman: I am going to move on to the next question but I am going to park the point about the recognition of patterns in the NHS, for public health and so on. We will want to explore that later.

Q130 **Viscount Ridley:** Good data is the fuel of AI research and application, and probably the thorniest issue we are wrestling with on this Committee. Should all publicly generated health data be made publicly available—subject, as far as possible, to anonymization, if that means anything—in order to encourage progress in AI research and innovation? Perhaps, if you get a chance, you could also comment on the proposals for data trusts and on the idea that the National Institute for Health Research should consider setting up an AI bioresource, similar to the approach taken in genomics.

Professor Martin Severs: I am happy to start with that. Speaking as the Caldicott Guardian for NHS Digital, I will first take the big question and parse it down, if I may. I would be very happy to make aggregate data with small number suppression publicly available, but what really counts is the degree of granularity in the data for AI. I would not want to make patient-level data publicly available but I would be very happy to make it generally available—for example, I am not going to publicise it but, in essence, I would be able to give it to people or allow people to access it under special conditions.

What are those special conditions? They are that it is anonymised according to the Information Commissioner's guide on that—the code of practice, as it were. There are a number of conditions. The first condition is that you would limit access for a particular purpose to a particular group. The second is that it would be de-identified; there would be data minimisation. The third is that there would be constraints around it which are organisational and technical. The fourth element is that there is a legal contract in place which would stipulate those criteria. The reason for that is, as you go down in your sophistication and granularity, the data becomes more and more identifiable. What you are trying to do is minimise the risk of re-identification, while enabling the AI company, researcher or commissioner to utilise the data to the best effect for society and for the individuals concerned. All the time you are trying to balance benefit from the data and the risks of privacy. That is the way we try to do it.

The Chairman: Do you want to follow up with that at all?

Dr Hugh Harvey: There is a good example recently in America. The National Institutes of Health released a dataset of 100,000 annotated chest X-rays for public use. There were no restrictions to access; therefore, it was freely downloadable. Within two months, researchers from Stanford had developed an algorithm that they claimed could diagnose pneumonia on chest X-rays better than a human could. I think the advantages of releasing data are quite clear.

The Chairman: Do you agree with the distinction that Professor Severs

made between public and general? I thought that was quite interesting.

Dr Hugh Harvey: Absolutely. I am coming on to that. I think we have to be a bit more careful in the NHS. I am of the opinion that all the medical data that the NHS holds is funded by the British taxpayer, and therefore any use of that data should generate benefit back to the British taxpayer—not directly to each person but back into the NHS. While we should open up as much data as possible to specific researchers for specific-use cases, within those data-sharing agreements there should be recourse for a return on investment of that data back into the system that created it.

Dame Fiona Caldicott: I agree with that. It brings me on to mention that I hope in these discussions, and as things develop, we will be taking the public with us. That is absolutely vital. One of the things that worries members of the public is what use their data might be put to that involves making a profit for somebody other than the health service. We have quite a lot of education to do, not least with the professions that look after patients and with the public themselves, in explaining the benefits of this and giving reassurance that it is not going to be profit for companies they do not feel comfortable having access to their data, and making absolutely clear that this is safeguarded through anonymization and that it comes back into the national or public good.

Viscount Ridley: Can I come back to Professor Severs briefly on one point? I hope I understood him right, but in the last session Dr Huppert implied that we are a long way from the point where we have to worry about the trade-off between privacy and usefulness of data. He is looking quizzical at the back of the room, so I must have misunderstood what he said.

Professor Martin Severs: I represent an organisation that has fundamental public assets, and we have to maintain that public trust. I try to stay within both the law and the Information Commissioner's policies. If the Information Commissioner has produced a policy that says "This is how to do it", then that is what I do. I am not prepared to take on the Information Commissioner. If the Information Commissioner says "This is the standard by which you anonymise", then that is the standard by which I anonymise. I am trying to be a good citizen in an arm's-length body and do the best for the public and those who wish to help the public through the use of data.

The Chairman: I am going to bring in the Lord Bishop in a minute. But where does portability sit within your plans under the GDPR?

Professor Martin Severs: I will be fully compliant with the GDPR, the Data Protection Bill and the Act by 25 May 2018, sir. That is my official answer.

The Chairman: Excellent. That could have quite an impact, could it not?

Professor Martin Severs: Looking a long way in advance, the bigger issue is that if a decision is taken automatically you have the right under the GDPR not to be subject to an automatic decision. That means that we have to demonstrate to the public that when a decision is taken automatically—given 10 years hence, or however far hence it is—it is

based on sound evidence and the public are getting a good deal. Although—and I do not want to upset Dame Fiona—it sounds soft to say you want to take the public with you, but it is not; it is very hard and you have to do it properly. You have to be robust and use good evidence to say, “This is the time we can go to the next step”.

Q131 The Lord Bishop of Oxford: I would like you to comment on how you see the readiness of the data we already hold for digital application. The impression I received from the evidence we took in the last session—and these are my words, not those of the witnesses—is that it is somewhat chaotic and massively unready, and that the process of digital readiness is not good and varies enormously. People are still communicating with faxes, data is being stored on paper and there is a huge amount of work to do. Would you like to comment on that?

Dr Hugh Harvey: I proposed a data scale which describes exactly this problem. It is in four parts. Level D data is existing data of unknown quality and quantity, and it is difficult to get to — it is behind firewalls and in silos. Even getting level D data up to the next level, level C, where it is available to start looking at and finding the omissions and clarifications that need to be made, is a huge expenditure. To get it all the way up to level A data, where it is perfectly ready for machine learning and deep-learning algorithms to find patterns, is even more resource-intensive.

There is a lot of focus in the media on the development of algorithms and very little focus on the preparation of data. Medical data, as you have quite rightly pointed out, is very chaotic at source. This comes down to a delay, specifically in the NHS but also across the world, in the technology that is available in healthcare institutions compared to the technology that is available on the high street. We are significantly far behind. My phone is more powerful than many of the computers in the hospital, for instance.

The Chairman: Do you accept the BYOD philosophy put forward by Dr Huppert earlier?

Dr Hugh Harvey: I was not here.

The Chairman: “Bring your own device”.

Dr Hugh Harvey: It depends what it is for. You are not going to bring your own CT scanner, I hope, to a hospital. If you are talking about iPads and mobile phones for communication or some form of data collection, then yes, bring your own device is one way of doing it.

The Chairman: That leads very neatly on to Lord Hollick.

Q132 Lord Hollick: You have all commented on the importance of the NHS and the public generally getting good value for the data when it has been cleaned up. How do we achieve that? Typically, you establish the value of something by having a number of parties on one side and a number of parties on the other side negotiating from a position of some transparency. That is not the case here. You have a small number of buyers who have very deep pockets and a unique understanding of the business model and the way that, over time, they will in fact exploit the

data. How do national health trusts and the National Health Service generally organise themselves to play a meaningful role in that attempt to establish a proper value? I think that that is one of the remits of your organisation, Dame Fiona.

Dame Fiona Caldicott: This is part of the dialogue we have mentioned that we need to have with the public. Members of the public understand that there is a value in the data, certainly to them as individuals in relation to their own care but also given its aggregation and the uses to which it can be put in improving the health of the community and so on. These are things we have to discuss with the members of public who are interested to do so. That is not all 60 million people in the country, but I would say that there is growing interest and concern among the public about these matters, partly because people realise that smartphones and all the other things that have just been discussed can enable them to do all sorts of things which they cannot do in relation to their health data. The difficulty some people have in seeing their record as held by the general practitioner is something about which they write to us, because they should have access and they have difficulty in getting access.

We have to improve the systems through investment, which is beginning, and the use of the technology, which is changing rapidly because the people researching it are able to make progress very rapidly. What we have not done is take the public with us in these discussions, and we really need their views. What is the value? Are they happy for their data to be used when it is anonymised for the purposes we have described? We need to have the public with us on it, otherwise they will be upset that they do not know what is happening to their data and be unwilling to share it with the people to whom they turn for care. That is the last thing we want to happen in our health service.

Professor Martin Severs: As an organisation, NHS Digital would support a national, open and consistent approach. On a personal basis, that is what I think we need. AI cannot exist without large amounts of data, and that data is public data. How do we do that in a sensible way? I do not think that charging for the data is necessarily the right way round. Neither do I necessarily think that IP is the right way round, but it might be. If the public saw that their data was used to develop an AI product and that AI product was discounted for the NHS for the benefit of society, I believe that—subject to Fiona’s discussions—most of society would think that was a fair deal; English society would get something back from that. My view is that I would not have a barrier for entry but I would have some mechanism of demonstrating societal benefit from the data as it is being used. NHS Digital is open to any of those which have a consistent buy-in by all the organisations.

Dr Hugh Harvey: I would agree with that. I do not think there should be a monetary barrier to entry to data access. After all, a lot of people trying to get access to the data will fail in their attempts to make algorithms that work in the clinical workplace. We need to encourage innovation and have failure, and they need to be allowed to fail at low cost. I also think that, as both the other witnesses have said, there needs to be that return on investment of the data back into the NHS.

You asked roughly how valuable do we think the data is. I found two figures in my research. IBM bought a healthcare company in America called Merge Healthcare for \$1 billion. I think that covered about 5 million or 6 million patients. Times that by 10, and that is roughly what we have in the NHS. The Royal Society commissioned a report, and it said that the direct value of data for the public sector is about £1.8 billion and the wider socioeconomic benefits are £6.8 billion at a minimum. We are talking about billions of pounds-worth of value in this data. It would be a shame for the UK economy, especially in times of strain on the funding of the NHS, not to be able to leverage the value of that data for benefit.

Lord Hollick: I have a very small follow-up question. Amazon has recently announced that it intends to move into the prescription business. That would give it details about customers' healthcare and conditions to add to all the other data it has. How would you react to that, if it came to the UK?

Dr Hugh Harvey: I am all for innovation. If anybody can make a big project succeed then it will be someone like the Googles and Amazons of this world. The flipside of the argument is that although we talk about the public sharing their data, Facebook knows more about us than the NHS does. People share that for free, and they are very aware that Facebook is selling their data to advertisers. With the Amazon prescription service, I do not think the general public are going to be particularly upset that their data is shared, if the service works and is better than the current system on offer. I share your concerns over what Amazon could potentially do with that data and whether or not it will benefit the UK taxpayer directly. I am not sure what the business model would be.

Lord Hollick: If it paid tax, would you feel more relaxed about it?

Dr Hugh Harvey: Of course. Absolutely.

The Chairman: That is a big "if". Professor Severs or Dame Fiona, do you have any contrary view?

Professor Martin Severs: No, I do not. I have not heard exactly what Amazon is doing, but if it is acting as a pharmacy outlet de facto—perhaps virtual, perhaps physical—it has to be registered with the CQC. It has to fulfil all the usual regulatory components of being a provider. Therefore, it has to act responsibly. NHS Digital is not a regulator, but Amazon will be regulated. If it wants to do research in medical devices, it will be regulated by the MHRA. If it is a provider of services, it will be regulated by the CQC. We would be entirely neutral in that, as with any other qualified provider.

Lord Hollick: Would the existing regulatory framework protect the issues you are concerned about?

Professor Martin Severs: Yes.

Q133 **Lord Swinfen:** How do you ask the public for permission to use their data? An awful lot of them will not understand, others may not care and some may want to be paid for it.

Dame Fiona Caldicott: Could I tell you a little about what we have done in the preparatory work for the reviews? We have had a number of focus

groups carefully constructed by people who understand how to have a cross-section of the public present in a group, and we have asked them the sorts of questions you are interested in today. People are very willing to have those discussions.

Another approach has been to have citizens' juries, where you take a group of the public who do not know very much about this sort of topic, give them some information about it and then work with them over a period of time—for instance 48 hours over a weekend—to see whether their views about how their data is used change as they learn more and interact with each other. Obviously, there are the historic ways of informing people and getting feedback. We should be using the technology for that purpose.

It is a complicated issue. As I have already indicated, many members of the public are not interested in these matters until they or a member of their family need the service, and then they become very interested. The NHS uses a lot of methods: questionnaires as people leave out-patients or in-patients, questions to their family, questions on patient experience and so on. It is not something I would say that we are expert at, but we are learning how to access the opinion of the public quite rapidly. We have to go on trying. We have to ask them how they want to be informed and to tell us what they want and think.

The Chairman: I am going to bring in Lord Puttnam. If either Professor Severs or Dr Harvey has something to add to what Dame Fiona has said, you can do it at the point where you reply to Lord Puttnam.

Q134 **Lord Puttnam:** Dame Fiona, this is an area where we might be helpful to you. I do not think anyone around this table doubts the need to carry the public with us on the benefits, but we are doing it in a media environment which finds it far more profitable to frighten people than explain the benefits. How do you suggest we even begin to address that?

Dame Fiona Caldicott: I am a psychiatrist by background, so I have worked for many years on the way in which the media portrays the mentally ill and related issues. One way in which I have found it very useful to proceed is to identify those in the media who are at least reasonably open-minded—I will not say positive—about the issues and build alliances with people who are prepared to give balanced and neutral reporting to the public. Obviously, we have to go on relating to those members of the media who are prejudiced and put forward different points of view. We have to keep giving out the messages, as responsible, professional people who have done some research and know the arguments. We cannot be flustered by, let us say, the horror stories told, but must pick them up when appropriate and counter them. Of course, we must work with those members of the public who are willing to be advocates on our behalf. There are many people who are willing to help the health service with some of these difficult moral and ethical questions. There are many ways of doing that, but any suggestions your Lordships are willing to make that we may not have thought of would be helpful to us.

Professor Martin Severs: I think it is important we separate personal data from data which is regarded as anonymous by the ICO. If the data is

regarded as anonymous by the ICO, as an individual I do not have any rights over that data. There is a story to tell to make society comfortable with the use of that data, and to show proper due diligence when there have been errors. When it comes to personal data, we need to move the dialogue to a different place. The public having our own personalised records—you, me, all of us—will change the debate and could start helping both the NHS and social care, because we will start to understand more about ourselves and more about what this data looks like.

Part of the discussion to take the public with us is about the benefits, but it is also about enabling the public's right to have their own records, so that they can have an equal relationship and dialogue with me, as a doctor, around what is wrong with them. That is an important cultural change that will inevitably happen as we move forward.

The Chairman: I am going to move on to Lord Levene's question. Dr Harvey, if you want to come back in, when we answer that question you can add your views.

Lord Levene of Portsoken: My question is about anonymization, and I think we have already answered all that.

The Chairman: I think it follows on rather well from Professor Severs, in a way.

Q135 **Lord Levene of Portsoken:** I will ask it then. Some of our previous witnesses have suggested that data can never be truly anonymised. Is this right? How could NHS data be used safely and securely for the benefit of society?

Professor Martin Severs: I entirely agree with that statement; it cannot be. However, let us be clear, when data goes out in the public—

The Chairman: By the way, your views about the new provision in the Bill—Clause 162: the new criminal offence of re-identification—would be very interesting, in that context.

Professor Martin Severs: I will cover that. First, can data be truly anonymised? No. The risks of re-identification need to be kept as small as possible. For example, if a public dataset goes out, the probability of re-identification would be 0.05 per cent, then 0.014 per cent of that dataset could be re-identified with that probability. In simple language, we take risks every day of our lives. The risk of re-identification needs to be kept small and proportionate to the benefit. We have to have a dialogue of what that proportionality is.

One of the issues here is what happens if things go wrong. I have already expressed how you anonymise, but if somebody transgresses—this is the point you raised about Clause 162—it is a criminal offence, or it will potentially be if the Bill goes through. In essence, a fine is not going to cut it for a big multinational. The big issue here is, if you transgress with the public's data, you will not get the same level of granularity or you may not even get the public's data again. That is a much greater deterrent: not to be able to participate in what you see as good activity is a much greater deterrent.

At the moment, one of the key things that NHS data is looking at is degrees of trust. If we give you the data and you look after it well, we can give you the data again. If we give you the data and you do not look after it well, then we cannot give you the data with as high a risk. We have to start thinking about proportionality and good behaviour. That is a much more powerful, rational way of working in this area. If someone transgresses, the ICO will take action, but the transgressor may not get national data again, or they may get it in a much less granular form which will mean that other companies and researchers will think twice about it.

Dr Hugh Harvey: I agree with those statements but I would like to add perhaps a flipside to the argument. Yes, medical data is not 100 per cent able to be anonymised, but I do not think it should be. If a researcher who has access to a large dataset finds something worrisome about a particular patient, and think that perhaps it has not been picked up before in the clinical realm, they should be able to trace that patient back to source or at least alert the source supplier of the data that patient X on their database needs a clinical opinion. I will give you an example. If I was doing algorithm development on brain CTs and found a brain tumour on one of my datasets, but I did not know who the patient was, I should be able to alert the person who gave that dataset, and say "I found a brain tumour I don't think you knew about". You should be able to trace it back to that patient for those reasons.

The Chairman: Dame Fiona, do you agree with that?

Dame Fiona Caldicott: Yes, I agree with that. We three seem to have a lot in common.

The Chairman: We are trying to stir it up here, you know.

Dame Fiona Caldicott: We will try to think of something controversial.

I want to add to the point made about Amazon. I can see that there are members of the public who would see the system of getting a prompt prescription, with the caveats that Martin Severs has mentioned, as quite straightforward. We have to bear in mind that in many of our medical records there are much more sensitive areas than getting a prescription for a drug that you take regularly. However, there are a lot of people who would not want to get into an interaction with a company such as Amazon. Yes, there may be partly an entry for some of these companies in that area; but I remind you that we are talking about very sensitive data. It is not the same as doing our online banking. These are things we really do not want to share with people we cannot trust.

Q136 **Baroness Greener:** Professor Severs, you said that a financial impact will not make any difference, but saying, "We won't share your data" will. But if a little trust in the NHS tells these huge global players that it will not share its data any more, will they not, as was argued in the previous session, go elsewhere? Why is it any different from the argument you made about the financial penalty?

Professor Martin Severs: Because brand is important. I have found in my dealings that several of the private companies take this as seriously, if not more seriously sometimes, than the public sector. For example,

Amazon underpromises and overdelivers, and Google has its own trend. Bad publicity across a whole population would be damaging to their company and damaging to their brand. If they lose trust in their brand, it starts to unravel.

Baroness Grender: You are not saying you would necessarily withhold data, but it would be public knowledge that they had misused the data.

Professor Martin Severs: Misused health data, yes. My personal view—it is not my organisation’s view because I have not discussed it—would be that they would take that very seriously.

Dame Fiona Caldicott: I want to add a point about the unique situation that this country is in. As I understand it, not being an expert in it, we have the most advanced artificial intelligence research in the world. We also have a National Health Service which is a unique, huge database. That is one of the reasons these companies are interested. I link it to the recent report on the life sciences strategy, because I think we have an opportunity here.

The Chairman: Sir John Bell’s report. That is very interesting.

Q137 **Lord Giddens:** You may feel that you have slightly answered this question, so please interpret it as you will. I will plough through it. Increasingly, AI applications are being developed in specific contexts, using unrepresentative datasets before they are sold in places with very different clinical contexts and patient demographics. What assurances does the NHS look for that these systems are appropriate for use in the British context? Obviously, this occurred with IBM-Watson in the US against a different backdrop.

Dr Hugh Harvey: Absolutely. This is the domain adaptation problem; if you train an algorithm on one dataset it is unlikely to work with the same efficacy on another dataset. We know this. It is a problem with the generalisability of the algorithms.

Going back to your point, Dame Caldicott, about the unique position that we are in in the UK, we have a nationwide nationalised dataset, and if we trained algorithms on the entirety of that dataset they should in theory work anywhere within that system. You could develop an algorithm using data from across the entire country and then use that algorithm in any hospital in the country. If the NHS is going to end up being a consumer of algorithms from abroad, we will have to retrain those algorithms on NHS data anyway, so we might as well take the lead and build them from scratch ourselves.

Lord Giddens: Excuse me interrupting, but surely you would want to make use of global data on health.

Dr Hugh Harvey: Not necessarily, because the epidemiology of the prevalence and incidence of all disease in the UK is different from that in many other countries. Europe and other first-world countries have roughly similar prevalence and incidence, but you could not build an algorithm in India and apply it in the UK. You would get lots of misdiagnoses of tuberculosis, for instance.

The Chairman: Can we get you to disagree with that statement,

Professor Severs?

Professor Martin Severs: No, sir. I think we are at the beginning of this. If you asked whether we had a robust framework yet, I would say that I do not believe we have but I am quite happy to say that NHS Digital would wish to participate. There are five areas that we would want to look at. First, we would want to look at research, and we would want to look at extensibility as well as generalisability: that is, the different data signals have been tested such that you can take it from one trust to another trust, and you would not get a different result, broadly speaking.

The products should be developed with a safety culture. There should be a safety standard for the development at the industry level and there should be a safety standard sign-off when it is deployed. These devices need regulation and registration with the MHRA. When we get into machine learning we need what is called post-marketing surveillance. That sounds MI5-ish, but if the machine is learning and the decisions are being made we need to make sure that the outcomes for patients remain good. What is the process for doing that over time?

Finally, we need to have called business continuity. Let us say that we are using these machines, we are reliant on them and we have a JCB that goes through the electricity line or there is a successful cyberattack on these machines. We need business continuity in the classic sense, and we need what I would describe as a clinical-minimum viable product. We need enough clinicians that if the machine goes down we can still look after the population. We have not worked through these new areas. I believe there is enough knowledge and experience in the UK to get that assurance framework in a more robust state.

The Chairman: You are almost anticipating our next set of questions.

Dame Fiona Caldicott: I have nothing to add to that.

Viscount Ridley: Can I come back to Dr Harvey on this unrepresentative data point? I wonder if we are exaggerating the problem here. I am willing to bet that most of us in this room have taken drugs that went through clinical trials in either India or America and never went near a clinical trial in the UK, for example. Obviously there are issues with drugs being tested on old white men and not young women or people of ethnic minorities, or whatever, but we have ways of dealing with that issue and making sure that it is not a problem. Is it any different for data?

Dr Hugh Harvey: It is slightly different for data. There is a good reason why the FDA, the regulators in America, only accept medical devices and drugs that have been tested on the American population. With artificial intelligence, because the data is so much more granular than a generic pharmaceutical, we do not know the extent to which the data is ungeneralizable. We would be prudent to make sure that we have tested it on as wide a variety of English patients as possible if we are going to deploy the algorithms in England.

You may well be right that we can develop an algorithm outside of England for a specific clinical-use case that works fine here in the UK, but I can think of many examples of diseases across the world that we simply do not see here. You have different genotypes and phenotypes of humans

across the globe that are not prevalent in the UK. If they are included at some representative proportion in algorithms developed abroad, you will get errors in algorithmic outputs when they are applied to a UK population. For reassurance I would like to see algorithmic development done robustly on a UK population, if it is to be used here. You may well be right, and I may be overegging the pudding a little, but that is the safest approach.

The Chairman: He is not always right.

Viscount Ridley: Harsh but fair.

The Chairman: I have to warn the Committee that we have to move on fairly swiftly through the questions.

Q138 **Lord Swinfen:** Does the National Health Service have the capacity to take advantage of the opportunities represented by AI technology and to minimise the risks? Are the clinicians and other healthcare professionals equipped with the necessary skills to take advantage? What could be done to help them?

The Chairman: You seemed to be rather positive on that score, Professor Severs, when you were talking earlier.

Professor Martin Severs: Take the onset of computing in general practice. Let us say that the computers took over helping the GPs with routine prescribing. If we are saying that these technologies will help clinicians with, let us say, the routine and the normal, such as lots of chest X-rays—the radiologist could check them so that you could do the normal scans much more quickly—the clinicians will embrace them. If they say, “We’re going to replace all radiologists next week with AI”, I think you may have a different welcome. If you are proportionate and sensitive to the well-being of the public and you start with the routine and the rather more dull, which computers do really well, you will be warmly embraced by the clinical community. That would be my strong advice.

Dame Fiona Caldicott: One of the areas where we are not so strong is the multidisciplinary approach. We need to develop this for patients and the public, computer scientists and information engineers, alongside the clinicians. One of the effects of the Health and Social Care Act 2012 was that for various reasons we did not demonstrate to those people with technical expertise in the system that they were still valuable to us and we needed them. We lost quite a lot in the intervening period and there is a shortage at the moment. The trusts that are global exemplars in using the new technologies as effectively as possible are finding it difficult to recruit, even though they have the money. The experts believe that a multidisciplinary approach is necessary. I am not sure we are so well equipped with that, although I take Martin’s point about the clinicians, many of whom are willing to embrace this innovation.

The Chairman: As a clinician, how do you feel about that, Dr Harvey?

Dr Hugh Harvey: I entirely agree. The medical syllabus needs to start incorporating not just medical statistics but some basics of data science. An unfortunate side effect of having lost a lot of these researchers to

industry is that now the NHS literally cannot afford to have these data scientists on six-figure salaries come back to the NHS. The answer is a collaborative effort among all the stakeholders in healthcare. Government, the NHS, the medical associations, industry both large and small, and independent researchers need to work on this together to create that true, multidisciplinary approach. If the NHS was to try to do it on its own it would fall short of the relevant skills and funding to do so at a scale and pace that is possible with collaboration.

Q139 **Lord Holmes of Richmond:** Is the use of chatbots and virtual medical assistants to triage patients in healthcare a positive development? Have recent trials in the NHS in this field been successful?

Dr Hugh Harvey: This is my area. I worked at Babylon Health, which I am sure you may have heard of by now, which has a triage app.

The Chairman: You are intimately involved with Babylon?

Dr Hugh Harvey: I worked there for a year. I was in charge of the regulatory department, I went through the regulatory CE marking process for the app and I introduced the post-market surveillance and vigilance system. Does it work? The technology is very safe and effective for triaging patients who can use a mobile phone and read English. It is for a targeted audience. Whether or not it improves patient outcomes and reduces burden on the NHS has yet to be seen. There are ongoing trials in the NHS. I am not yet aware, since I left the company, of the results of those.

There are a few other companies apart from Babylon, such as Ada, and I would imagine that they would also look very carefully at whether their technology improves outcomes. The technology helps people who have access to it, if they use it appropriately, to avoid seeing their GP, if the problems or queries they have are soft enough not to require medical attention. There is definitely potential benefit, but whether or not it has been truly proven yet remains to be seen.

Professor Martin Severs: I am accountable for pathways online, which is based on NHS Pathways. We have a three-line defence. There are 750 different pathways that are based on best evidence. The original pathways telephone trials service, not the online service but the telephone trial service, which supports 14 million cases, was externally evaluated by three universities. They came to the conclusion that the risk of under-triage was 0.034 per cent. That is a strong message to the public that if you use pathways it is going to be unusual for you to definitely end up in the wrong place.

Thirdly, we take external advice from a group called the National Clinical Governance Group, which is made up of the royal colleges. Should we go further? We are looking to explore how to go further with that assurance to the public. I can assure you that I have a three-line defence: best evidence, external researchers, and external advice on an ongoing basis from an independent group of experts who set the standards for doctors.

Dame Fiona Caldicott: Again, I agree.

Q140 **The Lord Bishop of Oxford:** Can you comment on whether you think

new ethical standards or principles for the use of AI in health are needed? Are existing codes of ethics in healthcare sufficient? If new standards are needed, what should they consist of?

Dame Fiona Caldicott: We have a very good system of ethical practice and research ethics across the NHS and related areas. At the moment, I would make a case for having another system. Indeed, you could argue that we have a plethora of bodies; Professor Severs has mentioned some of them. We certainly need to attend to the research evidence that emerges before we introduce things without a solid base. It may emerge as this develops that we need another body or a reconsideration, but my view would be that at the moment we have as much as is appropriate for the pace at which this is going.

Dr Hugh Harvey: I completely agree. The MHRA, the governing body for medical devices, is working on how it can incorporate artificial intelligence into the current framework. It is still a grey area. The FDA in America is also quite unsure and is working with big companies such as Apple and Microsoft to find a solution. The problem with self-learning and self-improving algorithms in a regulatory framework has not been solved yet. There is definitely a framework for a frozen algorithm—i.e. of known sensitivity and accuracy—to be on the market, but for one that self-improves I do not think a regulatory framework has been found anywhere yet.

The Chairman: That is a very interesting point.

Professor Martin Severs: I agree. One of your Lordships made a point earlier about AI developed elsewhere. The only area I am slightly worried about in particular is end of life care or where there are different legislative frameworks in different countries that deal with certain types. This weekend in the British Geriatrics Society newsletter there was clearly a different legal framework or a different societal norm for the treatment of end-stage dementia with PEG-tube feeding. There are areas that AI may get into that are very determined by examination, observation, legal frameworks and patient preferences. As Fiona says, in those types of areas they may evolve some need over time for a medical ethics test or something of that nature that should be part of the normal registration and regulatory framework.

Q141 **Baroness Grender:** What government interventions should there be in policy regulation or investment in order to help the NHS and society benefit from AI?

The Chairman: You have touched on this a little earlier.

Baroness Grender: While answering that could you also consider whether the Government should regulate any aspect of the use of AI in healthcare in particular, and if so, how?

Dr Hugh Harvey: That is a very broad question.

Baroness Grender: Yes, sorry about that.

Dr Hugh Harvey: What should the Government do? The idea of data trusts in Professor Sir John Bell's report is very exciting. I am sure you will have heard in previous evidence which areas to focus on, but in

healthcare in particular we should focus on two areas, and I hope Professor Severs will agree. One should be electronic health record data. The other should be already digitised physiological electronic data such as medical imaging and digital histopathology, and other digitised data such as electrocardiography. These kinds of data warehouses or data trusts would be very valuable in the UK, specifically because, as we said, we have such a nationalised, generalised dataset.

The Chairman: In answer to the Lord Bishop's question earlier, you touched on the algorithm that learns from itself, so to speak. Were you talking more about an ethical framework than regulation?

Dr Hugh Harvey: I was talking more about regulation.

The Chairman: That is what I thought. You have not repeated that. You would advocate that being part of regulation.

Dr Hugh Harvey: Yes, absolutely.

Professor Martin Severs: NHS Digital has its own statutory framework. We are a creature of statute. One of the areas that I would like clarifying is Section 122 of the Care Act, on a purely parochial basis. It says that I can distribute data for health promotion and health services. I want to make sure that if people are building algorithms from a commercial background, that would fall within the scope of Section 122. That is rather parochial but quite important to the facilitation of the future.

Dame Fiona Caldicott: I would emphasise the question of consent. There has been a lot of discussion about this and about legal decisions. In this new area we need, as I have said already, to take the public with us and to inform them about their rights in relation to consent and when, with anonymised data, that is not required, which the public seems to understand. I would like to see a review of the policy on consent and advice to the professions about where we stand on those standards.

The Chairman: That may overlap to some extent with Lord Holmes' final crescendo of a question. I would never want to undersell you, you understand.

Q142 **Lord Holmes of Richmond:** Like any good lawyer. Cutting to the chase, obviously we are doing this report to put recommendations to Government. If there were one recommendation that you would like to see in our report when we complete our inquiry next spring, what would it be?

Dame Fiona Caldicott: It would be to consider the unique potential of this country to develop artificial intelligence and to use this fantastic resource that we have to the benefit of our population.

Professor Martin Severs: For me it is about facilitating the flow of anonymised data, so there is a very clear articulation of what anonymised data is and what it is not.

The Chairman: You were very succinct. You are model witnesses, if I may say so.

Dr Hugh Harvey: I would like to see the Government announce collaborative funding of data trusts with three or four specific areas in

Dr Hugh Harvey, National Data Guardian for Health and Care Dame Fiona Caldicott and NHS Digital – Oral evidence (QQ 128 – 142)

healthcare to focus on—radiology being one of them, because I am biased that way.

The Chairman: Thank you very much. It has been a bit of a gallop, because there has been an awful lot that we wanted to unpack, but you helped us to understand this area. You are clearly grappling already with quite a lot of the issues. This is not all futuristic; it is about the here and now, to quite a large extent. Thank you very much indeed. It has been very interesting.

HM Government – The Rt Hon Matt Hancock MP, Minister of State for Digital, Department for Digital, Culture, Media and Sport (DCMS) and the Rt Hon the Lord Henley, Parliamentary Under Secretary of State, Department for Business, Energy and Industrial Strategy

HM Government – The Rt Hon Matt Hancock MP, Minister of State for Digital, Department for Digital, Culture, Media and Sport (DCMS) and the Rt Hon the Lord Henley, Parliamentary Under Secretary of State, Department for Business, Energy and Industrial Strategy (BEIS) – Oral Evidence (QQ 190 – 200)

Evidence Session No. 20

Heard in Public

Questions 190–200

Tuesday 12 December 2017

[Watch the meeting](#)

Members present: Lord Clement-Jones (The Chairman); Baroness Bakewell; Lord Giddens; Baroness Greender; Lord Hollick; Lord Holmes of Richmond; Bishop of Oxford; Lord Puttnam; Viscount Ridley; Baroness Rock; Lord St John of Bletso.

Examination of witnesses

The Rt Hon. Matt Hancock MP and the Rt Hon. the Lord Henley.

Q190 **The Chairman:** Good afternoon and a very warm welcome to the Rt Hon Matt Hancock, Member of Parliament and Minister of State for Digital in the Department for Digital, Culture, Media and Sport, and the noble Lord, Lord Henley, the Parliamentary Under-Secretary of State at the Department for Business, Energy and Industrial Strategy. Perhaps I should say at the outset that I will refer to Mr Hancock as the Minister and Lord Henley as Lord Henley, so that there is no confusion when I say “Minister”, if that is all right.

This is the 20th formal evidence session for this inquiry. As you know, we are a sessional Select Committee. This session is intended to help the Committee to discuss the Government’s current and potential policy responses to the development and use of artificial intelligence. As you know, we have been taking a keen interest in all the developments that have been taking place both in the industrial strategy and in the Budget on this front, so we are very grateful to you for coming along today.

I should say a few lines of the usual rubric. This session is open to the public and a webcast goes out live as is and is subsequently will be accessible via the parliamentary website. A verbatim transcript will be taken and put on the parliamentary website. A few days after the session you will be sent a copy of the transcript to check for accuracy, and we would be grateful if you could advise us of any corrections as quickly as

HM Government – The Rt Hon Matt Hancock MP, Minister of State for Digital, Department for Digital, Culture, Media and Sport (DCMS) and the Rt Hon the Lord Henley, Parliamentary Under Secretary of State, Department for Business, Energy and Industrial Strate

possible If, after this session, you wish to clarify or amplify any points made during your evidence or you have additional points to make, you are both very welcome to submit supplementary written evidence to us.

Perhaps we could begin with our questions. I will kick off with a very general one. I assume that you have a carefully rehearsed line on who is going to answer which question. Many countries at the forefront of AI development have published AI strategies in the last year. Do you see the AI-related policies announced in the industrial strategy as the UK equivalent of those strategies? If not, do you believe that the UK needs an explicit AI strategy? Perhaps you could also give us a bit of background on how the policies were agreed and what consultations took place in the AI area, as well as some of the ways in which the decisions about the particular focus of the strategy and the funding were thought through.

Matt Hancock MP: Thank you, it is a great pleasure to be here. As you say, we have a carefully scripted response to your insightful question, but that can be improved with a little judicious timing—I am delighted to say that today a report by Oxford Insights has put the UK at the top of the world readiness index for the adoption of artificial intelligence. We have had some confirmation that we are leading the world in being ready for artificial intelligence.

Of course, the development of our AI strategy and its underpinning of the wider industrial strategy was subject to very broad consultation. This is an area where the Government know that there is an incredible amount of work to do across a whole range of areas in order to make the most of the extraordinary new technology, but we also know very well the limitations of what we in government know, so it is very important that we consult broadly and that there are investigations like this Committee's. So I warmly welcome the Committee and I look forward with some anticipation and enthusiasm to its conclusions. Its evidence sessions, which I have followed, have been very positive and insightful.

Lord Henley: I will add a little to that, if I may. As you know, earlier this year we had the Wendy Hall and Jérôme Pesenti review with a number of recommendations that we are considering. You have heard from Wendy Hall already—

The Chairman: —and we are seeing Jérôme tomorrow.

Lord Henley: There were 18 recommendations, some for the Government and some for others. In particular there is the recommendation for an AI council and office, and we are in the process of setting that up.

The industrial strategy—you will see that I have brought my props with me—followed a Green Paper that came out around a year ago. This was followed by consultation all year, with over 2,000 responses. Obviously the Green Paper goes wider than AI. I think the final response came from Lord Heseltine only two weeks before the strategy was published, but it was still a valuable addition.

HM Government – The Rt Hon Matt Hancock MP, Minister of State for Digital, Department for Digital, Culture, Media and Sport (DCMS) and the Rt Hon the Lord Henley, Parliamentary Under Secretary of State, Department for Business, Energy and Industrial Strategy

In the industrial strategy White Paper, as announced, we have made something of AI in that we have produced four grand challenges, the four big changes that we thought were important and that we have to face up to. One is that we address demographic change. The other big one is AI, and we recognise its potential as one of the grand challenges that needs to be addressed. We have identified four priorities that aim to continue to make the UK a global centre for artificial intelligence, so I am very grateful for Matt's announcement.

The Chairman: Do you foresee packaging that up as a separate strategy, breaking it out, if you like, from the industrial strategy in the way that other countries such as Canada have done?

Lord Henley: We are not exactly breaking away from the industrial strategy, but it is one of the big challenges. We are facing a revolution of a sort that we do not know. Talking to someone yesterday I considered the fact that it is 180 or 190 years since the railways came in. At the time, we did not quite know just what sort of changes they would lead to. Similarly, AI mobility and what it might do in the way of autonomous vehicles is one of the other grand challenges we face. It will be an even bigger revolution than railways, aeroplanes and so on. We cannot predict which way it will take us, but we have to be nimble and ready to change in the light of that.

The grand challenge can identify priorities; we want to continue to make sure that the United Kingdom is a major global centre, and I think we are well placed for that. We want to go on supporting individual sectors to improve their productivity through AI; productivity is a major part of the problems that we identified in the industrial strategy. We want to lead the world in the safe and ethical use of data with AI, and we want to help people to develop the appropriate skills, given that people are one of the other major parts of the strategy. So, yes, the industrial strategy is a large part of this.

The Chairman: Did the consultations include looking at the experience of other countries or indeed the experience of our own country with things like the Alvey programme?

Matt Hancock MP: Yes, of course, we have looked right across the world. There is a challenge, because we aim to lead the world in this. The potential impact of AI is very important to the future economy, which is where its underpinning role across the whole of the industrial strategy is important. There is a reason why it has the place that it has in the industrial strategy.

It is also broader than that. If you think about the ethical considerations in relation to AI or its adoption in government and the public sector, which I know you want to talk about, while these things are of course important for improving productivity, which is the underpinning goal of the industrial strategy, there are also huge potential social implications of adopting AI. So, yes, we look across the world, but we also see gaps across the world that no one has yet filled. In particular, it is important to ensure that we have the structures in place to harness the potential of this technology to improve the lot of humanity. That means allowing for

HM Government – The Rt Hon Matt Hancock MP, Minister of State for Digital, Department for Digital, Culture, Media and Sport (DCMS) and the Rt Hon the Lord Henley, Parliamentary Under Secretary of State, Department for Business, Energy and Industrial Strate

and encouraging innovation, as well as making sure that it is used for good and that some of the potential harms are mitigated by thinking about what the ethical framework should be and how it should be put in place.

That was a sub-clause of a long and rather meandering sentence, for which I apologise. No one else is doing that kind of thinking. Looking around the world, you find that we need to get on with it.

The Chairman: We have read some of your speeches on the subject too, and we will pursue that later in the session.

Lord Henley: I wonder if I could add a tiny bit on consultation. Obviously we have discussed this throughout government. I want to make the point that you have invited just the two of us to this session but there is no reason why you could not have had Ministers from BEIS, the Department of Health—

The Chairman: We could have had an entire panel—and, who knows, even the Treasury? Maybe next time.

Lord Hollick: The Government have made it clear that one of the key ingredients of success will be a robust and universal digital infrastructure. In 2016, Ofcom reported that only 498,000 premises were connected to fibre optic, which placed us 57th in the world league of industrialised countries. We had 1.7% connected premises while Japan had 53% and China, which is a large country, had 36%. What is holding us back and what are the Government going to do? The Government have announced in the industrial strategy that they are putting £1 billion into this. How many additional premises will that link up, and how will we get to a position where we have a digital infrastructure that can underpin the very strong position that we have developed in AI?

Matt Hancock MP: On the existing infrastructure for today's needs, we are in fact one of the leading countries in Europe. That is because as a country we went for a superfast fibre-to-the-cabinet rollout, and we are on target for 95% by the end of the year of premises with the availability of superfast 24 megabits per second broadband. Some other countries have gone straight to fibre to the premises; you mentioned two. The challenge for those that do not have full fibre-to-the-premises broadband is that they have very poor broadband speeds. We have taken the approach of bringing superfast broadband to as many people as possible as quickly as possible.

We are now approaching the end of that process. It is very important that we complete the rollout of the existing technology and get decent broadband access to everyone. Now is the time to structure the market to drive out full-fibre delivery. As you say, we have a budget of just over £1 billion to get the full-fibre and 5G rollouts going. The vast majority of this connectivity will be delivered through the market, not through subsidised rollout.

Lord Hollick: Can you explain how the market will do it?

HM Government – The Rt Hon Matt Hancock MP, Minister of State for Digital, Department for Digital, Culture, Media and Sport (DCMS) and the Rt Hon the Lord Henley, Parliamentary Under Secretary of State, Department for Business, Energy and Industrial Strate

Matt Hancock MP: Yes, people will pay for their broadband connections. We will make sure that we have a competitive market that some would say that we do not with copper-to-the-premises technology. That is very important. Our aim is to have a competitive market with many players bringing ultrafast speeds over full-fibre technology. I agree with the thrust of the question: do we need to get on with that? Yes, we do.

Lord Hollick: How far does that £1 billion take us? Just to pick up on Lord Henley's point, we are spending £65 billion and rising on making a 190 year-old infrastructure go faster; it is called a train, HS2. When are the Government going to put serious sums of money behind getting us a 21st-century infrastructure?

Matt Hancock MP: Some £1.7 billion has already been spent. We have over £1 billion in the budget to spend, but I repeat the point that we will not deliver the ubiquitous ultrafast connectivity that we need purely through a government rollout. It is very important that we structure a competitive market properly to deliver these connections.

Lord Hollick: So what number of premises connected to fibre do you expect to rise to over the next five years?

Matt Hancock MP: We have not put a figure on it, and we want as many as possible.

The Chairman: We will move on from there.

Q191 **Lord Holmes of Richmond:** How do you see the recently announced institutions—the centre for data ethics and innovation, the Alan Turing Institute, the Government Office for Science and the industry-led AI council—interacting but not overlapping with one another?

Matt Hancock MP: This is an incredibly important point that we have thought quite hard about. I would start by adding to the list of institutions that need to play complementary roles, because some non-government institutions are also playing in this space. Each has their role. Starting at the hard policy enforcement end, the Information Commissioner has a very important role to play in making sure that artificial intelligence is used in a way that protects people's privacy. We have the Alan Turing Institute, which is very close to the universities; it is essentially the champion of artificial intelligence research and basically reports to BEIS. We have work from the Nuffield Foundation, the Royal Society and the British Academy, which work in a non-governmental way on promoting the ethical use of artificial intelligence. We are now building the centre for data ethics and innovation. It will not be a regulatory body, but it will provide the leadership that will shape how artificial intelligence is used.

All these bodies have a unique role, and defining their terms of reference so that they each add value is important. The proposed office for AI is being proposed as a government departmental organisation, because it will be the joint unit between BEIS and DCMS to ensure that we are joined up at the central government level.

HM Government – The Rt Hon Matt Hancock MP, Minister of State for Digital, Department for Digital, Culture, Media and Sport (DCMS) and the Rt Hon the Lord Henley, Parliamentary Under Secretary of State, Department for Business, Energy and Industrial Strate

We accept the challenge that is clearly implied in your question and we think that the institutional architecture that we have designed is the right one.

We have looked at rolling each of these institutions together to try to make it as efficient as possible, but in each case there is a good reason why not. We do not want the university-led champion of AI research also to be the body that does the thinking on the ethics and framework, because while those are important things, we want the Alan Turing Institute to be able to take on industrial sponsorship, as it does in a big way, and work directly for corporates in developing AI, of which it does a great deal. We want a gap between those institutions. It is tricky, and we have to make sure that we do not have any unnecessary overlap and that we try to avoid underlap. I hope that we have thought about it.

Lord Henley: As well as the office for AI and the Alan Turing Institute, UK Research and Innovation will be part of the landscape. It will be there to ensure that the UK maintains its leading position in research and innovation. Again, it is seeking to deliver something nimble—a word I used earlier—agile, flexible and able to respond strategically. But, again, as Matt made clear, we want to make sure that although DCMS and BEIS are leading here, all other departments should be involved. That, again, is the theme behind our whole industrial strategy in other matters as well as in AI.

The Chairman: There are two follow-up points to make there. There is also Innovate UK, which has skin in the game and money to grant, and the British Business Bank, which by the look of it will have an interest in some of this. That co-ordination will be quite formidable for the office for AI, will it not? What sort of resourcing do you imagine?

Matt Hancock MP: We think it will be resourced by civil servants reporting directly to Ministers. The office for AI is part of government. It is not independent. It is the team that will manage this policy development and architecture.

The Chairman: Where will it sit?

Matt Hancock MP: In between DCMS and BEIS.

The Chairman: Not in the Cabinet Office but in the ether somewhere?

Matt Hancock MP: It will be a joint unit. We have these all over Whitehall. The apprenticeship unit, for instance, used to be a joint unit between BIS and the DfE and it worked extremely well.

The Chairman: Will the Alan Turing Institute have the capacity to carry out the role that you have assigned to it?

Lord Henley: Yes. I do not think you should worry about a lack of co-ordination here.

The Chairman: That was a capacity question, really. I think you started to answer the co-ordination point.

HM Government – The Rt Hon Matt Hancock MP, Minister of State for Digital, Department for Digital, Culture, Media and Sport (DCMS) and the Rt Hon the Lord Henley, Parliamentary Under Secretary of State, Department for Business, Energy and Industrial Strate

Lord Henley: There is a degree of co-ordination, but, with the Alan Turing Institute and others, there is always a question of, “Let a hundred flowers bloom”, although I appreciate who said that.

Viscount Ridley: Chairman Mao.

The Chairman: The provenance is questionable.

Lord Henley: Yes, the provenance is questionable. However, in something like this where we do not really know what will happen, it is best to let a thousand things bloom so that the Government, as long as they remain nimble, can respond in the appropriate way.

The Chairman: You mentioned possibly seven departments. You do not think there should be a single Minister responsible in this area?

Matt Hancock MP: There is a reason why, for the first time, we have a Minister for Digital: right across Whitehall, this extraordinary technology is changing things everywhere. A large part of my job as Minister for Digital is working laterally across Whitehall, whether that is on autonomous vehicles with the DfT, on the better use of data for keeping people alive with the Department of Health, with the Cabinet Office on govtech or with BEIS on industrial strategy indications.

The flip is true for the industrial strategy, which is that it is intentionally a cross-government matter. That is because the technology is by nature breaking down the silos of government, and we have to have an open-minded attitude in responding to that.

Lord Henley: You want to make sure that other departments are thinking along these lines, whether it is in transport or health, because of the changes that are going to come.

The Chairman: So a duo will work, will it?

Matt Hancock MP: It is more than a duo, but I would say that we are the two lead departments on it, BEIS for the application and the wider economy through industrial strategy, and us for the AI sector itself and the digital strategy. We have a joint unit because it naturally falls into both departments, and, as you can see, we have an exceptional ministerial-level relationship.

The Chairman: You are finishing each other’s sentences.

The Lord Bishop of Oxford: Thank you. That is really helpful. One of the dangers that this inquiry has exposed is the gap between what is actually happening, the capacity of AI and the significant change that is coming and public understanding and trust in it. Therefore, developing public trust and confidence in the use of AI in health and many other areas is key. How will the strategies that are being set up enable leadership and good communication in enabling that public-facing trust and confidence?

Matt Hancock MP: That insight is at the core of the need for the centre for data ethics and innovation. The centre was proposed in the Conservative Party manifesto, because we, too, spotted that gap. Whenever any great new technology comes along, it is important that we

HM Government – The Rt Hon Matt Hancock MP, Minister of State for Digital, Department for Digital, Culture, Media and Sport (DCMS) and the Rt Hon the Lord Henley, Parliamentary Under Secretary of State, Department for Business, Energy and Industrial Strate

harness the opportunities while mitigating the risks. Doing so, especially with technology that raises some quite profound ethical questions, requires in some cases new institutional architecture.

The big-picture aim for the centre for data ethics and innovation is best described by an analogy. Two to three decades ago, human fertility technologies and stem-cell research came to the fore. They had great promise for improving healthcare and the treatment of sometimes very ill people, yet they raised quite profound ethical questions. The Human Fertilisation and Embryology Authority was put together to be able to lead not only on the practical questions and issues around the new technology but on the ethical thinking and public advocacy. We want to ensure that the adoption of AI is accompanied, and in some cases led, by a body similarly set up not just with technical experts who know what can be done but with ethicists who understand what should be done so that the gap between those two questions is not omitted. I am delighted that we have now been funded in the Budget in order to set it up. It is incredibly important to ensure that society moves at the same pace as the technology, because this technology moves very fast.

Q192 **Viscount Ridley:** Minister, you mentioned the Oxford Insights study that came out today, which showed that we are number one in AI preparedness. One of the reasons it gave for giving us that accolade was that the Government have done a good job of digitising the process of government services with GOV.UK, the kind of stuff that Francis Maude, Mike Bracken and others have been doing in government over the last few years. We have seen the announcement of £20 million for the GOV.UK catalyst team. So it is clear that the Government could be a big customer in driving AI. What are the opportunities for the deployment of AI in the delivery of public services, what barriers are there to this and what are you doing to remove those barriers? If you will let me, I will come back on the NHS in a minute.

Matt Hancock MP: There are examples of AI already in use in government. If you take the most straightforward usage, chatbots are already in use, including in local government, where you can get more responsive public services because services can give better answers to citizens. Machine learning is already very good at detecting patterns of complex environments to help to deliver answers better and to detect abnormal behaviour such as fraud. So there are already use cases in government. I think there are many more and we should continue to expand its use.

As you say, we in the UK Government have good experience of how to do this, with GDS taking a lead. I had the honour of being the Minister responsible for GDS for a year. It is absolutely clear that you need in government a group of people who are at the vanguard and proposing these changes. I would say that one of the biggest barriers to change is always culture, because it is always much easier to run things the same way you did yesterday rather than differently. You need a group of people who are out there to be champions for the technology, and then it needs to be delivered alongside the policy people who are responsible for

HM Government – The Rt Hon Matt Hancock MP, Minister of State for Digital, Department for Digital, Culture, Media and Sport (DCMS) and the Rt Hon the Lord Henley, Parliamentary Under Secretary of State, Department for Business, Energy and Industrial Strate

that business area who really know the ins and outs of the delivery of that policy. I have given you a couple of examples. I can see it impacting on planning and traffic management, and of course we can talk about healthcare.

Lord Henley: Matt will go back to healthcare later, but I will just add something from my personal experience. My last department was the DWP, in which I had served many years before when it was the DSS. When one looks at the changes in things like methods of payment, one can see that we have moved on enormously. What can come and what is already coming through blockchain and other things such as speed of payment and all that is very exciting, and there is more to come. Do you want to go on to health?

Viscount Ridley: Yes. The question is really about what the Government are currently doing to ensure that the NHS is at the forefront of developing and utilising AI for healthcare while also ensuring that access to public health data is not undervalued or undersold.

Matt Hancock MP: We have a significant opportunity in the UK to improve people's lives and the health of the nation by using the health data appropriately and carefully. It is important that this is done with privacy front of mind. The work that the National Data Guardian for health and social care, Dame Fiona Caldicott, is doing to ensure that we get the innovation yet are also four-square behind privacy is important.

The Data Protection Bill, which is currently before your House, considers these matters and whether primary legislation is in the right place—I know there was a big debate on that during the Committee stage. You have to make sure that the privacy elements are in a good place and that there is also space for and encouragement of innovation.

I have never been a Minister in the DoH, but my experience in the Cabinet Office and in my current role is that the best way to roll out new technologies in the NHS is to demonstrate it through objective clinical results. The model of healthcare that we have in the NHS means that a foundation trust or one of the various kinds of trust uses the data in a way that protects privacy but also demonstrates results. That is the best way of getting the technology taken up.

If you try from top down to drive a particular piece of technology, as we have seen time and again, that will be a struggle. But if from bottom up someone demonstrates that using data in such and such a way can reduce the morbidity rates of a particular problem, it spreads like wildfire through the NHS. At the top level, our job is to give the structure and the permission. Then you want to see clinicians taking up the technology to save people's lives.

Viscount Ridley: Sir John Bell has been fairly scathing about the NHS's attitude to innovations. The NHS is often very slow, particularly in the diagnostics field. Do we think we have solved that, or is it different in the digital field?

HM Government – The Rt Hon Matt Hancock MP, Minister of State for Digital, Department for Digital, Culture, Media and Sport (DCMS) and the Rt Hon the Lord Henley, Parliamentary Under Secretary of State, Department for Business, Energy and Industrial Strate

Matt Hancock MP: I do not think we have solved it at all. I was describing the best way to make it happen from my years of experience at ministerial level. Can it be improved? Absolutely.

Lord Henley: The only thing I would add is to underline the importance of data protection, but a great advantage of having a National Health Service is the quantity and the quality of the data that we have, which other countries do not necessarily have.

The Chairman: We will come on to the data aspects fairly shortly. Will you have procurement specialists within the GovTech Catalyst team, or will they be purely technical in this respect?

Matt Hancock MP: Most of GovTech is about procurement, so it would seem odd not to. Indeed, getting procurement rules right is one of the most important parts of driving improvements in technology through government, because you need the leadership and the permission from the top to drive the change. Then you need to ensure that when new technologies are available they are bid into a solutions-oriented procurement system—my God, I cannot believe I have just used that phrase—that actively encourages new ideas to come forward to solve problems rather than a procurement structure to deliver the same outcome as last time. That is a huge part of driving digital as a whole, including the take-up of AI.

Q193 **Lord St John of Bletso:** It is well known that almost 45% of SMEs have not embraced the potential benefits of AI and are still using pen and paper records. How are the Government's AI-related policies going to make a difference to the non-technology-focused SMEs?

Lord Henley: As I said, the industrial strategy set out the grand challenge of putting AI at the forefront of what we want to achieve. AI will obviously transform business models across many sectors. We cannot direct the SMEs in how we do this, but with the various sector deals that we are hoping to set up in which government and industry can work together, in certain industries, particularly those with large quantities of SMEs, we can provide help and guidance to support SMEs.

We want to continue to build on the research that we already have in place by working within industry to develop new uses of AI, whether it is analytic technologies or whatever. That does not imply that there is a magic wand that BEIS can wave, but I have talked about the focus in the industrial strategy on the four grand challenges and from that on to deals with individual sectors. For example, in the sector deal that we have already announced in construction—sorry, it will be announced just before Christmas, I hope—not only can we look at ways in which we can help the industry but the industry can look at ways in which it can help the many SMEs in that sector.

We also announced in the industrial strategy a degree of funding for our business basics fund. Again, that can look at promoting productivity through boosting individual technologies. One would look at a test-and-learn approach, using sector deals, to see how we can help in one way or another. Obviously, it has to be for the individual business that is, as you

HM Government – The Rt Hon Matt Hancock MP, Minister of State for Digital, Department for Digital, Culture, Media and Sport (DCMS) and the Rt Hon the Lord Henley, Parliamentary Under Secretary of State, Department for Business, Energy and Industrial Strate

put it, still using the quill pen or whatever to think how it can move on. The ones that do not move on will probably go out of business. The ones that bring in new ideas will thrive and flourish.

Lord St John of Bletso: The numbers are quite substantial. One of our witnesses from Sage talked about the lack of digital adaptations leading to what they called a productivity gap. Their research found that businesses are spending 120 days a year doing admin—mundane paperwork tasks—which equates to £34 billion in GDP that we could create or unlock for the economy. It is a substantial opportunity. Obviously government cannot force SMEs to embrace AI, but a lot more can be done to encourage them.

Lord Henley: Productivity was the big weakness that we identified in the UK economy. Part of the problem with productivity is that we have very high rates of employment compared to, say, the French, with their rather high rates of unemployment. The unemployed are not counted in productivity figures, but we would prefer to have people in work. It means that they need to address that, and one way of addressing productivity is going to be to adapt to new technology. If they do not adapt, they will not be there in future.

The Chairman: Can I ask you about the AI council? It may or may not be part of the solution in terms of SMEs, but what sort of representation do you envisage on it? I am assuming that this is going to have users of AI, or potential users, who will not simply be in a conglomeration of US tech giants.

Matt Hancock MP: Yes, it will be broad in representation. There has to be small and medium-sized business representation but also users and developers, and in many cases companies are both. Think of some of the big industrial users and their supply chains: often the big users of AI are the people who are themselves procuring the development. So yes, we expect broad membership.

Q194 **Lord Puttnam:** Are you concerned that the majority of the most promising AI start-ups are being acquired by large US technology companies? If so, what do the Government intend to do about this?

May I just contextualise the question? We have taken a lot of evidence, and quite clearly in the area of basic research we are good. You mentioned earlier that we are number one on readiness take-up, and that does not surprise me at all. This is a very good country into which to sell innovation and it always has been, so take-up has been very rapid.

On the negative side, I would argue that historically we have been very poor at implementation, marketing, brand-building and so on, and at driving and keeping big businesses. The industrial strategies that we have heard about—we had a particularly impressive presentation from Germany—make it clear that Germany, Japan, Korea, France and China all have policies to retain the businesses that they build. Historically, that has not been true here.

HM Government – The Rt Hon Matt Hancock MP, Minister of State for Digital, Department for Digital, Culture, Media and Sport (DCMS) and the Rt Hon the Lord Henley, Parliamentary Under Secretary of State, Department for Business, Energy and Industrial Strate

Last summer the Secretary of State told a cross-party group of us that a Bill was being prepared to clarify the public-interest element of future takeovers. That, from my point of view, is very good news. To take an example that I think is tragic, the sale of ARM, a business built entirely on the back of defence expenditure is now gone. It is hard to imagine an American Government letting a business that had been built on the back of their defence expenditure get sold off to China. When are we going to re-evaluate our capacity for retaining the small businesses that we build and turning them into large and profitable businesses in the UK?

Matt Hancock MP: That is a very important question. The challenge that I would throw back is that one person's sale is another person's investment. What really matters is that the UK is the leader in the world in the activity of using and adapting AI, and that we use that to better the lot of our citizens. That is our job, and you have to keep your eyes on the prize. The investment of Softbank into ARM was one of the biggest foreign investments into the UK in history. One of the consequences of that investment is that it has brought many billions of pounds into the UK investment ecosystem, because some of the founders of ARM who remain based in the UK—indeed, the company is expanding in the UK—are now able to invest in start-ups and growth companies, and there is more liquidity in the market to grow businesses here as a result of that investment.

While it is always nice if that can be done domestically, there are big advantages to inward investment from overseas. One reason why the Silicon Valley ecosystem is so strong is that they got going a generation ahead of us in developing these start-ups, and when people start and then can sell down part of their business they tend to reinvest in tech. This is true in AI and in other parts of the tech industry. That gives you a deeper market. Then the second-generation and third-generation entrepreneurs, who have been through that process and made money out of the investments that they have made out of the money that they made from their original investment, deepen and broaden the available finance to start and grow these businesses.

In the UK, over the last decade we have seen a significant deepening of the investment pool available. A lot of it has come from money that has been invested in the UK in return for the sale of a business to a bigger international player. In a way, we would like to catch up with where the States is, so that we can get these investment into the UK from UK entities. I think it takes a generation to do that. One of the biggest government impacts that you could have on this is making sure that that finance, particularly growth capital, is available.

The strongest policy tool that we have in this area is to have the British Business Bank investing side by side with private investment organisations to back British start-ups and growth businesses and use the enterprise investment scheme to incentivise investment into this space. I was in BIS when we started the British Business Bank. It has gone from strength to strength and is increasingly in the venture capital and start-up growth funding area rather than just infrastructure, where it first started. We have backed it in Autumn Statements and Budgets in a very

HM Government – The Rt Hon Matt Hancock MP, Minister of State for Digital, Department for Digital, Culture, Media and Sport (DCMS) and the Rt Hon the Lord Henley, Parliamentary Under Secretary of State, Department for Business, Energy and Industrial Strate

big way, and we are clear that if we do not remain in some of the European schemes then we will more than match that with BBB investment. So there is a role for the Government. Approximately 40% of capital at some stages in the market is backed in some way by the British Business Bank, so we should not underestimate the role that we are playing in this space.

Lord Puttnam: Before Lord Henley answers, I think I understand that, but why do so many of our western European and indeed Asian competitors have a different strategy? Why do we not have things like golden shares that would allow us to hold on to companies where we think the public interest is not well served by them being sold off to the highest bidder?

Matt Hancock MP: Because where there is a good case for investment from overseas, we should welcome that investment into the UK. We have benefited enormously and are a leading player in this space because we are a global leader in welcoming investment, and sometimes welcoming investment means allowing people to buy the equity.

Lord Henley: Then you have to think, as Matt Hancock said, about the amount of inward investment that we get in order to understand that.

On our review of patient capital and the most recent Budget, I saw Sir John Bell after the Budget, and his excitement at the apparently small changes to the investment that pension funds could make was enormous. He suddenly realised the amount of money that could come out of those into sectors such as this. However, it is not just sectors such as this that Sir John Bell is interested in; it is also pharma and others. I am not sure what words I should use to describe Sir John Bell's excitement, but he thought this was one of the best things that could happen.

It is not necessarily for the Government to decide to go into those businesses; it is for other people with money to decide. As you and others have said, we are jolly good at creating small businesses. Let us now hope that they can invest in those small businesses and let them grow.

Lord Puttnam: Yet our economy is nothing like as successful as Germany's.

Matt Hancock MP: That is a moot point and probably the subject of a different inquiry.

Lord Hollick: Rupert Murdoch is usually regarded as rather a dominant player in the markets in which he plays, but he has recently decided that he is too small; that when it comes to competing with the Silicon Valley behemoths, he needs to be bulked up. He has pointed out that they seem to be lightly taxed, lightly regulated or unregulated, and not subject to the usual rules of market dominance. For instance, Google has 70%-plus of the search market. That is a fairly good indication that the power that Silicon Valley digital companies have is very significant, not just in relation to buying up small companies but in relation to putting established companies at a competitive disadvantage.

HM Government – The Rt Hon Matt Hancock MP, Minister of State for Digital, Department for Digital, Culture, Media and Sport (DCMS) and the Rt Hon the Lord Henley, Parliamentary Under Secretary of State, Department for Business, Energy and Industrial Strate

The question is: are the Government thinking about the steps that they need to take to reduce the competitive dominance of these companies, companies that often exist not nationally but internationally? Do you see that as a problem and, if so, what might the solution be?

Matt Hancock MP: It is a live question. What really matters here is ensuring that the markets are competitive. This also goes back to the question about SMEs, where ultimately competitive pressures will improve the adoption of good and positive new technologies. The question is: can a market be entered by a new entrant? Many of these markets are highly competitive, even if the nature of the networks and the way the data works mean that there are some very large players. These companies are both very large and very competitive at the same time, but having a strong competition policy is important. That is the direct answer to the competition point that you raised.

However, you raised a broader point about our attitude to the really big companies, especially where they have very big sets of data and a big impact on people's lives. We think that the founding principles of many of these companies, which were essentially libertarian and said essentially that connecting more people is unambiguously a good thing, need to be tempered by the principle of not doing harm to others. This applies both to the area that you mentioned and more broadly. It applies with respect to the protection of intellectual property rights online, to some of the potential harms to children, and to the discussion about protection against terrorist communication online. Our attitude to the big companies should essentially be liberal: that is, "We promote your freedom but not the freedom to trample on the freedom of others".

Thankfully, we have thought over the last couple of hundred years about this balance offline, so although there are some significant brains on that side of the table, in this discussion we do not need only ourselves; we can rely on great thinkers of history, because a lot of political philosophy is about how you tackle this problem. We need to draw on that and apply it to a new technology so that we can promote and encourage the amazing freedom that is brought about by this technology while also mitigating the harms. I think that applies in this area as well as elsewhere.

Q195 **Baroness Rock:** We have heard from a lot of people in this sector about the desire to hire the best people they can for the sector. If I may raise the question of Brexit and immigration, what are the Government doing to ensure that Brexit does not have a detrimental impact on the development of AI by UK businesses and universities? In particular, what are they doing to ensure that the most skilled individuals can come to the UK and find work within the sector?

Matt Hancock MP: Thank you very much. Thank you, Lord Henley.

The Chairman: We like to see competition between our witnesses in answering the questions.

Lord Henley: On Brexit, we are making progress. We had the announcement from the Prime Minister last week and the European

HM Government – The Rt Hon Matt Hancock MP, Minister of State for Digital, Department for Digital, Culture, Media and Sport (DCMS) and the Rt Hon the Lord Henley, Parliamentary Under Secretary of State, Department for Business, Energy and Industrial Strate

Council this week. Obviously nothing is certain. You talk about hiring and firing, and these things will have to be sorted out.

Take something like AI or other successful parts of the UK economy; I am thinking of an inquiry into the advertising industry that I did before I went back into government. At the moment, advertising attracts a lot of people from the EU. It also attracts just as many people from all around the world, because the advertising industry is one of our great successes. As the Minister has explained, and I fully agree with him, it is one of our great successes and it could be even more successful. People always come to successful industries and we will find ways of letting them in, just as we did before we were in the EU and as we will after we leave the EU. I still hope that we will make sure that we get a good deal. Last week's activities and what happens at the Council this week are very important and will take us further.

I am an eternal optimist in these matters. I hope I do not sound Pollyanna-ish, but we will get there in the end. It is just that we do not quite know where it is at this stage.

Matt Hancock MP: We have also announced in the last month that we are doubling the number of exceptional-talent visas. That is a very practical example of policy backing up our wish to continue to attract the brightest and the best from around the world. You will also note the section on attracting and growing our own talent in the Hall-Pesenti report, which was important. Of course, growing our own talent and attracting the brightest and the best go hand in hand. So Brexit is very important, but I share the noble Lord's optimism.

Lord Henley: There were also Budget announcements of further money going into training and other matters. We will have the right people here, and we want to attract people from all around the world.

Viscount Ridley: I have a quick supplementary question, if I may. It has been suggested to us by Eileen Burbidge that AI-related roles should be included in the tier 2 shortage of occupation list. You mentioned tier 1 visas. Might this be thought of by the Home Office Migration Advisory Committee as another sector for tier 2 visas? Would you be interested in passing that on to the Home Office?

Matt Hancock MP: I am always happy to be a messenger.

The Chairman: I am sure you always have great influence with the Home Office.

Matt Hancock MP: We were delighted by the doubling of the exceptional talent route, and the Pesenti-Hall report makes it clear that this is a very important area. After all, to remain the best in the world, as we are delighted to say we are in terms of readiness, we need to continue to attract the brightest and the best.

Q196 **The Lord Bishop of Oxford:** Minister, you spoke very eloquently earlier about the leadership role of the Centre for Data Ethics and Innovation in building public trust and confidence. Could you expand a bit more on the vision for the centre? In relation to the Royal Society's request for a data

HM Government – The Rt Hon Matt Hancock MP, Minister of State for Digital, Department for Digital, Culture, Media and Sport (DCMS) and the Rt Hon the Lord Henley, Parliamentary Under Secretary of State, Department for Business, Energy and Industrial Strate

stewardship body, is it filling a similar role? What role might it have in the regulation or formulation of regulation of AI? What is your timetable for further announcements about its remit and about when it will be set up?

Matt Hancock MP: The answer to the last question is shortly. We got the money through only three weeks ago, so we are proceeding at pace in setting it up. Quite rightly, it aims to fulfil the requirement demanded by the Royal Society and the British Academy, as well as the manifesto commitment that I mentioned. I am always slightly cautious about responding to flattery by going further, but there is a really important set of roles here.

The first is doing the thinking with the right people who combine the understanding of the technology and the understanding of the requirement for an ethical framework and pushing and publicising that work, including being public advocates of the benefits of this technology within reasonable bounds. That is foursquare at the heart of what it needs to do. We have chosen to make it not a regulatory body but rather a body that recommends changes to policy as and where appropriate, rather like the National Infrastructure Commission, which, by the way, is publicising a report on the use of AI in infrastructure this week; I have not seen it but I know that it is coming. It is quite similar to the Human Fertilisation and Embryology Authority, which I mentioned earlier, but it is more like the NIC, which recommends changes rather than executing them. We propose that it has a role in the development of data trusts. This is a key recommendation from the review: to propose trusted mechanisms to make it easier for organisations to understand how to use and share data for AI safely and securely, because the bigger your dataset, the more powerful your AI can be.

The Chairman: Will it oversee the data trust? Is that what is envisaged?

Matt Hancock MP: Potentially. Certainly governance has been recommended for that. I do not want to condemn it by saying that it must run them itself, but it should certainly oversee their development.

There is also the promotion of standards, because much of the modern use of data involves different datasets talking to each other in a secure way and returning results with minimal transfer of data itself. That is good for innovation but also good for data privacy. That requires standards, and it is best that those are not written by regulatory organisations, because then they become crusty. They can be rather more fleet of foot if they are written by an organisation at arm's length that is not itself a regulator.

These are all things that we find industry players want to know the answers to but do not necessarily themselves want to answer. Developers of autonomous vehicles do not necessarily want to answer the ethical questions required in an algorithm for autonomous vehicles that are implicit in a human driver but are made explicit by the fact that this driving is done by an algorithm. It is perfectly reasonable for the corporates to say that that is a social decision, not their commercial decision, but we need an answer on how we should structure it.

HM Government – The Rt Hon Matt Hancock MP, Minister of State for Digital, Department for Digital, Culture, Media and Sport (DCMS) and the Rt Hon the Lord Henley, Parliamentary Under Secretary of State, Department for Business, Energy and Industrial Strate

The same is true, at a more practical and prosaic level, for standards for the transfer and use of data. People need to know what the standards are, and sometimes you need a convening body herding the cats and taking a leadership role in that, rather than just waiting for different companies to come together.

The Chairman: But you still see a continuing role in that space—for the ICO, presumably—once the standards have been formulated. Would it be the ICO that would enforce them?

Matt Hancock MP: Only if they are statutory standards. In many cases data transfer is based on standards that are of course within the law and compatible with the ICO regulations but are drawn up in a way that benefits the industry. If an industry wants to get together and ensure that data is shareable according to certain standards, we do not want to be writing the standards ourselves, nor do we want the ICO writing them. It is important that they meet the ICO's regulatory requirements, but it is reasonable for the industry to come to a conclusion on exactly how the data is shared and used. In most industries everyone has an angle, so it is best if another body can come forward with these.

The Lord Bishop of Oxford: On education and curricula, we heard in our previous evidence session about the importance of combining ethics with computer studies and programming, and primary and secondary skills. Would its remit extend so far?

Matt Hancock MP: That is more directly a policy question for the DfE. We do not want to encumber it with so much at the start that we are trying to boil the ocean. However, it is certainly a very important issue that I am sure anyone in this area would have a view on.

Q197 **Lord Giddens:** How would you assess the possible impact of AI on the labour market and the future of work? Are you considering how to cope with this now, or is it still too early? As the Minister will know, there is a discussion about this across the world and there is already a massive literature on this. As an economist, he will know that there are articles with titles like, "Will Humans Go the Way of the Horse?" It would be good to hear you comment on this. I do not think one should just repeat the homilies that we have full employment now or that new jobs have always been created in the past, because this could be structurally quite different.

Lord Henley: It is very worrying for a great many people. I will put myself into that category. Two of my children are employed in the insurance world. I recently saw some people from the insurance world, who talked about the changes that could come, and they seem to involve an awful lot of job losses in certain bits of insurance. So one thinks, "Gosh, they've only just started there. Are they going to have to move on to something else, or will there be other things coming on?" Obviously the Government should be doing something now. That is why we asked the CEO of Siemens UK to do the Made Smarter review, looking at that question. He looked at the impact that the digitalisation of industry would

HM Government – The Rt Hon Matt Hancock MP, Minister of State for Digital, Department for Digital, Culture, Media and Sport (DCMS) and the Rt Hon the Lord Henley, Parliamentary Under Secretary of State, Department for Business, Energy and Industrial Strate

have and saw that although there would be job losses there would be a net gain of a great many jobs in his particular field.

Yes, there will be job losses. However, when the car came in a lot of ostlers lost their jobs but people went on having horses—in fact, I think the number of horses went up. But there were changes. Changes happened with the railways. With every industrial revolution there will be changes. It is right that we should look at the possibilities, hence the Made Smarter review. What those changes will be we do not yet know. Like you, I think we should be worried about certain areas, but we should also be optimistic about what will also come about. It is not for us to hold back the changes that are inevitably going to happen; rather, we should welcome those changes and ensure that the new opportunities, such as those that were seen in the Made Smarter review, arise.

Matt Hancock MP: I agree entirely with the view expressed. The question of ensuring that as a society we support those who are disrupted as well as supporting the disruption is incredibly important. The change is happening anyway. The question is not, “Will we lose the jobs that can be replaced by automation?” They will go. The question is, “Are we doing what is necessary to build, develop and attract the new jobs that will come?” We need to focus on supporting those whose jobs are disrupted. There is a wonderful phrase to describe what we should aim for, which is to automate the work and humanise the jobs. I am a profound believer that there are things that humans do that only humans will do, and it is human tenacity and capability that will generate new ideas and new jobs.

This is a concern that has run through the ages. Keynes talked about technological unemployment. Before Harold Wilson got on to the “white heat of technology” while he was in government, in opposition he talked about the end of employment because of technology. In fact, when I was researching this I discovered that I am descended from a leader of the Luddites in the early 19th century: Richard Hancock led a mob of 1,000 who smashed up looms throughout Nottinghamshire because he came from a family of hand-weavers. I hope that while we Hancocks remain as sensitive to those whose jobs are disrupted by technology, we also recognise these days that violence is not the answer; the best thing to do is to make sure that we are also leading the world in the development of the new technology.

Lord Giddens: I think it is important that government considers the possibility that things may be different this time. It does not follow that because it happened in the past it will happen again. We had 60% of the population working in agriculture. Now it is 1%. We had 40% in manufacturing in this country. Now is 8% or 9%. We only have service occupations left. If a large number of those jobs are automated, it is not completely clear that the traditional things that Keynes and others said will hold. Therefore, I hope that the Government will look at the extensive literature and seriously consider the risks as well as just repeating the idea that there will always be jobs.

HM Government – The Rt Hon Matt Hancock MP, Minister of State for Digital, Department for Digital, Culture, Media and Sport (DCMS) and the Rt Hon the Lord Henley, Parliamentary Under Secretary of State, Department for Business, Energy and Industrial Strate

Lord Henley: Can we send Lord Giddens a copy of the Made Smarter review? He might find it interesting reading.

Matt Hancock MP: The review is important, but I can reassure you that we take the literature very seriously and are engaged in the debate on it. Whether or not this time is different and this technology is materially different does not change what we need to do now, which is to ensure that we are developing new jobs and that people get the training they need to be able to cope with the changes that are coming about. That is all part of the productivity challenge that my noble friend spoke so powerfully about earlier. The big question is: do we as a country understand that this is happening and have policies in place to ensure that the technology benefits people? I would say that that is very much what we are focused on.

Q198 **Lord St John of Bletso:** You mentioned that it is our objective to lead the world in AI.

Matt Hancock MP: Did I mention that?

Lord St John of Bletso: Yes, many times. What discussions have you had with other countries about AI? You also spoke about the need for standards. Do you think that there is a need for an international AI effort to agree standards on such things as ethics and the use of AI? How could Britain lead that effort?

Matt Hancock MP: Yes, yes and yes we can. Ministers from across government repeatedly have discussions on this question internationally. I have attended G20 and G7 discussions on this subject. It is interesting that in those discussions the UK has a leading part to play in how the developed and developing world comes to terms with and maximises the benefits of this technology. As I think I said earlier, it is interesting that, while of course we learn from other countries, one thing that we have learned is that there are gaps that we need to fill.

Lord St John of Bletso: It is not just the world of AI; we are also leading the world in the use of robotics. When I was in University Hospitals Birmingham recently, I saw that it was embracing robotics in a major way. One company in particular has done extraordinarily well on the stock market in getting a lot of robotics contracts, not just in the UK but right across Europe.

Matt Hancock MP: Yes. This is a subject of constant discussion with Ministers around the world.

Q199 **Viscount Ridley:** We had a fascinating—spooky, I should say—session last week on autonomous weapons. We understand that in 2011 the MoD defined autonomous weapons as needing to be capable of understanding higher-level intent and direction. We have heard that this is an extremely high threshold, out of step with those used by other countries, and is likely to exclude almost any weapons system into the foreseeable future. Is this still the Government's working definition for autonomous weapons? If so, should it be adapted to reflect definitions used by the rest of the world? I am sorry that this MoD question comes at you.

HM Government – The Rt Hon Matt Hancock MP, Minister of State for Digital, Department for Digital, Culture, Media and Sport (DCMS) and the Rt Hon the Lord Henley, Parliamentary Under Secretary of State, Department for Business, Energy and Industrial Strate

Matt Hancock MP: Well, it is an MoD question, but in a way it demonstrates how broadly the technology is impacting government and indeed life. There is not an internationally agreed definition of lethal autonomous weapons systems. We think that the existing provisions of international humanitarian law are sufficient to regulate the use of weapons systems that might be developed in the future. Of course, having a strong system and developing it internationally within the UN Convention on Certain Conventional Weapons is the right way to discuss the issue. Progress was made in Geneva by the group of government experts just last month. It is an important area that we have to get right.

The Chairman: Thank you. That was very concise. Our final question is from Lord Hollick.

Q200 **Lord Hollick:** AI throws up some new regulatory challenges, particularly for explainability and liability. Do you believe the existing sectoral regulators, insofar as they exist, can take on the responsibility for understanding and regulating AI, or do you think there is a need for a specialist AI regulator? To tack a question on to that, you talked earlier about the cornucopia of public data, and the NHS has been a very good example. How do we ensure that the public and the NHS get proper value for that?

Matt Hancock MP: Do you want to go first?

Lord Henley: On regulation? No.

Matt Hancock MP: Technology is changing the way things happen in so many different industries that right across the piece we need to ensure that regulations, which were often designed before this technology even envisaged, are updated. There are areas where we have had some very good successes. The FCA's regulatory sandbox is a good and often-cited example of forward-looking regulation. In a way, there is already a regulator that regulates the core of AI, which is of course the Information Commissioner, because it regulates data usage and ultimately AI is a big-data technology. So I would argue that that exists and that in fact developments in AI underpin a lot of the discussions that we are having on the Data Protection Bill that is currently before the Lords.

In the same way as we are going through the regulation of autonomous vehicles, which is also in front of Parliament, we will have to go through many different areas of life and ensure that the regulations are updated to take account of AI. One of the biggest challenges is that technology is moving fast, so we have to have smart regulators who can have a dialogue with the cutting-edge industries, and this is what the sandbox approach is all about: having an open-book approach between the regulator and the regulated, with high-quality people in the regulator essentially saying of a new innovation, "Yes, so long as you show me how it's working, and then if we need to put in place a more formal structure then we'll do that as we see the results", rather than the default being, "Not unless you can prove that it works", which sometimes leads to a chicken-and-egg problem.

HM Government – The Rt Hon Matt Hancock MP, Minister of State for Digital, Department for Digital, Culture, Media and Sport (DCMS) and the Rt Hon the Lord Henley, Parliamentary Under Secretary of State, Department for Business, Energy and Industrial Strate

An attitudinal shift is needed in regulators right across the piece, and in many cases also a legislative or regulatory shift. We had made a start on that, especially in finance and transport, and we have made some progress in health, but there is more to do. A fund was set up in the Budget, which BEIS is going to lead on: the Regulators' Pioneer Fund is a £10 million fund enabling the regulators to support start-ups and small companies with pioneering, innovative approaches, building on the excellent model set up by the FCA sandbox.

Viscount Ridley: Can I come in with a supplementary question? We have heard a lot of evidence either way on the question of whether there should be a specific regulator or whether an existing regulator should be used, but I want to get your reaction to this: I suggest that AI is much less a sector needing its own regulator—as is the fertility industry, which was mentioned earlier—than a pervasive technology such as, say, electricity, which goes across the whole piece. That is an important distinction to bear in mind.

Matt Hancock MP: That is mission-critical, and it is why in the industrial strategy there is no separate section for digital or AI; it underpins everything. While the ICO will regulate data and the data underpinnings of AI, actually every regulator has to understand the impact of this new technology on their sector.

Lord Henley: All I can add is that the Data Protection Bill is going through at the moment, but that is not going to be the last word on all this. We will have further data protection Bills in a few years' time, because things are changing. Similarly, AI is an area where things are changing and are going to change very fast. There is nothing that we can do to stop that. It has happened and it is happening, and we will have to change.

The Chairman: And there was I thinking that you were going to accept all my amendments, but there we are.

Matt Hancock MP: Actually I thought this Data Protection Bill would be enough.

Lord Puttnam: I was thrilled to bits by what the Minister had to say about the quality of regulation and of regulators. We have tabled an amendment to the Data Protection Bill to remove the salary cap from the ICO. Would he support that? I cannot see how the quality of regulation that he envisages can possibly be achieved if there is a salary cap on the wages paid to the best possible lawyers that you are going to need in order to get the regulation you want.

The Chairman: To add a supplementary to that, is the ICO sufficiently resourced to do the job that we are expecting of it?

Matt Hancock MP: The ICO is an incredibly important part of getting the new Data Protection Act into place and supporting companies through the changes, and we need to make sure that the Information Commissioner has all the support that she needs to do that. I think you make a very good point.

HM Government – The Rt Hon Matt Hancock MP, Minister of State for Digital, Department for Digital, Culture, Media and Sport (DCMS) and the Rt Hon the Lord Henley, Parliamentary Under Secretary of State, Department for Business, Energy and Industrial Strate

Lord Puttnam: It is a cross-party amendment, incidentally. It is not politicised. It is a practical amendment.

Matt Hancock MP: Indeed.

The Chairman: Your reply is carefully noted. Thank you both very much indeed. It is good to see that the new generation of Hancocks is much more enlightened than the old and making up for the sins of the past.

Matt Hancock MP: Richard Hancock got sent to Australia. Still, these days we work very well with the Australians on this and other subjects.

The Chairman: Thank you very much. It has been a very good session.

HM Government –The Rt Hon the Lord Henley, Parliamentary Under Secretary of State, Department for Business, Energy and Industrial Strategy (BEIS) and The Rt Hon Matt Hancock MP, Minister of State for Digital, Department for Digital, Culture, Media and Sport

**HM Government –The Rt Hon the Lord Henley,
Parliamentary Under Secretary of State, Department for
Business, Energy and Industrial Strategy (BEIS) and The
Rt Hon Matt Hancock MP, Minister of State for Digital,
Department for Digital, Culture, Media and Sport (DCMS)
– Oral Evidence (QQ 190 – 200)**

[Transcript to be found under HM Government – The Rt Hon Matt Hancock MP, Minister of State for Digital, Department for Digital, Culture, Media and Sport \(DCMS\) and the Rt Hon the Lord Henley, Parliamentary Under Secretary of State, Department for Business, Energy and Industrial Strategy \(BEIS\)](#)

Dr Julian Huppert, PHG Foundation and Understanding Patient Data, Wellcome Trust – Oral evidence (QQ 116–127)

Dr Julian Huppert, PHG Foundation and Understanding Patient Data, Wellcome Trust – Oral evidence (QQ 116–127)

Evidence Session No. 13

Heard in Public

Questions 116–127

Tuesday 21 November 2017

[Watch the meeting](#)

Members present: Lord Clement-Jones (The Chairman); Lord Giddens; Baroness Greener; Lord Hollick; Lord Holmes of Richmond; Lord Levene of Portsoken; The Lord Bishop of Oxford; Lord Puttnam; Viscount Ridley; Lord Swinfen.

Examination of witnesses

Dr Julian Huppert, Dr Sobia Raza and Nicola Perrin.

Q116 **The Chairman:** Good afternoon and a very warm welcome to our witnesses. This is the 13th formal evidence session for our inquiry. This session is intended to help the Committee discuss the application of artificial intelligence to healthcare in the UK. Welcome to Dr Julian Huppert, the chair of the Health Independent Review Panel for DeepMind Health; Dr Sobia Raza, head of science at the PHG Foundation; and Nicola Perrin, the head of Understanding Patient Data at the Wellcome Trust. Welcome, to coin a phrase.

I have a little rubric I need to go through before we ask you to give evidence. The session is open to the public. A webcast of the session goes out live and is subsequently accessible via the parliamentary website. A verbatim transcript will be taken of your evidence. This will be put on the parliamentary website. A few days after this evidence session, you will be sent a copy of the transcript to check for accuracy. We would be very grateful if you could advise us of any corrections as quickly as possible. If, after this session, you wish to clarify or amplify any points made during your evidence, or have any additional points to make, you are very welcome to submit supplementary written evidence to us. Would you like to introduce yourselves, starting with Nicola?

Nicola Perrin: I am Nicola Perrin and I lead the Understanding Patient Data initiative at Wellcome. Understanding Patient Data was set up to support better conversations about the uses of health information. Our aim is to explain how and why patient data is used, how it is kept safe, what is allowed and what is not allowed. We work with patients, charities and healthcare professionals to champion responsible uses of data. We are set up and hosted by the Wellcome Trust, as I have said, with a broader range of funders, including the Department of Health and Public

Health England. One of the issues that we are looking at is the impact of new digital technologies, including artificial intelligence, and their implications in healthcare and the implications for public confidence.

Dr Sobia Raza: I am Sobia Raza. I work at the PHG Foundation, which is a non-profit independent health policy think tank with a special focus on how emerging technologies are translated to deliver more effective healthcare and deliver improvements in health for patients and citizens. As the head of science, I lead the foundation's scientific and technical review of subject areas and work within a multidisciplinary team with ethics, regulatory and public health experts.

Dr Julian Huppert: Thank you. It is a pleasure to be back here. I am Julian Huppert. My day job is as the Director of the Intellectual Forum at Jesus College in Cambridge. We have looked at a number of issues including AI and various consequences of it. I am not speaking for the College. I am here, I think, because I also chair the Independent Review Panel for DeepMind Health—a very new approach to governance and accountability—which I hope I will have the opportunity to talk a bit more about with you later on. It is something we are still developing.

Q117 **The Chairman:** Thank you all very much. I am going to start with a fairly broad question. To what extent is AI already used in healthcare and where in health do you see the biggest potential for the use of AI? You may along the way in answering this give some idea of what you think about the Government's recent AI review, if that is relevant to the sort of developments that you will be talking about. Shall we start with Nicola?

Nicola Perrin: There are huge opportunities for AI in healthcare. We are beginning to see AI used in a few places, in pockets, but there is definitely potential for much more. One of the key things to emphasise is that all the applications at the moment, and for the foreseeable future, complement but do not replace the clinician. It is important to be clear about that. This is not about putting doctors out of work.

One of the biggest potentials initially will be in image analysis where AI is already being used. I know you are going to hear a lot about that later today so I will not dwell on it, but algorithms can already differentiate both more quickly and more accurately cancerous and healthy tissue. That will be one of the things that we will see most quickly.

There are also some other uses of AI already, such as clinical-decision support tools and helping clinicians by spotting patterns and making predictions so that they can calculate risk and reach a diagnosis more quickly for patients. There are also some examples in the self-management area such as helping patients choose their best care path. For example, Arthritis Research UK is using a new virtual assistant. That helps answer patients' questions about arthritis, learning from the questions they are asking to tailor the answers to them.

There is also potential in drug discovery. We are beginning to see research to mine disease databases. You can use algorithms to identify new targets for drug development, and to look at how one might re-purpose drugs—taking drugs that have already been through clinical trials and looking at using them for other conditions, or drugs that did not

succeed in a clinical trial and thinking about whether there are other diseases that they could be used for. We are just at the beginning, I think.

The Chairman: Obviously you will talk through some of the developments that you are particularly interested in, Dr Raza, but do you agree with the analysis that it is complementary rather than about replacing in particular?

Dr Sobia Raza: In the immediate term, yes; I would concur with Nicola. Some of this comes down to the processes that are used for implementing these tools, but generally, for now, it is complementary and assistive rather than a replacement technology. Some of the rhetoric in the press would certainly suggest otherwise, but right now I would reflect on some of the points Nicola made; as a technology it is emerging within healthcare and it is by no means expansive or at the stage where it is ready to replace human decision-making in interventions.

The Chairman: Are the examples that Nicola has given the ones you would put forward yourself?

Dr Sobia Raza: In the near term at least, I would concur that medical imaging analysis is certainly the lower-hanging fruit where AI could have quite significant impact—for example, in radiology and radiotherapy in terms of planning for oncology treatments. Further down the line—although this will take some time—as we begin to collect richer datasets, there is great potential to apply that to refine our understanding around disease. That could be used to stratify populations and, essentially, provide more tailored care, but that is some way down the line yet.

The Chairman: Dr Huppert, you are on the advisory board for DeepMind Health. Does that give you an insight into the sorts of areas that it is going into?

Dr Julian Huppert: Yes. It is a review panel rather than an advisory board and our job is slightly different from that. We are there to criticise or to praise, whichever is appropriate as it comes. I would agree first with what Nicola and Sobia have said. For example, although DeepMind is very much a machine-learning AI company, none of what it has in clinical use is using machine-learning AI. Streams for example, which is the product that is being used, is not an AI process; it is algorithmic. A lot of what is described popularly as being AI is algorithms.

The Chairman: Would you like to unpack that slightly? I am not sure that we have had that distinction made before.

Dr Julian Huppert: For example, Streams is looking at acute kidney injury (AKI) and, essentially, it uses the NHS's standard protocol, which is to compare the measurement of one blood component to a previous measure of the blood component, and, if it is higher than a certain amount, it predicts AKI and, if it is lower, it does not. Comparing one number to another number is not really AI; it is a machine that does the comparison rather than a human. I would not think of that as machine learning or AI. That is following a simple algorithm.

The Chairman: There is nothing autonomous about it, so to speak.

Dr Julian Huppert: No, there is nothing autonomous about it. It tells a clinician who looks at a device, “Is this number bigger than that, according to the standard protocol?” There are a lot of buzzwords. There is a huge amount of hype around the whole AI area, as I am sure you have discovered. There is nothing like as much as there could be and will be. I agree that things such as image analysis are where the breakthroughs will be. DeepMind Health has a project with Moorfields Eye Hospital looking at retinal scans. That was announced last year. There was a published protocol for how the studies would be done. The results are promising—I have seen them—but they are currently under peer review for publication, which is the right way to analyse whether they are as good as they suggest. Diagnosis is very good and that could be a massive timesaver for skilled clinicians, who currently spend a lot of time analysing these images.

One thing which has not been mentioned yet is that one use that AI could have is trying to break disciplinary boundaries within medicine. We have a system of hyper-specialisation, with extremely good clinicians who generally specialise in one system, one organ or one disease. Sometimes there will be comorbidities or drugs will influence other systems, and the clinicians are not experts in every other system because they could not possibly be, so one could well imagine that a role for AI would be picking up some of those interactions in a way that humans cannot because they cannot study every single system.

Lord Swinfen: I have a supplementary for Nicola Perrin. How do the patients react to the virtual assistant and does it vary with the age of the patient?

Nicola Perrin: As far as I am aware, and it has only been trialled recently, they have been very supportive of it. They have the choice of either a chat line with a human or 24 hours a day using a virtual assistant. It is giving them tailored advice around arthritis and exercise, which is a particularly relevant issue to the patient population. Arthritis Research UK did a significant amount of trialling of it to make sure that it worked well for patients and that it was giving them the information that they wanted, in an interface they wanted. I do not know if they have yet analysed different demographics but, clearly, they are trying to ensure that the information they provide is accessible in different ways to different populations.

Lord Giddens: Could I add to this? It seems to me there are two uses of AI in medicine. One is in the content of medical research and genomics and genetics, et cetera; the other is organisational. The issue is how you put those two together and the questions that then arise from trying to do so, because you have two very complex datasets there and somehow you have to find a way of not losing yourself in them. Those two aspects of AI in medicine are not the same thing.

Dr Julian Huppert: There are a number of aspects. You talked about research and there is some really interesting work being done there. Diagnosis is the other one, which is not research but is not quite organisational either.

Lord Giddens: That would be somewhere in the middle.

Dr Julian Huppert: Indeed. The organisational point is a really good one. In our first annual report we looked at how, if you have systems such as Streams or anything else, that will change the way the NHS operates. The NHS is not configured to deal with much of this. It will change the way that doctors interact and it will change the accountability. For example, if you have a digital log of every time a doctor did something, that makes a big difference in terms of accountability. If there are, let us say, court cases on whether a doctor acted according to the best protocols, having recorded to the second when they did what will change the nature of that. It will change the productivity and have some really large organisational consequences for the NHS, which still need to be thought through.

Lord Giddens: It will change in two directions because there are two areas of digitalisation, which I find myself struggling to put together, and I think if I was running a hospital I would find the same, because you have a surge of advance in both of them, and you have DeepMind's experiments with blockchain trying to try to find a bridge between them.

Dr Julian Huppert: There are a number of things. DeepMind is using a blockchain-like, but not actually a blockchain, technique for data audit. The idea of that is to make sure that any time somebody accesses some data using the Streams app, or anything else I guess, there is a permanent undeletable record made of it so that you can always go back and have a look at who accessed what data and say, "I understand why that person accessed my data and why that person did; why did that person look at my data at that time?" That is a slightly different issue, but I agree that there is a problem with digital capacity within the NHS. I am a lay member of our Clinical Commissioning Group in Cambridgeshire and Peterborough. I am not speaking for the NHS or the CCG, but it is quite clear, looking across the whole NHS, that there is a lot more work required to e-enable all the hospitals. I think it is still true that Cambridge University Hospitals—Addenbrooke's—is the only fully digital hospital at the moment. It had huge problems implementing that but it is a lot better now, and there is a lot further to go across the whole country.

The Chairman: Would either of you like to comment on the supplementaries you have heard so far, because I am just about to bring in Lord Levene? Do either of you disagree with what Julian had to say in response to Lord Giddens?

Nicola Perrin: I would agree. The other thing to say is that there need to be very clear governance processes, and the governance will look slightly different for research purposes and for individual care, so that may help navigate the landscape slightly.

The Chairman: GDPR willing, or whatever it is. Lord Levene.

Lord Levene of Portsoken: I want to know whether AI is now used in the interpretation of MRI scans. If it is, is that interpretation regarded as the last word or is there any human review of it to make sure that it has got it right?

Dr Sobia Raza: This may be a question for the second panel because I believe there is a radiologist on it. As far as I understand, there is some

confusion as to the extent to which it is being used within medical imaging and for MRI scans. I believe there is a certain pathway that radiologists use where the image would be analysed by two individuals, and perhaps there is a role for AI that would alleviate some of that burden. It is not a particular specialist area of mine but could be one to explore with Dr Hugh Harvey.

Dr Julian Huppert: I am not aware of it being used. If it were I would be very alarmed if there was some sense that the results were automatically correct. I would very much expect a clinician to have a look.

Lord Levene of Portsoken: It would take some of the donkey work out of it.

Dr Julian Huppert: Exactly. A clinician would never look at the raw data from an MRI scan, the actual bits that come out; a certain amount is always processed ahead of time. I would expect it to look like that. From seeing the Moorfields' results that were shown by DeepMind Health, my recollection is that there were three machine-learning tools which had been trained slightly differently and which therefore would make three different suggestions, with a confidence estimate. They would say, "Having looked at this, we are pretty sure it is that, but it might be that or that". The idea is that you could show that to the clinician, who could spend most of their time looking at those things and, where it is not sure, that is where they will spend their time, rather than spending a huge amount of time doing, as you say, the donkey work.

Q118 **Lord Swinfen:** In your experience, how do the public view the use of artificial intelligence in healthcare? Are they indeed aware of its use? How does one improve the public perception of its use? If and when a medical AI application goes wrong and, for example, makes a decision or provides advice that adversely affects the patient, how should liability and possibly compensation be handled? Do we need a new mechanism completely?

The Chairman: Do you want to start with that, Nicola, in view of your job title?

Nicola Perrin: There has been very little formal analysis of what people think and how they feel about the use of artificial intelligence in healthcare, but the one thing we know is that there is very low awareness. People do not realise where algorithms are being used at the moment. Understanding Patient Data is working with the Academy of Medical Sciences to undertake some work to unpack public attitudes in more detail and to look at attitudes of the public, patients and, importantly, healthcare professionals, so hopefully we can begin to fill that gap a bit.

In the meantime, one piece of work by the Royal Society looked at people's attitudes to machine learning. It took a number of different areas and case studies—crime, education and driverless cars, as well as health—and it found that the views really depended on the purpose for which the data was being used. People seem to make a case-by-case assessment according to the context. They were particularly supportive of the opportunities in healthcare. They really seemed to get that machine

learning could be used in a positive way to support clinicians. Their concern was that they did not want to lose the human interaction. That doctor/patient relationship was really important to everyone.

The slight caveat with that study, in answer to your question, is that that was looking at machine learning, not AI. That was the term that it was using. The phrase “artificial intelligence” may be received in very different ways by people. There is a real risk that it is a sensationalist word that could lead to misperceptions. People are concerned about the use of patient data anyway. They are concerned about companies being involved in using data. If you put that together with this hyped-up artificial intelligence, I think there is a real risk of a very confused landscape. It is important to be able to explain clearly to people how data is used and why and what machine learning and AI is about.

The Chairman: What about the redress issue that Lord Swinfen raised?

Nicola Perrin: That is a much more difficult one to answer. It will come into some of the discussion later about accountability and regulation. Fitting it into the system in relation to liability, is the accountability to the person who has developed the algorithm, the person who has put the data into it or the clinician who is using the results of it? It comes back to your previous question as to whether an algorithm gave the decision on an MRI scan on its own. At the moment there will be some kind of clinical involvement and that will be part of the discussions around liability. It comes down to transparency as well. The difficulty with the black box element of an algorithm is not being clear how a decision has been made. The transparency and the explanation of how a decision has been made are going to be crucial.

The Chairman: You are not saying there should be a different regime for health rather than other forms of liability in that sense.

Nicola Perrin: I do not think so—yet—but there needs to be a clear regulatory framework and it needs to be very clear who has oversight and who is accountable.

Lord Swinfen: You said there was clinical involvement. Do you envisage a time when there may need to be no clinical involvement or there is no clinical involvement?

Nicola Perrin: I think it depends on what one is talking about. A very simple analysis of an X-ray could be done with much less clinical involvement, and routine ordering of diagnostic tests could be done without any, but I think it is 80 per cent to 20 per cent. The issues where there is much more uncertainty are where you really need the clinical involvement and specialist expertise.

Dr Sobia Raza: With reference to the public view of AI within healthcare, we have not done any specific work to examine this at the PHG Foundation, but I would reflect that AI developments within healthcare are largely contingent on using large patient datasets. We know already through work that has been commissioned through the Wellcome Trust²

² Note by witness: This refers to the report: Ipsos Mori for the Wellcome Trust, [The One-Way Mirror: Public attitudes to commercial access to health data](#) (March 2016)

that patients and people are generally unaware how patient data is used within the NHS, let alone by researchers and commercial organisations. The discussions around AI should include a broader conversation about how health data is vital to these uses in terms of improving healthcare and services. It ought to be as much a conversation about patient health data and its uses as it is about AI per se as a technology.

The Chairman: When you talk about a broad conversation, where and how and so on?

Dr Sobia Raza: Some of this work is currently being undertaken by Understanding Patient Data and the National Data Guardian has a role to play within this as well. I guess the point I am trying to make is that AI is one of the many useful technological applications that can be done to patient data. There are lots of simple things we could be doing better right now if there was improved use of patient data.

Dr Julian Huppert: There is little real public understanding of what is going on. It is hard for the public to separate out the hype when some things are over exaggerated and some things are underplayed. It would be nice if the public had a better appreciation of the benefits and the risks, because there are both. Future Advocacy, which gave evidence to you, did an interesting opinion poll which found that most people who had an opinion wanted AI used to diagnose diseases but most did not want AI to take on tasks usually performed by doctors or nurses. There is an interesting tension which I think goes to what Nicola was saying that the purpose is quite important. There needs to be a lot of effort around public understanding so that people can make more informed decisions. Our report recommended that that should happen and DeepMind Health should do some of that. Obviously, it cannot be the one telling people all about it because, quite rightly, it would not be seen as impartial, but it would be really good to see that level of education happening.

To pick up issues about responsibility and clinical involvement, I am not a lawyer, but my understanding of GDPR is that there would be real problems with making decisions about people in a fully automated consistent way. That will change in some special circumstances here, but I think it is quite relevant that there should be a human in the loop at some level.

In terms of responsibility when things go wrong, I do not think there is a need for a fundamentally different approach to what there is already with other medical devices. If you have a device which measures something in a patient, scans them, whatever it may be, if something goes wrong, it depends whether the system was designed badly or whether the clinician misread it. There is a whole series of things. I am not an expert in the area, but I think the same principles that currently apply could apply here. It will depend. If you have a machine that is giving unhelpful results, there should be a liability on the people who supplied that in a hospital, or wherever else it may be. It is one reason why I would be very keen in a Bayesian approach, I guess, of having outputs for these things saying, "This is what we think and this is how sure we are". I would much rather an AI said, "I am pretty sure it's this" or, "I'm not really sure but I think it might be that". It would be very helpful if there could be some

sense of when the device is not sure of the answer. Knowing that is really valuable.

Viscount Ridley: Can I push Dr Raza a little further on one point in particular? Given the background of the PHG Foundation, you have obviously done more on genomics than most of the people we have spoken to, and that is a case where you made the point that this is about health data rather than AI specifically. What lessons can be learned from the exploitation of genomic data in terms of public perception—for example 23AndMe and those kinds of things—about the mistakes that were made and the benefits that were garnered? If that is not your area, do not worry.

Dr Sobia Raza: With genomics there has often been a big focus on the technology and perhaps that is occurring with AI as well. We often need to consider what the uses and downstream applications are and what the potential consequences of that could be, rather than getting carried away with the technology and the nature of the data. The danger of exceptionalism is often spoken about in genomics and it could be the case with AI as well.

Viscount Ridley: What do you mean by “exceptionalism”?

Dr Sobia Raza: Potentially treating genomic data quite differently from other types of data because it might be seen as more sensitive. Increasingly, it is about thinking about the context and the use of the technology rather than just the very specific processes of how the technology works.

The Chairman: Why is the argument made for exceptionalism? It is not obvious, except to those who are the experts, as to what the data means. Why is the case made for it to be exceptional?

Dr Sobia Raza: Because genomic data is seen as inherently very personal to the individual from whom it comes. Increasingly, with the datasets that might be used for AI, we may require richer and richer datasets that link lots of different information together—biological markers as well as other demographic information—and, in and of itself, that can also be quite personal to the individual from whom it comes.

Q119 **Lord Giddens:** Should all publicly generated health data be made publicly available—if anonymous—to encourage progress in AI research and innovation? How could this best be done?

Dr Julian Huppert: My short answer would be no. The word “anonymous” that crept in there is a challenge. Anonymisation is not a magic spell where once you have made data anonymous, it stays permanently anonymous. Particularly with complex datasets, it is very hard to ensure that the data will remain anonymous. I am not sure that is very meaningful in the case of individuals. Re-identification is, in general, much easier than people realise. Lots of academic projects have taken apparently anonymous datasets and worked out how to re-identify. It is worth saying that the proposal in the Data Protection Bill would make it illegal to re-identify people. However, it is not clear that that is all that helpful. It will cause some serious problems in the research community. Also, people who wish to do that with malicious intent will continue to do

so regardless of whether it is illegal. We cannot, of course, do very much with people re-identifying overseas. Legislation does not work here. The analogy I draw on that point is that saying it is illegal to walk in somebody's back door and steal all their things is certainly useful, but encouragement to lock back doors is probably a more effective way of preventing it. We should try to make useful datasets available, but the idea that all publicly generated health data could be made publicly available just would not work.

Nicola Perrin: I would completely agree with that. You are absolutely right that access to data is key. To develop good algorithms you need as comprehensive a dataset as possible and it needs to be as high quality as possible, but I do not think the answer is to have a publicly available database to do that, for all the reasons that Julian has just set out. If we look at access to other health data, there are a lot of very good examples of mechanisms for controlled access, where you have an independent review process, so somebody is checking the purpose that the data is being used for and who the user is. You can have a data-sharing contract that sets out very clear controls over how the data can be used. NHS Digital has that already with the IGARD committee. There are other examples with clinical trial data. There is an independent review panel for CSDR—the ClinicalStudyDataRequest platform. Biobank and a number of cohort studies also have data-access mechanisms. It would work much better to develop a workable controlled access mechanism rather than a publicly available dataset.

Dr Sobia Raza: I would add that anonymisation diminishes the utility for some developments. It is not always a possibility. In addition to that, it is often quite challenging to robustly anonymise certain datasets. If we are thinking about health datasets here, it could be quite a small pool of data that is made available publicly, as robustly anonymised. In terms of machine learning and AI, ideally these algorithms work on large datasets, so, again, it calls into question the utility of a potentially very small dataset that is publicly available.

The Chairman: What about data trusts as repositories that would be not publicly available but be a halfway house? They were recommended by Dame Wendy Hall in her review. Would that be useful to have in the health field?

Dr Sobia Raza: In principle, I would concur with the proposals that were put forward by Dame Wendy Hall. It is just how we would see those working in the context of healthcare. There are also streams where some of this occurs now, so it is a question of avoiding duplication and asking whether there are already organisations which could undertake some of these roles.

Nicola Perrin: NHS Digital obviously has a data repository but there are also proposals in the Life Sciences Industrial Strategy for digital innovation hubs which may be able to act in the same way. There are also other NHS proposals for local and regional hubs to integrate care records. There may be opportunities to develop those into the data trust approach provided that there is clear governance and controlled access.

The Chairman: It is under active discussion and you think that there are

benefits if one takes that forward.

Nicola Perrin: Yes.

Dr Julian Huppert: I would agree with all that. One interesting approach which applies in some areas, particularly to the development of projects, and which DeepMind has also been doing, is to use synthetic data. To test whether something works and to develop it, you can generate data which looks like it would come from a human but in fact does not. You have to make sure you do not contaminate everything you are doing from it. For some purposes, if you test on synthetic data where a clinician has said, “This is roughly what a hospital output would look like”, for example, you can do what you like with it because there is nobody involved in the first place.

Q120 **Baroness Greender:** I want to ask you about the value of data as a commodity almost, particularly in the context of the NHS. Should the NHS be in some way recompensed or incentivised when it makes data available to companies for the purpose of AI development? This has come up in several submissions that we have had, and it is the issue of an algorithm being developed and built on the basis of data from the NHS and ending up in the west coast of the US in terms of any value. Given that we have all spent so much taxpayers’ money—and our parents and grandparents—to build up the NHS, should that value not come to us, and how?

Nicola Perrin: It is a really important question and a crucial one to get right because of the implications for public confidence. We know that what people worry most about in relation to access to data is when companies access it. They do not like the idea of the NHS selling data, but they are even more concerned if companies are making a profit at the expense of both the NHS and patients. We have to get much better at working out how the NHS can realise the value. There is not yet clarity over what that business model might look like. One of the reasons people are so suspicious of Google and DeepMind Health is because they have not made clear what their business model is, and that has exacerbated the difficulty in that conversation.

There are already approaches looking at how you reimburse for the price of a drug where a company has conducted a clinical trial with patients in the NHS, so there will be some parallels. I am not sure we have got it exactly right in relation to reimbursement for drugs. It is even worse for data because of the ongoing value of the data and the potential of taking the data and continuing to do more and more with it.

Should the NHS get access to a service at a reduced rate if patient data has been used? Probably yes, but what are the implications at a local trust level versus the NHS as a whole? There is a huge amount to be thought through. We were very pleased to see the Life Sciences Industrial Strategy flag it as an issue and recommend a working group to explore it in more detail. It is absolutely essential to get the benefit for both patients and the NHS if people are going to have confidence.

Dr Sobia Raza: I would concur with Nicola, but it is also worth reflecting that while the NHS has the essential ingredient here—the data—it does

not have the compute and the machine-learning expertise. Essentially, the development of these algorithms is going to have to be a collaborative effort. In order to build these processes, it is important to develop the appropriate frameworks to support cross-sector data sharing and ensure that there are sufficient incentives on both sides, benefits for all and fairness in how those benefits are distributed. I would again reflect on Sir John Bell's Life Sciences Industrial Strategy report and say that it is not clear how the benefits of existing data-sharing agreements, which are made locally between an NHS trust and a company, can be distributed more widely and shared across other trusts across the NHS.

Dr Julian Huppert: This is a really crucial question. Absolutely, the NHS should get some sort of return, but getting it right is quite tricky to do. We have already identified through our second annual report the DeepMind business model as being one of the things we want to particularly concentrate on, both in terms of the relationship with the NHS and how it interacts with SMEs. There is a real issue in this area of AI, not just in healthcare, about whether we will see a small number of over-powerful organisations or a large number of smaller organisations. One thing we have been pressing is to make sure that DeepMind Health does not force other companies out but tries to welcome them in. We will see where that goes.

How one structures it is quite tricky. It is worth saying—there has been some misconception of this in the press—that Streams is not a machine-learning tool, it just applies the NHS algorithm, so in that case there was no use of NHS data to train anything. It is not quite the same question that happened there, but it will come up, I would expect, fairly shortly. Free access is sometimes part of the exchange. For example, Moorfields is getting some help with its data storage systems, and it gets to control its own data and who gets to see and who does not, and that is part of its return. With Streams, there is benefit to DeepMind Health from the NHS mostly in terms of feedback from doctors and nurses on how a system works, and that is sort of recompensed by five years' free access to the tools. It is not easy to get it right. There is a real challenge as to whether we want them—

The Chairman: I do not know whether you know the basis, but I think in the arrangements between DeepMind and the Royal Free, a figure of £1.6 million was put on the value of those patient records.

Dr Julian Huppert: That was the number of records, not the value.

The Chairman: Sorry, yes, so how was the value of those records attributed?

Dr Julian Huppert: I am not sure there was any formal valuation of those records. In that context, DeepMind Health was acting as a data processor rather than a data controller. It was running that data through a system. The ICO has been quite clear about that and our independent legal advice said much the same thing. It was not using it to train any sort of AI. Can I come back to what I said? If we make it very hard to apply AI to healthcare here, companies will simply train any AI they use in another country around the world which may have slightly laxer rules.

Baroness Greender: Let me come back on that very quickly because the

NHS is almost unique. What you are describing are some piecemeal arrangements that are happening now, but I guess the question we are asking is: are we short selling what is very significant, unique-to-the-UK asset?

Dr Julian Huppert: At the moment there is not that transfer in AI particularly and we are not yet seeing very heavy use of training from that. I absolutely agree with you that we should make sure that there is a return to the NHS, but the question is what the best format is. People can develop and learn from these algorithms based on datasets in the US and other places around the world. I know there are places that are trying to do that. There is an interesting question about the best way of getting the most return to the NHS which does not make it so prohibitive that people just do all the learning in the US, or wherever else it may be, and then come along and say, "The bill is now this much". That would be the worst-case scenario from my perspective.

The Chairman: I am going to bring in Lord Holmes and then Viscount Ridley, Lord Hollick and the Lord Bishop.

Lord Holmes of Richmond: Is there any such thing as NHS data?

Dr Julian Huppert: The public tend to believe that the NHS is one institute which has all the data in one place. I think your question is well put because it is absolutely not like that. There are real problems with data storage, availability and flow throughout the NHS at pretty much every level. It is very much in silos at the moment. That may change and it is a question of how you change it to make it better rather than worse.

Viscount Ridley: I have a similar question. Your annual report said, "The digital revolution has largely bypassed the NHS which, in 2017, still retains the dubious title of being the world's largest purchaser of fax machines". I was stunned to read that sentence. Are we being a bit unrealistic about the value of the NHS's data?

Dr Julian Huppert: First, I should say that I have been unable to find a primary source for that figure. It was quoted by the Secretary of State for Health so I am sure it is entirely accurate, like all his pronouncements, but I would like to have found the primary source. I remember when I was in my CCG role and we were upgrading to NHSmail 2, we were told what the benefits were and it was said, "If you need to send confidential information, you can do it by using an email rather than a fax". Thus only last year or so, the benefits of email over fax were still being made out. Yes, there is a huge amount of work that is still needed to make the NHS more digitally savvy. It is definitely the case that NHS data would be far more valuable if it was in a better system. There is a long history to some of the issues with DeepMind and the Royal Free. When the conversation with DeepMind and the Royal Free started, the expectation of DeepMind was that NHS data was in a much more sensible pattern. The conversation—I am hypothesising here—was, "Great, we can learn some things from this; what format is the data in?" "It's a pile of paper over there". There is a mismatch there. We would get much more value from NHS data if it was in a secure but properly electronic modern digital form.

The Chairman: I am going to bring in Lord Hollick, but if you want to add anything, by all means do.

Lord Hollick: The use of public data is clearly a very valuable asset to DeepMind and its confrères. Dr Huppert, you are in a rather unique position because you have looked inside the beast, you know what makes it tick and you are on the journey to understand its business model. Could I ask you, because this is a very important question for this Committee, if you were advising the Royal Free now, knowing what you know, what would you say were the principal ingredients that it should focus on to ensure that proper value is received by the public realm—in this case the NHS—for the very valuable data that it is providing?

Dr Julian Huppert: Again, the data that was provided in that case was not used for any sort of AI. It is an important and slightly different conversation from the value in training algorithms on data. There certainly is some value in terms of testing Streams. In exchange for that, and one could argue about whether that is the right exchange rate, the Royal Free is getting five years' free use of the system. That is some sort of return. According to the ICO, our independent legal opinion and so forth, DeepMind Health is definitely not acting as the controller of that data; it is a processor, from a legal perspective. There were serious problems with that deal. I would not dissent from the ICO's commentary about how patients were informed about what was happening and the assumption that direct patient care was being set up, and I think that was in error. It is worth noting that even before the ICO report, DeepMind Health and the Royal Free had produced a revised contract which was definitely better than the original one. The new contract with Taunton and Somerset NHS Foundation Trust for Yeovil District Hospital is better again. I hope that will continue to be the case.

One of the big learnings for DeepMind was that it had assumed that if the Royal Free as an NHS body said, "It is all fine, we can transfer all of this data over and it will be under direct patient care," that was a correct interpretation. I know that DeepMind Health has now brought in its own in-house NHS information governance people so that it can develop its own views on how that would work. There were definitely errors.

Did the Royal Free get good value by providing some advice and feedback to have five years' free use of this? That is a question you would probably have to ask the Royal Free. The feedback I have heard from clinicians there is that they are very pleased with it.

The Chairman: I am going to bring in the Lord Bishop, but, again, if you want to add anything to what Dr Huppert had to say, please speak now.

The Lord Bishop of Oxford: I want to check out what I am hearing because it seems to me you said something very important a few minutes ago in this area. You seemed to be describing an emerging market in which NHS trusts are separately, more or less entrepreneurially, making different reward arrangements with different companies to use datasets that are very variable in their worth, suitability and application, which I have to say is a much more chaotic picture than I was assuming was the case, and not one that commands public confidence, I would suggest. I would like you to confirm whether I am correct in that assumption.

The Chairman: Can I just see if the other members of the panel have a view on this and then we will come back to you, Julian.

Nicola Perrin: That is absolutely the situation and it is an extremely good summary of it. What happens at a local level is very different between different trusts. They all have different systems. Some hospitals have hundreds of different databases which do not talk to each other. Absolutely, arrangements are happening at a local level and there are different conversations with companies and individual NHS trusts. From a company perspective, it is very difficult for them to know how to access the NHS which is a big beast and some hospitals have much easier conversations than others. There is a piece about how industry and the NHS can work together more effectively. There is, though, a national system as well. There are local-level data-sharing agreements and there is also NHS Digital, which is the national repository of different NHS datasets and where there is some linkage across different datasets. If one is talking about NHS data, perhaps one is talking about data in NHS Digital. There the conversation about value is equally important because NHS Digital has to operate on a cost-recovery basis; it cannot sell data. If you are concerned about the value being given away, it is positively set up in a way that means that, in order to ensure public confidence, it cannot sell data; it is cost recovery only. We need to look at what the model looks like at a national and local level.

Dr Sobia Raza: I would concur with that. There is a real opportunity here to realise the true value of NHS data at a national level. In addition to that, it is worth remembering that when these agreements are made at a local level, we are talking about local pockets and datasets, and a population, say, within London is quite different from a population within a rural area. From a technical perspective, a huge opportunity arises when you can capture the differences in demographics and, essentially, collate a more enriched dataset which is more reflective of the wider population for which you want to develop tools and which it could serve. There are different dimensions to this: realising the benefits in terms of negotiations with companies and developing a dataset that could provide more opportunities for accurate tools and algorithms.

The Chairman: Do you have a final word before I move on to the next question?

Dr Julian Huppert: Your characterisation of the state of NHS data is worryingly accurate. There are lots of different providers in lots of different trusts and the system is very chaotic. There has not been very much work to look at some of the providers whose standards are not very high. At DeepMind Health we hold them to a very high standard. There is a role for people to look at all the providers of all data and all the systems within the NHS because there are some very real problems.

Q121 **Viscount Ridley:** You have already briefly touched on unrepresentative data in answer to a couple of questions. We read the report about how IBM Watson was trained on Memorial Sloan Kettering data only and is not applicable to the rest of the world. Is this a problem and how do we deal with it? What are the ways around it?

Dr Sobia Raza: When it comes to medical devices and tools there are already principles and approaches in place which ought to be applied. It

comes down to assessing the validity of the tool—knowing that it does what it claims to on the tin and its clinical utility.

Viscount Ridley: That is already a problem addressed in normal medical technologies.

Dr Sobia Raza: There are assessment processes and pipelines in place for those. If a tool has been developed and trained on a certain population, you want to know that it is going to still be effective in the population that you want to apply it to. It is really understanding and having access to the evidence base in terms of how that tool has been validated. When it comes to AI algorithms, there are two important aspects we would want to know: the training data that has been used to develop the algorithms and how representative that training data is of the population within which you want to use that tool. We need some simple information as to whether it has been trained on data from 20 people or 20,000 people, and, increasingly, we need to know how it [the tool] has been tested, so, again, whether it has been tested within the population on which you want to use it. An analogous comparison would be autonomous vehicles. They might be trained and tested in California where they work perfectly, but if you want to use them on windy country roads within Britain, you want to have the evidence to show that they can also perform there. It is not much different.

Viscount Ridley: Presumably, there is an ethnic aspect to this too.

Dr Sobia Raza: Yes, and the other point is that, as a country as we seek to use our datasets to develop AI algorithms, we need to account for the fact that we are diverse in terms of population so, again, in order to mitigate against any disparities, we want to work with datasets that are sufficiently representative.

Nicola Perrin: Again, some of these issues are not completely new or unique. There are similar kinds of issues with clinical trial populations. Most drugs are tested on white, middle-aged males who meet very strict eligibility criteria for a trial and are not on any other drugs. They will be rolled out in the real-world population to females, young people and elderly people on multiple drugs. We have got used to testing and to translating from a test environment to the real world with drugs. It is exacerbated, though, with algorithms because of that black box element and the fact that the algorithm itself incorporates any biases that are in the dataset in a way that is very difficult to account for, or to explain, because it is a black box. It is exacerbated with AI, but it is not necessarily a completely new issue.

Viscount Ridley: So transparency is a big part of the answer.

Dr Sobia Raza: Absolutely.

The Chairman: Do you agree with that, Julian?

Dr Julian Huppert: I would agree with that. Explainability of AI is a big problem. As Nicola said, some of these challenges are not new in healthcare. There are very serious problems across the whole of AI with unrepresentative data propagating prejudices and biases that were there already. I would add very briefly that one advantage of getting some of

the systems trained on NHS data rather than overseas data is they are much more likely to be representative of NHS populations.

Q122 Lord Levene of Portsoken: Does the NHS have the capacity to take advantage of the opportunities represented by AI technology and to minimise the risks? Are the clinicians and other healthcare professionals equipped with the necessary skills to take advantage of AI technology in their practice? What could be done to help them?

Dr Julian Huppert: The short answer is no, the NHS does not have the technical capacity to make the most of it and to minimise the risks. There is a real risk that we are just not in a fit state in the NHS to do the best on that. Clinicians vary: some of them are very technologically savvy and very keen and eager and some of them very much are not. In some areas, I am sure this could be seen as a threat to employment, and if there are systems which can do some tasks which people are currently doing, they will be reluctant to implement them. This is gradually changing and the NHS is realising the need to be current. However, there are challenges with that. We are seeing more and more people working in the NHS who are used to having everything on a phone with an app. They expect a simple, clear interface in the rest of their daily lives and they want similar things. We are seeing more use of bring-your-own-device—BYOD—in hospitals. That brings its own challenges, particularly around security, but we will see a transition there. For now we need to do much more to help the NHS modernise. For example, WannaCry brought out just how archaic some of the systems are.

The Chairman: BYOD is a new one on us. It is like pick your own, is it not, or whatever the acronym is?

Dr Julian Huppert: Yes.

Dr Sobia Raza: I agree that there is an important need here for healthcare professionals to have knowledge about the technology, to be aware of what it is capable of and to understand its limitations and gauge an awareness of how it might change or influence clinical practice in years to come. At the same time, it is equally important for early engagement between those developing the machine-learning tools, the machine-learning practitioners, and healthcare professionals in order to prioritise and identify those clinical and research questions to which machine learning would generate most utility right now. This is as much about a multidisciplinary collaborative piece as it is just about healthcare professionals and their awareness around AI.

Lord Levene of Portsoken: If I am a healthcare clinician and I have heard about this and think it is wonderful, but I know very little, what is there to help me to learn? Where do I go and how do I find out?

Dr Sobia Raza: That is an important question and not one I am able to answer right now. Perhaps there is work that the Royal Colleges could do around this to raise awareness.

The Chairman: Are you more or less saying that the NHS is inherently conservative and therefore we have to look for other mechanisms?
Nicola.

Nicola Perrin: There is the new NHS Digital Academy, which has just been set up for CCIOs. It is very much in its infancy so it does not yet have any resources that would help answer the questions you are asking, but it is beginning to look at how to support CIOs and CCIOs within NHS trusts.

Going back to the question you were asking before about how we can help the public to have confidence, the public trust clinicians and their doctors, so we need to tackle healthcare professionals first. GPs and clinicians need to have confidence in how data is used. At the moment, they do not and they are very risk averse. They are confused and uncertain over the data protection regime, so their concern is to protect data rather than share it. That was exacerbated perhaps by the care.data fiasco. One of our starting points needs to be helping clinicians understand more about how and why data is used and how and why algorithms and AI might have potential benefits both for them and for their patients.

Lord Levene of Portsoken: If there is little or no resource for them to go and learn, how are they going to be able to do it?

Nicola Perrin: Health Education England has some e-learning resources. It needs to be part of CPD for clinicians.

Lord Levene of Portsoken: Presumably we are talking about a lot of money and the NHS is clearly lacking money. Is there anything even approaching what they need to do it properly?

Nicola Perrin: Not at the moment.

The Chairman: Somebody needs to take a leadership role presumably.

Nicola Perrin: Yes.

The Chairman: And there is a bit of a vacuum at the moment.

Nicola Perrin: Which is where the Digital Academy could come in.

The Chairman: Thank you. We will move swiftly on to the Lord Bishop.

Q123 **The Lord Bishop of Oxford:** Do we need new ethical standards or principles for the use of AI or are existing codes of ethics sufficient? If we need new standards or principles, what should they consist of?

The Chairman: This is in the healthcare field. Shall we start with Dr Raza for a change?

Dr Sobia Raza: I am probably not best placed to comment on this. My first point is that we need to be careful not to duplicate initiatives. I would be happy to reflect on this, consult with colleagues and provide a written response.

Dr Julian Huppert: New fundamental principles are not needed, as the core standard principles we have still apply, but there needs to be a lot more work on how they operate in practice because what were sometimes small edge cases are now core issues. Trying to have clarity on how you take the various different sets of rules and principles and put them together would be very helpful. The ICO report on the Royal Free, for example, was helpful in saying, "We would like to help develop clarity

in this". My centre at Jesus College is trying to do some work, together in fact with Nicola, to see if we can get some of that to work.

Transparency is going to be really important and there are lots of ways of achieving it. With DeepMind Health I talked about the verifiable data audit so that there is a clear log of everything that has happened for an auditor to check. There are points about interpretable AI. We have been pushing them very hard to publish all of their contracts with the NHS. They are up for doing that but the NHS is a little more reluctant. If people can see everything that is being done, it will help.

Can I say a bit about the role that we have as independent reviewers? It is a new model. I am not saying we have done it perfectly the first time. We are a year and a half through it and I wish I knew a year and a half ago what I know now about how to do it. The concept is quite interesting and very bold of having nine of us, including Matthew Taylor from the RSA and me, to review what they do. We are not under any confidentiality obligations. Remarkably for an Alphabet organisation, we are specifically not under any NDA. I am free to set up in competition with them using what I know, although I am not planning to. We have a budget to investigate them, so, for example, we commissioned independent legal advice about the contract with the Royal Free. We chose a lawyer and briefed him. I think they only knew who we had chosen when they got the bill. We commissioned a company to undertake a detailed security audit of all it was doing and we published everything that was found. I do not know of many technology companies in the country or the world that would be happy to have that done, particularly given that they got the bill for doing it. As it happened, it found fairly minor things, but if it had found major vulnerabilities, we would have said so. It is quite a bold approach. Some people in Alphabet are slightly nervous because there are nine of us essentially holding a gun to their heads. It also means that if we find anything problematic, we will be saying so.

Whether that is the exact model other organisations could use, I do not know, but it would be nice to see some sort of system like that for transparency so that, as soon as possible, concerns can be aired and looked at.

The Lord Bishop of Oxford: I appreciate that and I really appreciated your report—I think there are great strengths in the model—but everything you have said just now reinforces for me the immense value to the company of building public trust and the huge potential of the development of this technology for the future to make that investment worthwhile. However, that is a Catch-22 and I think the level of scrutiny is very good.

Dr Julian Huppert: Baroness O'Neill would of course say that the key thing is to demonstrate trustworthiness rather than just build up trust, and I think she is largely right on that. Part of our role, hopefully, is to demonstrate that they are trustworthy, if they are, and where they are not, and if they are not, to say that so there is not false trust.

The Chairman: I am going to bring in Lord Puttnam but first, Nicola, do you want to add anything?

Nicola Perrin: I do not have anything to add to what Julian has just said.

Q124 **Lord Puttnam:** In a sense, you have half answered the question, but I am taking advantage of you being here. As part of our day job, the Lord Chairman and I are struggling with the Data Protection Bill. One of the proposals we are advancing is a committee not unlike your own but a government committee. What advice could you give us that you have learned that might help us in trying to formulate legislation on data protection?

Dr Julian Huppert: I might need to think about that to give you a more informed answer.

Lord Puttnam: By all means drop us a note.

The Chairman: You are no less informed than we are.

Dr Julian Huppert: Perhaps I could send something separately later because then I could give you a more sensible, properly thought-through answer.

Lord Puttnam: That would be very helpful, thank you.

Dr Julian Huppert: I would be delighted to.

The Chairman: Do you have any thoughts on the Bill or GDPR? Do not worry if you do not; it is fairly crunchy. Lord Hollick.

Q125 **Lord Hollick:** One question you have already held to the head of Alphabet is to ask the rather intriguing question at what stage an organisation such as Alphabet becomes simply too powerful. Unfortunately, you do not give the answer. Let us move on. What are the criteria that you would select to make a judgment on whether or not they are too powerful?

Dr Julian Huppert: That is a fascinating question and indeed is one of the things we are looking at more for the second review. I am a liberal and one of my core concerns is the overconcentration of power in any form. That is one of the things that drives my political philosophy, I guess. I have a concern about whether there will be a handful of monolithic tech organisations. That worries me just as much as monolithic press organisations or anything else. I was quite keen to look at that. I do not necessarily have a clear answer, but I think there is a real risk that we will find monolithic organisations. Government has a role in challenging that and trying to make sure that there is diversity. For example, one thing that could be done is to make sure that, where appropriate, data is in relatively open standards. For example, if DeepMind Health produces a tool of some kind which is good at analysing retinal scanning, predicting acute kidney injury, or whatever it is, it should use a standard data interface system, such that if somebody else comes up with something better, they could easily put that in and displace DeepMind Health. I would like to see that. DeepMind Health would only get to do things if it is, in fact, the best at doing that. What we have seen for many decades, as I am sure you know, is in many cases tech companies getting lock-in because it is too hard to change the infrastructure. That is a very worrying system.

Lord Hollick: With a 70 per cent share of search, is Alphabet not already in that position?

Dr Julian Huppert: Indeed, and I am concerned generally, which is why I put that piece in the report. How one tackles the control of some of these very large companies is a real challenge. I do not want to tread too much into the current debate, but being able to work as part of a larger cluster of countries is perhaps a more effective way of standing up to companies such as Google. That may be a different conversation you do not want to wade into.

The Chairman: It was very neatly introduced, if I may say so.

Baroness Grender: We hear you.

Lord Hollick: That is for the Committee next door.

Dr Julian Huppert: There is a general question about how one stands up against very large organisations in the press world or tech world or any other world and whether one can take legal action and have it matter. If a tiny country such as Iceland were to say, “Google must do the following or else they cannot operate here”, ultimately, I suspect Google would say, “Well, we don’t operate in Iceland”. If the US or the entirety of the EU, or whatever, said, “You have to do this”, they cannot take that risk. It is a very interesting challenge. I do not have the magic answer to it.

Lord Hollick: Will you have it by next year’s foreword?

Dr Julian Huppert: I am not sure I have any great answer to it. I will continue to raise the concern. It is not just about Alphabet; there is also, for example, real concern about Facebook. One can imagine a future where Amazon—or Amazon and Alibaba, or whatever it might be—controls a very large fraction of the world’s shopping. That would also worry me because it leads to possible abuse. That applies if you have malicious intent from a large powerful organisation; then you have problems and, in the future, the people running them could be more malicious than the current ones. It also means there could be inadvertent error which causes great harm, even if there was no intention to do so.

The Chairman: I would love to pursue all this, but we have move on. Lord Holmes.

Q126 **Lord Holmes of Richmond:** What form should government intervention take in terms of policy, regulation or investment to enable the NHS and society to gain the benefits from AI?

Dr Sobia Raza: In my opinion, there are two areas which we have already touched on which ought to be the focus. The first is to not lose sight or momentum when it comes to NHS digital infrastructure in terms of having the appropriate frameworks for storing, collecting, transferring and processing data, because that is a large overhead when it comes to enabling data collation for AI. Secondly, there should be a continued focus on building the trust of patients, the public and healthcare professionals when it comes to sharing patient data.

Nicola Perrin: That is right. AI technology is advancing quickly and both the policy and the public conversation have to catch up fast. There is a

real risk that there could be a public backlash against AI and we need to make sure that does not happen, and the Government definitely have a role to play in ensuring that.

There are examples of how one can introduce a new technology and keep public confidence. If we look in particular at mitochondrial donation or stem cells, that is a really good example of where it was done well. It was done well because a number of things were brought together at the same time. There were the advances in the science and technology, a robust regulatory framework, a wide discussion of the ethical issues and the implications for society, full public consultation and, most importantly, there was benefit for the patient. I think we need to bring all those things together again in AI. With the regulatory framework, there needs to be sector-specific regulation, as we have been talking about. Regulation for driverless cars will not necessarily be applicable to the healthcare sector.

The Chairman: Are there some common principles?

Nicola Perrin: Common principles and learning from each other and oversight across the board, but sector-specific regulation.

The Chairman: Actual regulation is required.

Nicola Perrin: I think so. Ethical issues, as we have said, are not new or unique, but they are magnified with AI. Transparency is going to be essential and that is one of the key parts about the public dialogue and starting with examples of where there is benefit for the patient and for the clinician, because that is the best way of ensuring the public have confidence and it is showing them why it is important for them.

Dr Julian Huppert: I have three suggestions and one principle. The first suggestion would be to bring all the policies and regulations together so there is one workable system that fits with all the regulators, with the DH, the NHS, patients, companies and so forth, and there are no grey areas or conflicts in the rules and people know what they should and should not be doing. It should be quite simple, with quick responses when something has gone wrong. I would support, for example, the Royal Society/British Academy proposal to have a stewardship body. That could play a crucial role in that. The second is investment in NHS IT systems and knowledge throughout the system to be able to do these things competently. The third one, as Nicola said, is about public and patient engagement, so they know what this is and what it really means.

Regarding the principle I would put forward, essentially, there are questions about the benefits of doing AI in health and the benefits of, for example, protecting data and privacy. People often say there is a tension between these. At the risk of trying to draw a graph at a Select Committee, I see these as large orthogonal axes. There is a limiting position. There is a maximum you could do while protecting certain amounts of privacy, but we are nowhere near that limit; we are well inside it. There are options which would mean that patient data, for example, would be better protected than it is now because the system now is far from perfect and we get the benefit of learning and improving from it as well. We could have the conversation about the limits once we are a lot further along.

The Chairman: We have to move on rather swiftly to our last question.

Q127 **Lord Puttnam:** If there was only one recommendation you would like to see the Committee make at the end of this inquiry, what would it be?

Nicola Perrin: There needs to be much greater clarity in the governance landscape. A proliferation of different bodies has been proposed. Julian has just mentioned the Royal Society/British Academy proposal for a data stewardship body. DCMS is looking at a council of data ethics and the Nuffield Foundation is talking about a convention on data ethics. There are also regulators in the space already. There is the ICO in the health sector, the MHRA and the National Data Guardian, from whom you are about to hear. I am not entirely clear how all those different proposals fit together, or even if one needs all of them. If you could give clarity over what is needed rather than suggesting yet another new body, it would be very helpful. The new solution needs to include horizon-scanning—short, medium and long term—given how quickly everything is evolving. It needs to be able to convene cross-sector discussion and bring together all the different stakeholders, including industry, to build consensus and identify any regulatory gaps, the kind of oversight that you have mentioned, and it needs to be really clear what kind of authority and accountability it has and whether it has teeth.

The Chairman: Terrific. We are taking furious notes at this point. Dr Raza.

Dr Sobia Raza: Mine would be that we ensure we get the value and potential from NHS data for patient benefit, first, through a continued drive towards digitisation and embedding appropriate and modern digital infrastructure and, secondly, through developing perhaps an NHS-wide strategy around health datasets that are used for algorithm development to ensure they can be sufficiently representative of our diverse population.

Dr Julian Huppert: I would agree pretty much with both of those. To pick up a different one, I would hope that you would say more about how to invest in getting more understanding of technology, and what it can do and how it does it, and that should be everything, from schools, to journalists, to the general public, to Parliament.

The Chairman: Thank you all very much indeed. It has been a very interesting session, and challenging in some cases. We will be taking to heart quite a lot of what you have said. Thank you.

IBM, Sage and SAP – Oral evidence (QQ 76–84)

Evidence Session No. 9

Heard in Public

Questions 76–84

Tuesday 7 November 2017

[Watch the meeting](#)

Members present: Lord Clement-Jones (The Chairman); Baroness Bakewell; Lord Giddens; Baroness Greender; Lord Hollick; Lord Holmes of Richmond; Lord Levene of Portsoken; The Lord Bishop of Oxford; Lord Puttnam; Viscount Ridley; Baroness Rock; Lord St John of Bletso; Lord Swinfen.

Examination of witnesses

James Luke, Andrew de Rozairo, Kriti Sharma.

Q76 The Chairman: I welcome our three witnesses to our first session today. We intend to explore in particular the impact and use of artificial intelligence in business-to-business interaction. Welcome, James Luke, chief technology officer for the public sector and master inventor at IBM; Andrew de Rozairo, vice-president, customer innovation and enterprise platform, SAP; and Kriti Sharma, vice-president of artificial intelligence and bots, Sage. They are all very good job titles, if I may say so.

I have a little rubric that I always need to read out, at the risk of boring my colleagues. This session is open to the public. A webcast of the session goes out live and is subsequently accessible via the parliamentary website. A verbatim transcript will be taken of your evidence and will be put on the parliamentary website. A few days after this evidence session you will be sent a copy of the transcript to check for accuracy, and we would be grateful if you could advise us of any corrections as quickly as possible. If, after this session, you wish to clarify or amplify any points made during your evidence or have any additional points to make, you are welcome to submit supplementary written evidence to us. First, would you like to introduce yourselves for the record? Then we will ask the questions.

Andrew de Rozairo: I am vice-president of customer innovation at SAP. I run a team of experts on big data and machine learning, and I work with customers across different industries to use data-centric approaches to improve their competitive position. Our team works with people such as Vodafone, BP, Burberry, AstraZeneca and ARM on a range of different topics. I am very excited to be here today. It is great that we are focusing on the business-to-business side, because while there is a lot of hype on the consumer side about driverless cars, personalised medicine and all those things, there is a quiet surge of things happening out there

in the B2B space that are much more fundamental and are bringing real advances in productivity today, which is especially important for us in the UK.

Kriti Sharma: I am VP of artificial intelligence at Sage, which is a FTSE 100 company. We build software for anything from money to people for about 3 million small and medium-sized businesses. We are a little different to SAP; we are bringing that perspective. I lead the practice of AI across Sage with customers in 23 countries with a focus on driving productivity, similar to what Andrew said. These are small businesses which are wasting a lot of time doing mundane activities. A lot of the AI we do is not as sexy as robots and self-driving cars, but it adds value on a day-to-day basis. I also spend a lot of time in building the next-generation talent pipeline across the industry, also focusing on the ethics and the issues in diversity with AI. Since last year I have been building a technology community in London of about 1,500 people who are developers who are bringing AI technologies to share those practices, based on what we have learned.

James Luke: I am an engineer; IBM calls me a distinguished engineer. I am the CTO for our public sector business, and I have spent most of my career working in the public sector developing artificial intelligence solutions for government agencies and for commercial organisations that solve interesting problems. I emphasise that, because I am an engineer and I am about the practical application of the technology to real-world problems, and I hope that I can bring that to the Committee today. I have also worked in IBM's development organisation, where I have a number of patents for my work on text analytics. I worked on the Watson solution and capability, and I lead the IBM Academy of Technology's core team on cognitive computing, where we inform and influence IBM's technical strategy in that area.

Q77 **The Chairman:** Thank you very much indeed. I am going to open with a very general question. What, in your opinion, are the biggest advantages and disadvantages that AI could bring to UK businesses over the next 10 years? If you like, you can put that in the context of the Dame Wendy Hall review, focusing in particular on what your own businesses are doing in investment and where you think businesses need to be encouraged to take up AI systems.

Andrew de Rozairo: The biggest benefit we see is integrating artificial intelligence into business processes as they exist today. That is where we see the biggest benefits coming to businesses around the country.

On the challenges that we see or the areas where we would encourage more adoption, in the UK we have a fantastic set of resources, a pool of skills, for artificial intelligence—some of the best in the world. There are two areas where we are lagging behind. One is the number of people who can translate artificial intelligence algorithm-concept science into real business benefits. The biggest one is the data side. You need two parts to make artificial intelligence work: one is the algorithm and the skills in order to build those programs, but the other is the data, whether it is enterprise data or machine data such as the internet of things. From Eurostat's studies it looks as though we are really lagging behind there.

The Chairman: This is really interesting. I will bring in Kriti and James in a minute. Are you saying that your customers do not yet have the products they need in this area and that their own data is not adequate for enabling them to adopt AI solutions?

Andrew de Rozairo: Specifically, we see two challenges. Some customers—I would say the less mature customers—have not collected enough data to be able to do something interesting with it. A manufacturer will say, “We would like to be able to take a look at the productivity of this production line. What data have you been collecting?”, and they will say, “Ah, yes, we have not really started”. The more mature customers have data, but it is very siloed and they have never brought it together before. They have real struggles doing that. If they want to bring three sets of data together such as customer complaints, what was happening on the production line at the time, and the training that the people on the production line had, it is hugely challenging for them. They do not have this unified data platform.

The Chairman: Your customer base is larger companies than Sage. Is that your experience, Kriti, but even more pronounced?

Kriti Sharma: Our issue is a little different. Our small businesses are in two categories. One is the new generation cloud-based technology-driven companies, the small businesses that have truly embraced this technology. For them, the move to AI is encouraging and positive. They are the early adopters. In the UK we have some issues with digital adoption. Some 55 per cent of small businesses are still using pen and paper, Excel spreadsheets, fragmented datasets. When I talk to them about AI, they say, “Hang on, let me get my data in the cloud or somewhere online”.

The Chairman: They have not even gone digital yet.

Kriti Sharma: Exactly. That is where we see a challenge. We have a lot of early adopters and customers who are truly embracing AI. We have several of those since we launched our first iterative AI product, Pegg, which automates the mundane tasks, such as chasing invoices, setting up expense reports, tax returns, VAT filing and payroll. They are benefiting from it. For a lot of the others there is still a journey to go on, and we need to do a lot more work in educating the segment of small businesses to make sure they are not left behind in our drive to AI.

The other challenge is that there is a lot of talk about AI, but we need to make sure that we talk about automation. They are not the same things. You can automate processes without needing AI, without having to go through the worries of robots stealing jobs and taking over. We need to be very careful about it. The lack of digital adoption is leading to what we call the productivity gap in the UK. We did research and found businesses are spending 120 days a year doing admin; mundane paperwork tasks. That is massive. This equates to £34 billion in GDP that we could create or unlock for the economy. That is the biggest challenge for small businesses.

The Chairman: I will come back to what can be done to incentivise greater take-up in a minute. James, do you have a more sophisticated customer base?

James Luke: I would echo the comments of my colleagues here. I see a huge opportunity for AI across business. It is potentially massive. We see some very successful projects. The big challenge is if you take AI beyond improving what organisations already do. There is this huge efficiency gain if you take AI and apply it to automation. I would pick up on the point that you do not need AI to do a lot of automation. One of the features of AI through history is that a problem is seen as an AI problem until it is solved, and then it is seen as just ordinary computing. It displays intelligent behaviour.

However, we also have a challenge that AI is much more than that; it is an enabler for huge new business opportunities. I worked with an insurance company that wanted to look at insurance analytics, and they saw a massive new way of delivering life insurance underwriting. It was a huge potential business opportunity to them. The challenge is: will the AI work? We have to evaluate it. How do we evaluate it? We need the data to evaluate it. We do not have that data yet because this is a completely new business process.

How do you get over that hurdle of demonstrating the return on the investment before you get the data that justifies going out and getting the data? It comes down to skills. A key thing that we have to do is lead our customers through this adoption of emerging technology; teach them how to begin the journey. Often the destination you reach is not the one you originally set out for. That whole skills base of understanding the technology, how it applies to the business, how the business needs to evolve and how to run an emerging technology programme is critical to the successful adoption of AI.

The Chairman: That is very interesting. Andrew, you mentioned skills earlier, did you not? Do you accept that that is the key issue? It is not a matter of special incentives; it is a question of giving people the skills so that they then understand what can be done and needs to be done.

Andrew de Rozairo: They are skills, absolutely, and they are skills on three levels. The first is the skills to create the programs. Also, we talk about translators. How do you take these capabilities that are delivered by artificial intelligence and translate them into something that drives business benefit? The third area of skill is better education of the executives at the companies so that they have the leap of faith to say, “We believe that AI is the way forward. Even if we don’t have a rock-solid business case today, we need to make the investment. This is the way things are going”. We invest a lot in building digital prototypes—quick ways of showing customers, using this data, using this approach the value you can derive from it. It is a very quick approach.

The Chairman: Thank you. Could we have a word from you, Kriti, on how you get from pen and paper, as you described it, towards AI and greater productivity?

Kriti Sharma: I think it is an opportunity to leapfrog. Now you can have superpowers, as they would call it. For businesses that do not have a finance function or an HR function at all, or medium-sized businesses scaling up to enterprise level, during that growth phase, AI can help them get to that and set the basics right. Going back to the skills point,

absolutely you need people who can develop the AI technology, such as James or me, or engineers like us; people who can apply it to solve problems; and, thirdly, those who understand that there is a lot in it for them. The point we need to address is how we make people or businesses realise that there is a lot of value to be gained at a productivity level and in expansion or scaling.

Baroness Bakewell: The phrases that have cropped up are “struggles to do that”, “we need to educate” and “we need to teach”. Who is going to do the educating, the teaching? Do we have a cadre or cohort of people already equipped to do the teaching?

James Luke: It is down to industry and practitioners in the field. In IBM, we run education programmes with universities. All our apprentices who come through are sent on training in cognitive skills now and how to build chatbots. Every IBM-er is required to do 40 hours of education a year, and we have over 300 AI learning modules. As the cognitive lead for the IBM Academy of Technology, I am starting an initiative on explaining “cognitive” where I am looking at how we move people’s understanding of cognitive away from just artificial intelligence and deep learning. Everybody thinks it is algorithmic and it is not; it is about building solutions. It is about understanding how to define the problem, how to cleanse the data, how to understand and fault-test the system as you go through it and how to continually improve it. There is a whole new set of skills there. It is the sort of thing we are rolling out across our company but, also, into industry and into reachback. You can even go to an online academy and do work in this area.

Baroness Bakewell: It sounds as if your company is doing well. Are the universities doing enough, do you think?

Kriti Sharma: I finished university seven years ago. I studied computer science. The way AI or computing was done seven years ago is so different to how it is done today. The part that worries me the most is that in schools and the early days of university education we are still teaching something that was relevant a few years ago. By the time these kids graduate, the nature of AI will have become very different. Thanks to people from IBM-Watson, for instance, there are tools you can apply to solve problems. It is like the early days of websites, where you had to be a web developer to build a site, and now anyone can get started in an hour. That is the generation we need to prepare for this. We need, first, to build the core skills in AI. We are not doing enough to bring in a diverse skillset of people who are the problem-solvers who can take the tools that have been developed and apply them to solve issues that the customers of SAP or Sage have today.

The Chairman: I will bring in Lord Giddens and then we will come to you, Andrew.

Lord Giddens: Thank you for some very interesting opening comments, which I much enjoyed. I would like to take up the issue of automation. If you look at manufacturing in this country, over the past few years it has had one of the most amazing secular changes that has ever happened in western economies. About 45 per cent of the labour force used to work in manufacturing. Now, depending on how it is calculated, it is down to

about 12 per cent. That is quite extraordinary. Very little has come from Donald Trump's overseas competition; most has come from automation. It seems to me there were two phases of automation in most of these industries. In the first phase it was mostly mechanical but in today's phase it is pretty integrated. Certain levels would not be possible without digital technology. That is where the cutting-edge changes are. Since you have more or less got rid of manufacturing jobs, the only jobs left are white collar and service jobs, and we are now talking about the automating of a lot of those jobs. It would be interesting to hear you amplify your comments on these things.

Andrew de Rozairo: It is an important thing to think about. What we see happening now is that this AI approach is not about replacing human effort; it is about augmenting human effort. It is about refocusing people to spend their time on value-added activities. I will give a very simple example. Someone working in finance in a large company may have a monthly task of taking a look at all the payments that come in from their customers, and trying to map it to all the outstanding invoices that they have. It should be easy: one payment, one invoice. It does not usually happen that way. Usually, a payment is the first two lines of this invoice, 40 per cent of this invoice and these three things on that invoice. Today they have huge amounts of spreadsheets and they are spending their time trying to mix and match. This can be easily done by computers. You can provide a shortlist or a recommended or likely match between these two things, and then the finance people can spend their time on the exceptions or understanding how is it that we change the way we issue our invoices or get our payments in order to make it better. We see this phase of artificial intelligence as absolutely about augmenting the human efforts in businesses rather than replacing them.

Lord Levene of Portsoken: Andrew, you were saying earlier that companies have to look at the opportunities there, decide whether they are going to start making this investment and then take the plunge. I can understand that for large companies that are used to taking lots of such decisions; they can make investments and if some pay off and some do not it does not matter. How do you deal with the medium or even smaller companies that would probably benefit from the use of AI? How are they going to know how much this is going to cost them before they plunge in? Is there a way you can help them with that?

Andrew de Rozairo: We see two patterns. The first is running something called a digital prototype. How can we quickly take some data against a specific business outcome and prove to you in a relatively short period—we are talking four to eight weeks; low investment—what we can do with your data? That is one thing. The second thing—and SAP is very focused on this—is we are going to learn from some of the projects we do with large customers and take that artificial intelligence capability and embed it in the business processes. The other customers do not have to invest in the same way; it is already built in. If someone says, "Okay, I'm now posting a new position", the artificial intelligence system will say, based on similar positions and the CVs you short-listed last time, "Here's a shortlist of the CVs this time". They do not have to think. They do not even recognise that it is artificial intelligence helping them to shortlist

this; it is embedded. The huge productivity engine that artificial intelligence provides is when you are able to embed it.

The Chairman: We must move on. I want to ask James or Kriti whether they have anything they want to say in response to what Lord Levene or Lord Giddens had to say.

James Luke: I would make a couple of comments. First, picking up on how this will change jobs, we know throughout history that technology changes the way we work, and people move out of the fields into the factories, and so on. We have never seen a reduction in jobs. I see this as a huge opportunity for people to move into new areas, for professionals to be enabled and a bigger focus on creativity. In certain areas, if you look at healthcare, this is not about replacing what people do; it is about enabling them to deal with the sheer volumes of data they now have. Fifty years ago a cancer specialist would have had to review one or two hundred papers a year, if that. Nowadays, they see thousands of papers every year coming through, and they would not be able to do their job in the future.

I would also pick up on the point Andrew made about this need for an iterative approach, where you are prototyping and evaluating very quickly. Part of that is creating an ecosystem where employees, business partners and members of the public can access your data and access APIs. I had a situation with a major bank where they wanted a question-answering system. It was part of my education with the new technology. I sat down as an experienced engineer but with no skills in the programming language and no skills in cloud, at that time, and in four hours I put together a tool that would have taken me 12 to 18 months and a team of people a couple of years previously. A lot of the emerging technologies out there, the APIs, can allow you to do very powerful things. What you need to do is empower your employees and your business partners so that you can benefit from the innovation.

Kriti Sharma: I have some very quick comments. First, on automation, I believe that roles will change and there might be a significant shift with regard to certain jobs that exist today, but new ones will be created. We are already seeing that. For instance, in my team we have a conversation designer who is a linguistics expert and is training the personality of the conversations of the machines. That is a role that did not exist today. Accountancy is an area that has often been talked about when we discuss automation. We spoke with the accountants and the accountancy practices. A lot of them are ready to embrace the change and they see their role moving to more of an advisory or consulting position. There will be a shift in roles and new opportunities will be created. We are also educating those people, at the same time, to be ready for the next generation.

To address the small and medium-sized businesses point, in addition to all the great ideas we have heard from James and Andrew, there is also a responsibility that the software providers, such as Sage, SAP and IBM, have to take to create platforms for these small or medium-sized companies to scale up, or for enterprise companies to get most value with minimal investment. We can do that because we can build platforms

that function at scale. We need to raise our hand and take that responsibility.

Q78 Lord Holmes of Richmond: Good afternoon. What sorts of problems have you faced when developing AI for business applications? We have already gone into some of this space, but you might want to say something about commercial issues, the readiness of clients to be keen and enthusiastic about this area and any issues around data regulation. Feel free to cover any of those issues, or indeed none of them.

Kriti Sharma: One of the biggest challenges that we often see is over-expectation. The expectations of what AI can do have been raised so much higher and, thanks to Hollywood, there is a very different view of what this technology can deliver. First, there is a need to explain what it can or cannot do today. The awareness and education piece is very important. Secondly, AI has to be designed in the right way. It is a new field for a lot of companies. Although in one way you could argue it has been around since the 1960s, quite recently we have seen its adoption and explosion in the consumer world.

Designing it the right way includes the bias that exists in algorithms and things that machines learn, or even design principles. When you are building a chatbot, do you give it a male or a female personality? In our case, we do not give it any. Should it pretend to be human? When you are talking to a customer service agent, should you know who you are talking to? Is it human? Is it a bot? We believe you should be transparent about it. These are issues we had to tackle early on, which are more about the design or doing things the right way; explaining to customers what the technology is capable of doing and showing them the first interaction. That is the most powerful. When they see it working for the first time, it changes the users' mindset completely.

One more thing is to explain that this is not a new technology; it has been around for decades. AI did not appear last year; you are using it when you use Google Search or when you see something on your timelines. This is nothing to be scared of.

Andrew de Rozairo: We see a range of problems with our customers. Some of them still see this very much as an IT problem, whereas it is not; it is a business issue. Some of them are focusing their IT budgets on keeping the lights on and on improving data quality, and some of our customers have the perception that they have to get their data perfect before they can start mining it. We spend a lot of time trying to talk about needing to do these things in parallel. Yes, you absolutely need to get the foundation of your business and your transactional environment correct but, at the same time, you need to be innovating at a different speed.

The data issue is a big one, especially for us here in the UK. If you look at the kinds of proxies we have for gathering data, if it is IoT data or sensor data, according to Eurostat, we have the second lowest adoption of RFID in Europe. We have the second lowest adoption of electronic invoicing in Europe. Both these sets of data are incredibly important for us to be able to build rich algorithms on top. That is a key issue for us and one that we

have to help businesses overcome; the understanding that the data is going to be critical for them moving forward.

James Luke: I would agree with most of what has been said there. The only thing I would emphasise further is this challenge of defining the problem you are trying to solve. AI is a huge enabler for changing the way you work, for new capabilities. You do not know where you are going when you set out on the journey. Certainly, through the last two decades, we have been conditioned with an up-front business case, a return on investment and strict project milestones, and we need to recognise that sometimes a project may go in a different direction and still be very successful. I use a very simple analogy with my clients. I point out that pre-World War II the British Government first approached Watt not to buy radar, they wanted to buy a ray gun. What we got out of that was radar. The end destination is not necessarily where you had set out to.

The Chairman: is this the reason for the more agile project management taking place now across the board in this area?

James Luke: Definitely. We find that we do a huge amount of work with organisations now on very iterative agile programmes. It is important to remember that agile does not mean we throw away all the governance. That is not how agile works. Agile is very rigorous and we are constantly reviewing the progress and the direction. It gives us huge amounts of flexibility in how you design the ultimate business process and the user experience. We have gone through a massive transition ourselves in the last 10 years. We have recruited huge numbers of designers who specialise in working with our customers to understand the end client, the personas we are dealing with and designing completely new business processes. We do rapid prototyping behind that so that we can show and test with customers the impact of the technology and whether it works. Then we have to put in the hardening to make sure this stuff works at scale, is robust and has the good engineering behind it.

Lord Holmes of Richmond: In a nutshell, because I think it is a critical part of what you have just said, how would you define agile, as compared with more traditional governance approaches?

James Luke: I am at risk of giving a definition of agile and being exposed by the agile evangelists. I will not give you a formal definition; I will give you my personal definition. An agile approach is one where we are highly iterative; where we are led by the end user, the customer and the business problem; where we deliver to fixed timescales but not necessarily to fixed levels of functionality and where we are continually reviewing progress against the overall business goals in a highly user-centric approach. I would also say that just as in AI there are multiple different AI techniques, and you should not rely on one and you have to build an entire system, there are many different ways of building a project. Something I stamp on very hard is when I see agile proposed for a project that should not be delivered as an agile project, or vice versa. You have to use the right method for the problem.

Q79 **The Lord Bishop of Oxford:** You are giving a very hopeful picture, in contrast to some of the people we have interviewed, but a picture of significant change. One thing that can sometimes be left behind in

changing processes is ethics. Can you talk to us about the role that ethical principles play when your companies develop new AI systems?

Kriti Sharma: At Sage it was not an accident or an afterthought for us to build ethical guidelines into the system; it was part of the design and inception process to begin with. The way we approached it is slightly different. We learned a lot during our first few experiences and we created what we call the ethics of code: a few principles that every developer, every product person—so people who are solving day-to-day problems using AI—will adhere to and so forth. These are simple. For example, AI should reflect the diversity of its users, which means algorithms should not just be trained on data that is built with certain kinds of users in mind; it should be for a diverse user base. I mentioned some points earlier, such as AI should not pretend to be human because humans start to trust other humans, and it leads to all kinds of issues. AI should level the playing field. We talk about working with ecosystems and other start-ups or creating sharing mechanisms and holding algorithms and machines to account. We do not let humans run wild without holding them accountable for decision-making, so we try to do the same with AI. Going back to the skills and jobs point, if AI is replacing it must also create. We focus a lot on bringing a diverse group of people together.

We published this guideline, we shared it with the industry, and I have shared it with my 1,500 AI developer buddies in London. We are getting very good feedback. We would love this Committee today to encourage those kinds of guidelines that the industry shares, because this will not work if it is just Sage, SAP or IBM doing it in silos alone.

The Lord Bishop of Oxford: I like them and I like the fact that there are only five, and they are very succinct. Would you say that you are developing to a higher ethical standard than the consumer use of AI? I am thinking particularly of the “pretending not to be human”. Siri and Alexa seem to depend on anthropomorphic tendencies.

Kriti Sharma: Yes. A big reason for that is because we are a business application. We are dealing with customers’ financial information, their accounting data, their human resource people and payments-related information. Customers hold us to a different level as well. They expect a higher level of security and data protection. Ideally, it should be the same in the consumer and business world, but because of the information and the critical nature of the applications we hold and perform for our users, we need to be held accountable to a higher standard. That is how we have started it.

Andrew de Rozairo: In addition, similarly, we have a code of conduct for good artificial intelligence governance and best practice. We would look to try to get that endorsed by policymakers. We need to have two other discussions. One is at an industry level. There is something called the partnership on artificial intelligence, which is about how artificial intelligence will benefit people and societies. It was co-founded by IBM, Microsoft, Google, Facebook, Amazon—some of the big names. But it also has representation from Amnesty International, human rights, UNICEF and the Electronic Frontier Foundation—people who are talking about civil liberties in computing-type environments. I think there is an industry

layer. Each company needs to come up with their code of conduct as well. There is also a multi-stakeholder dialogue needed nationally. We are currently working with the France AI strategy group that is building a response to this, we are working with the German stakeholder group and we would welcome an opportunity to do the same here in the UK.

James Luke: Ethics are core to our business. We have been in business over 100 years and you do not last 100 years if you are not ethical. Data ownership, privacy, and trust around data are areas where we have published our own standards on data responsibility. It is critical that trust is front and centre in the development of AI. There is an evolutionary aspect to this as well. We work with a lot of businesses and organisations around the world, and we all know how important ethics are and the ability to trust the systems we deliver. If those organisations did not trust us with those systems and did not trust our systems, we would not survive. Trust is critical to the adoption of AI. Ethics is front and centre, certainly in our organisation and, you have heard, in the other organisations here.

The only point I would ask the Committee to consider is: is AI different here to any other new IT system coming along, where we are delivering new capabilities? Ethics should apply across the spectrum. I do not know if there is a particular reason to single out AI, but we are delivering new capabilities based on trust and data, and we need to have trusted systems across the board.

Lord Hollick: IBM has said that all AI systems should include explanation-based collateral systems. We have heard from a number of witnesses that that is not possible. Are they wrong?

James Luke: It is a question of whether you looking at this at an algorithmic level or at a systems level. Let me give you an example. We may look at a system that takes image data from one data source and performs classification on it—let us say facial recognition—it may take data from other structured repositories and perform another type of analytic on that, and it may fuse that together. There are certain types of artificial intelligence where it is challenging—and this is the bone of contention—that it is impossible to do explanation-based. In image classification using a neural network it is not possible to say at present why a neural network made that decision. It is hard in many of those cases for a human being to know why they made that decision. However, we have to look at the overall system and the information you put together. It is the fact that we have taken an output from one system, which may be deep learning, but we have fused it with data from another, and we may have had a human being in the loop. Ultimately, it is the system as a whole we have to consider, not the specific algorithm.

Lord Hollick: The example given was of an autonomous driverless car driving into a bus queue, and that it would be impossible to reverse-engineer it to find out why it took that decision.

James Luke: There is an overly simple view there that it would be making the decision based on a single algorithm. It goes against the message I am trying to communicate today. There is a view in the hype out there that if you take enough data and throw it in a machine-learning

algorithm the system will suddenly learn how to do things such as drive vehicles. My experience, as an engineer, is it will not.³ We need to engineer very complex systems that have multiple AI components that interact together and work. If you have a system like that, rather like a military command system or an air traffic control system, it will be taking data from multiple sources, there will be multiple sub-processes going on, there will be weighted voting and other forms of fusion going on in there, and you need to look at the overall system. That overall system should be able to explain why it has made that decision.

The Chairman: I am going to ask each of you, briefly, whether you expect your business customers to adhere to the ethical principles. When you supply them with product, do you expect them to adhere to those ethical principles?

Andrew de Rozairo: Absolutely. Part of the code of conduct we are putting forward, and I am assuming it is the same for Sage and IBM, is talking about taking a look at the training data you use, and making sure there is no inherent bias in that data; taking a look at the output of all your algorithms, and making sure that it is well tested and that you can document cases around it. We all sell platforms that enable artificial intelligence to be implemented at our customers. We cannot mandate what they do with these systems, but we can encourage them to behave in certain ways and we can provide education to say, "This is good AI practice".

The Chairman: It is not a contractual requirement.

Andrew de Rozairo: No.

Kriti Sharma: When it comes to developing on our platform, it should conform to the standard practices that we follow for APIs. Apple has a process of screening every app listed on the App Store, and they need to follow the security protocols of data protection. This should not be any

³ Note by Witness: If we consider a driverless car, there will be multiple sub systems. There may be a vision system that identifies and classifies vehicles. There may be a laser range finder that calculates the distance to objects around the vehicle. There may be a GPS system that identifies the location of the vehicle and this system will be supported by mapping data. There may be a noise classification system that identifies the sound of sirens. I would expect all this data to be fused together and fed into a decision engine and I would expect the output of the decision engine to be explainable. I would expect the system to state, *"I made the decision to pull over to the side of the road because the noise classifier detected a siren approaching from behind, the vision system identified an emergency vehicle approaching from behind, the road was straight and it was a safe place to stop without causing an obstruction"*. Within this explanation there are "classifications" that may have come from an AI system that cannot fully explain its reasoning. For example, the vision system may not be able to explain how it identified an emergency vehicle approaching. However, there are many situations when human beings cannot explain image or sound recognition. Even then however, it would be possible to design a Machine Learning system so that it retrieves the training data that influenced its classification decisions. The vision system may not be able to explain its classification but it should be able to say, *"here are the closest matching images in my training data"*. In summary, whilst certain algorithms cannot explain their decision making, overall systems should be engineered to be explainable and machine learning systems should be able to support their decisions with training evidence.

different. We believe it should be the same. I wanted to touch on this ethics conversation. First, I am very pleased that we are all doing it but, at the same time, I worry that at the moment it is based on individuals within organisations who have volunteered, who really care about it and are consciously making an effort. This should be the norm. It should be the practice in every company. That is what we should be driving at. Partnership on AI is a great organisation, but it is probably fit for large companies with huge budgets to sign up to. What if I am one developer in one company just exploring AI? Do I get access to it or not? We need to enable and empower everyone who is experimenting with AI or building new systems to have the same access as everyone else.

The Chairman: James, is it large company, big budget?

James Luke: I would agree with what my colleagues have said. Ethics and how we approach our projects with our customers are front and centre. When we review projects that we are undertaking we consider the ethical aspects, but it is not just a process, it is embedded in the culture. Every IBM-er every year has to go through a set of training we refer to as business conduct guidelines, where we look at the ethics of how we work and operate. It is embedded in the culture. I certainly have been in one client situation where we were asked could AI not deliver a certain capability, and we turned round and said, "We're not prepared to do that. It's not ethical".

The Lord Bishop of Oxford: That was hugely helpful, Kriti. You make a really important point that it should not be left to individuals and individual companies. Do you have a view on where it should rest and which body should regulate it?

Kriti Sharma: I would not even call it regulation. It reminds me of social media policy: "Hey, guys, don't post anything on Twitter that's not going to look good", as an employee. It needs to be along those lines; a corporate governance model where there are senior people—

The Chairman: "Comply or explain" type corporate governance.

Kriti Sharma: Yes. People at the executive level should care about it. I am concerned about developers or product people who are somewhere in large organisations, or small and medium teams, working on it and building systems. They will go to the executives at some point with a business case and say, "Will you fund it?" and that is probably too late. It needs to have that senior stakeholder ownership, and then be deployed it across the industry.

Baroness Grender: How would you ensure that that happens?

Kriti Sharma: The Government could help us to facilitate that. We have all created a certain ethics code. We could collaborate and identify that ultimate checklist and then share it with the industry bodies and have the executives and boards consider that as the things they need to care about in addition to all the other risks they care about in the business.

Q80 **Baroness Grender:** Thank you for that, Kriti, and for telling us not to be scared. I am going to ask now about the prejudice issue. Thank you so much for everything you have said about this in advance already. I have

heard about guidelines and sharing good practice from you all, yet I feel quite impatient about this issue. We know that there is a lack of diversity among engineers and we know that there is a significant danger that the current bias will be built into whatever AI looks like in the future. I would like to hear some very specific proposals from you. Guidelines feels a little passive to me. Do you have anything else in your armoury to solve this?

Kriti Sharma: It probably bothers me as much as it bothers you. Algorithms are making decisions everywhere in our lives. They are deciding who gets hired, who gets fired, who gets promoted, pay rises, the bonus numbers, and it is learning from historical biases. We have experimented with creating diverse datasets giving specific actions. I do three things when I build AI systems. First, I create diverse datasets. It is not that hard. We now have the ability to create these. If data does not exist we need to work hard, we need to work together and focus on open datasets. That is what we are doing.

Secondly, I look actively for bias in the system. It is not going to pop up on your screens saying, "Your algorithm is biased"; you need to make an effort to detect whether systems are biased or not. Some interesting research work has been done, but it has not been commercialised at the same level. There are people like me trying to actively look for it, but it needs more funding from government—possibly through the Challenge Fund or Innovate UK; bodies that are currently funding challenges on robotics, AI and automation systems in the real world, but not enough on the bias detection and ethics work. The good news is it is possible to identify it and tackle it, but we need to be actively looking for it.

The third way I think about it is the design of the AI system itself; specifically, when it has a personality and a tone of voice. It bothers me that we have all these servile female assistants turning lights on and off, booking your cabs and scheduling appointments for you, and then there are bankers and lawyers with very male personalities. AI does not need to have a gender. Why do we even create these machines? Why do we give them so much power? It needs to be tackled at these three levels.

In addition to that, my wild card, the fourth, is that we have an elitism problem. I am a computer scientist. This is what I have done. I have studied it and invested all my time and energy into it, so I am in some ways the AI community member. That is insane. We need to enable more people to come into this industry who will do great work. I spend a lot of my time, when it comes to skills development, with young people who may or may not even be able to go to university. The good news is that Dame Wendy Hall and her report are addressing the master's and PhD level problems. They have asked for researchers. That is awesome, but we need to do a lot more work for those young people still in high school, making it interesting and fun for them. As I said earlier, the day-to-day work of an AI engineer is so much more fun today than it was seven years ago. You can focus on problem solving. It is going to be even more interesting in future, encouraging more people to come in. We all need to play our part in that.

Andrew de Rozairo: Some great points were raised there. There are two parts to solving the bias issue. One of them is recognising if the data that you are using is getting your AI system to bias something one way or another. That is about taking a look at the training dataset. If you provide better tools to identify that bias, it should be easier for companies and individuals to start taking a look at that.

On the flipside, we also have a big opportunity here. As soon as you recognise bias in this training data, you can fix it; you can choose a different training set that becomes unbiased. Translating that into action in terms of behaviour, the machine now behaves consistently. We know how difficult it is to get someone, once you have shown there is bias in their decision-making, to try to change. That is much more difficult. Getting the machine to act in a less biased manner is just a matter of addressing that initial training dataset. There is an opportunity for us as well as a challenge.

James Luke: As the father of a 17 year-old girl who was the only girl in her GCSE computer science class, I am passionate about this. I also have a colleague—another IBM-er—who has developed Machine Learning for Kids, which is online and is encouraging young children to get involved in AI and machine learning. I fully agree with your point. We need a much more diverse workplace in engineering, as a whole. We are continuing to strive towards that, for one simple reason: it makes better solutions. We have better products when we have diverse teams. We have products that better reflect the needs of the market. It is the right thing to do. We need to bring more people into engineering from more diverse backgrounds. There is a whole raft of things we can do around that. We have talked a bit about education. Even the move towards more open APIs and more open datasets creates more innovation and gives more people access to be innovative and creative.

From an engineering perspective, I agree with the point Andrew has made about representative data and making sure there is transparency and explainability in the data so that you can identify bias and correct it. It is about addressing issues across the spectrum. Fundamentally, encouraging a more diverse engineering community is core to this.

Kriti Sharma: Can I make a very quick point?

Baroness Greender: Yes.

Kriti Sharma: AI can help us fix some of the bias as well. Humans are biased; machines are not, unless we train them to be. AI can do a good job at detecting unconscious bias as well. For example, if feedback is given in performance reviews where different categories of people are treated differently, the machine will say, "That looks weird. Would you like to reconsider that?" Many times a woman makes a point and a guy makes a point right after that, and people just remember the second one, and vice versa. We can train AI to detect and solve issues with these little prompts. It is a system that can run and help solve some of those issues.

Baroness Greender: Who should detect the bias? Should it be with you or should it rest elsewhere? Should you do it?

Kriti Sharma: In the systems we build we should do it. At the same time, there needs to be a lot more work done. If I have discovered something or a researcher has uncovered something or some research has been published, we should be sharing more. We have not done that very well. Also, when we publish AI research reports, or publish papers, we should be publishing the datasets used for the papers as well. It goes into the reproducibility element, where anyone should be able to validate and build on top of it. Dame Wendy Hall's data trust idea is great. Perhaps organisations such as the Open Data Institute could help with that.

The Chairman: Sometimes it is a bit more subtle than going in, finding the prejudice and solving it. For instance, Watson Oncology is based on a US dataset, as I understand it, the Sloan Kettering database. That is not quite prejudice, is it? It goes to the fundamental nature of the way the product has worked on the dataset, which then comes up with a set of solutions that are appropriate for one set of circumstances but not for another. That is a kind of prejudice, but it is a rather subtle form, is it not?

James Luke: From a technology perspective, when we talk about machine learning and we talk about bias, we mean exactly that. We mean that the training data you have used is not necessarily representative of the real world or the world you are applying it to. It is good engineering to make it representative, and the system will perform better if it is. You are correct: if you were to pick up a dataset from one part of the world and try and apply it in a different part of the world or to a different organisation's data, you need to ensure that the training data accurately reflects that change in circumstances. That comes back to the point we have all been making all along: this is not about throwing data blindly at systems; this is about engineering systems, understanding the data, understanding how the systems are performing, making sure there is transparency and explainability in the systems, and delivering solutions.

Q81 **Baroness Rock:** We have talked quite a lot about skills and broad education. Perhaps we could move on from that. As the mother of a teenage girl who wants to be an engineer, I am delighted to hear about your daughter. I would love to know more from you on the next generation talent pipeline. What more can we do to encourage young people? Dame Wendy Hall talks about the PhD, university level, but going into primary and secondary education, what more could we encourage the young to do to broaden their knowledge and understanding of the broader technology questions?

James Luke: I am really passionate about this, simply because I hark back to Victorian times when engineers were up there with footballers or great people. Engineering is a fascinating career and we need to get more young people into engineering. In IBM, we do huge amounts of reach out to schools. I am going into a school in January to talk about careers. We do a lot of work with hackathons, bringing children into schools. I mentioned Machine Learning for Kids. We bring children into the laboratory in Hursley, where we expose them to all the great things that

engineers do. In our education system, as a whole, we need to show people what science can do—the breadth and range of opportunities available. We need to build on all the great work that has been done in the universities. The UK has an incredible track record in artificial intelligence at the university level. I hope we will take a lead in the future in stepping out to this concept of a systems approach, and move beyond algorithms to how you apply whole systems. I hope the UK will become a thought leader in that area and create better role models for engineers to attract young people into the field.

Andrew de Rozairo: Those are really good points. I would love to see us celebrating some of the AI successes that we do see out, whether it is in industry or the public sector, so that people get more enthusiastic about it. A lot of the press focuses on the challenges and the downsides to this. There is a huge amount of welfare, citizen well-being and benefits that AI can bring. We need to celebrate that for people to have the aspiration of wanting to pursue an education or a career in this space. I would love to see UK Government taking a more proactive role in spearheading some of these AI initiatives. I would love to see more being done, perhaps around transport, where we can clearly articulate the benefits to all citizens and make a showcase of it. It is about creating demand for it rather than trying to push supply.

Kriti Sharma: We need to do 10 times better than we are doing. Today it is optional; you can choose to do GCSE computer science if you are one of the lucky kids whose parents are in this room. Every child should have the opportunity to experience this technology and do things with it before they choose whether they want to continue to study it or not. It needs to be done in the real way, not in a way where the teacher is trying to do something theoretical or in a computer lab with boring old-school technology. There is really fun stuff out there, which is what we are using every day when we teach kids on next generation—or Bot Camp, in my case. I am sure there are a lot of people out there who would love to help. My vision would be that every child in this country goes through that experiential AI technology programme and gets to play with robots, and there is a hackathon, and then decide what they want to do. It needs to be done in a much stronger manner. If we need to work on the GCSE computer science curriculum, make it agile or iterative, we probably need to do that.

Lord Levene of Portsoken: To follow up on that, Kriti, clearly, children are taught various basics—English, maths, et cetera. Do you envisage they ought to be taught this as well, and at a very young age, so it is one more building block in their understanding of what is going on? If so, is there any movement towards doing that and bringing it in for children so that, first, they are not frightened of it when they grow up, and, secondly, that they have a basic understanding of it?

Kriti Sharma: Yes. That is a very good question. The way to think about it is that it is not in isolation as yet another building block. We have talked about diversity in various different backgrounds coming together: the application of AI to solve mathematical problems, or writing, or artistic work. Bringing those fields together would make it much more fun, rather than writing a piece of code for the sake of writing a piece of

code. When I think about coding and some of the bot camps we run, it is about problem solving. You are not just writing a piece of AI software. No, you are solving a problem, and the problem could be in any of the other disciplines they are learning. It is not in isolation as yet another building block; it should be connected to everything else they are learning.

Viscount Ridley: To follow up on Lord Levene's point, is there an opportunity to retrain people into this area and fulfil the needs much more rapidly? I know of a business that takes physics PhDs and retrains them in short order into data scientists. There are lots of people who get to the end of degrees and think, "Oh my God, I'm not as employable as I thought I would be", but are very smart and could, with relatively short training, be turned into AI or data scientists, or something like that.

James Luke: We employ IT architects and specialists who come from a whole range of backgrounds, from apprentices up to degree-qualified, with a whole range of degrees. You do not have to have come from a purely computer science background. I have an engineering background but we have mathematicians; there is a colleague of mine, a distinguished engineer, who has an English degree; you can come from all sorts of backgrounds. I would pick up, though, on a point from the previous discussion which I think is important, about the curriculum, which is about creativity. We tend to talk about sciences and we talk about the creative arts. Science, engineering and mathematics are wonderfully creative subjects. We often do not appreciate that or communicate that to children in schools. Perhaps teaching creativity rather than maths or science would be a huge opportunity for us. I go home every day looking back on a day where I have been creating wonderful things, and it is a very stimulating life. I wish we could communicate that to our children.

The Chairman: Do you all take apprentices, by the way? That is an absolute given. I will not move on to T-levels, and so on and so forth, but we may come back to that.

Q82 **Lord Swinfen:** Can you tell us how we can improve access to data for small and medium-sized enterprises and for researchers? What do your companies do to facilitate this? Is it available to them now? Do you believe that the monopolisation of the kinds of data needed to train artificial intelligence systems is a problem? How can this be addressed, if it is? Do you believe that the present market is competitive?

James Luke: We make data and APIs open and accessible. We are very committed to open systems. We have to recognise that from a commercial perspective data is also something that other companies invest in and data is a product going forward. There is nothing wrong with companies doing well out of managing to secure large datasets and put analytics on top of it. It is very important, and part of the trust we talk about in building AI systems, that when commercial organisations work with us they know their data is secure and that their privacy is respected. More sensitively, across the community as a whole, we have to respect the privacy of individuals and, again, we have to recognise that in how we make data available. I do not think it is as easy as saying

everybody should suddenly see all the data; we have to respect the privacy and the commercial interests of our customers and business partners. The commitment to make APIs open and to open up data where that is possible is an important step in taking AI forward.

The Chairman: Is that something you recognise, Kriti, as an issue?

Kriti Sharma: Yes. A lot of the partnerships happening in the industry today with academia or start-ups are very much on a piecemeal ad hoc basis. You happen to find the right person in the large organisation, they are excited by what you are doing and will get you access to the datasets. I believe we need to do something more strategic than that. Any organisation that is hosting data should publish clear guidelines on their sites: "If you are a researcher or a start-up and you want to work with us, these are the steps to follow or you submit a request or you tell us a little bit about yourself". It is not an easy problem to solve. We cannot say, "Everyone open your data", because, in addition to the commercial sensitivity, there are also compliance and security reasons. For example, at Sage we process payments data, and not every organisation we work with will be compliant to even deal with that data. We must be very transparent about it, that "This is the process you need to follow if you want to work with us". My understanding, at least from the people I have spoken with in the industry, my peers, are very open to working with other companies, with start-ups and with academia to develop more solutions. We do not want to be building everything ourselves. There is a lot more value in sharing. If they want to do that, it is about ensuring there is a clear process for us to know where to go.

Andrew de Rozairo: We can do a couple of things to improve the openness of data. As companies start to use the data and realise the benefits they get, not just from their own data but from other sets of data, they are more likely to open it up. You are more likely to contribute if you have also received some benefit from it. Continuing with more public sector open data is a great way of showing people, "Here's all the additional benefit you can get by combining your enterprise data with some external data". If they see the benefit in that, they will reciprocate.

The other thing that stops some companies from sharing data is they do not necessarily have, let us say, the legal resources to put the right framework agreements in place to share data. That is something where, again, a government-led approach could be very useful; some guidelines around IPR, privacy, and so on, that are made available to small businesses who do not necessarily have the legal resources to draft this from scratch themselves.

Lord Swinfen: Do you think the Government should set up a general data library that firms put their data into for other people to use when they need to do research or whatever work they are doing? At the moment it does not seem as though there is. Is the wheel continuously being reinvented?

James Luke: I would say it is an interesting idea. We would need to examine it in more detail. The devil is in the detail about the sorts of issues we have raised. It is certainly an interesting idea, but I do not think we can give a general answer yet.

Q83 Lord St John of Bletso: In your evidence today, you have spoken mostly about the opportunities of AI. Andrew, you mentioned that the media tend to sometimes talk about threats. Of course, we have an Information Commissioner, but do you believe that we need a specific AI watchdog or regulator to protect consumers with respect to AI-based goods, products and services?

Andrew de Rozairo: I would say that AI is one of the many tools we are using within business. It is already integrated today into business fabric, and it is going to be much more so in future. I do not think that necessarily taking a look at, let us say, the narrow view of artificial intelligence makes as much sense as keeping the broad view and recognising that artificial intelligence is one of those tools. If, for example, we are seeing bias in an artificial intelligence algorithm that is shortlisting candidates for a job, that is not about AI, that is about bias in the system, and that should still be very much the focus of ICO. I do not see a separate regulator necessarily because AI is so heavily intertwined with the general business processes.

Kriti Sharma: I would agree. A lot of it comes from the applications of AI in the industry you are talking about. It is hard to establish generic guidelines, as you will have to do a custom break for specific applications of AI in financial industries, with taxation and accounting, in one state versus the other. Healthcare is completely different. We need to look at it at a data level and at the application level. Data has certain controls in place and the same with healthcare, the financial sector or insurance.

James Luke: I would agree. I do not think there is a need for a separate form of regulator; it is already covered by many other industry regulations and other policies. We have GDPR coming through and it should be embedded within those.

The Chairman: It is the ICO, at the end of the day, that would cover the responsibilities for data, which then leads to other matters in the GDPR such as algorithms. Nods all round, for the record.

Lord St John of Bletso: I have a small supplementary question. James, you mentioned earlier that trust was critical to the development of AI. The question is: do we need to update data protection laws?

James Luke: My understanding is they are being updated with GDPR.

The Chairman: Are they being updated enough?

Kriti Sharma: I spent some time reading GDPR documentation, which was fun. There are some very interesting points there. Over the next couple of years, with the new technologies—and we have seen some massive disruption in the AI industry itself with AlphaGo Zero where they could train AI with no data whatever, just a few weeks ago, and other similar breakthrough technologies—we might have to look at it again. On my analysis of the current GDPR, I think it covers it quite well. There are some guidelines, specifically annex A, which I studied, and it had some very interesting guidelines on how to make algorithms explainable, or give people the right to get an explanation using a snapshot in time. For now, we are fine, but I would recommend that we look at it in the near future.

The Chairman: Andrew, you nodded and so I will take your answer as a nod that you agree with Kriti.

Andrew de Rozairo: Yes.

Q84 **The Chairman:** I come on to the last question. If there were one recommendation that you would like to see the Committee make at the end of this inquiry, what would it be? I am afraid you are confined to one only.

Andrew de Rozairo: If I had one recommendation, it is finding ways of getting businesses to recognise the benefits that AI will bring and recommend the adoption of it. If we adopt AI, given the strong skillsets that we have in the UK, we have a huge opportunity to boost productivity. If we do not adopt it, we are going to lag further behind. It is, more than anything, about adoption.

Kriti Sharma: He has my card, so I have another one.

The Chairman: It is this embedding point, is it not, that is so important for the future? I take that. You have another recommendation.

Kriti Sharma: Yes. We need to do a lot more with skills. We need to create the best curriculum in the world for kids going through the education process. We need to do it in collaboration with the industry, which is developing the latest tools to bring it back to schools. I would also highlight that AI is now getting to a point where it is capable of writing its own code. That means that the way we teach kids and other young people to learn AI or build AI solutions is going to change in a big way.

The Chairman: That was a very neat way of having two recommendations.

James Luke: We should understand that we are building, that we require end-to-end solutions and systems that have multiple components, and that therefore we need the skills and expertise to build these overall systems. That is reflected in all the comments we have made earlier today.

The Chairman: Thank you very much indeed. I am afraid we have overrun because it has been such a good session, for all the right reasons. Thank you very much to all three of you for a really interesting session. We will benefit greatly from having talked to you.

IEEE-Standards Association and Professor Alan Winfield – Oral evidence (QQ 18 – 28)

Evidence Session No. 3

Heard in Public

Questions 18–28

Tuesday 17 October 2017

[Watch the meeting](#)

Members present: Lord Clement-Jones (The Chairman); Baroness Bakewell, Lord Giddens; Baroness Greender; Lord Hollick; Lord Levene of Portsoken; Viscount Ridley; Baroness Rock; Lord St John of Bletso; Lord Swinfen.

Examination of witnesses

Professor Alan Winfield, Dr Ing Konstantinos Karachalios.

Q18 The Chairman: Can I very warmly welcome Professor Alan Winfield, who is professor of robot ethics at the University of the West of England, Bristol, and Dr Ing Konstantinos Karachalios, managing director of the IEEE Standards Association? We are delighted to see you, and thank you very much for coming today.

I have a little rubric that I need to go through every time we have a witness session which I will go through now. This session is open to the public. A webcast of the session goes out live and is subsequently accessible via the parliamentary website. A verbatim transcript will be taken of your evidence. This will be put on the parliamentary website. A few days after this evidence session, you will be sent a copy of the transcript to check for accuracy. We would be grateful if you could advise us of any corrections as soon as possible. If, after the session, you wish to clarify or amplify any points made during your evidence, or have any additional points to make, you are welcome to submit supplementary written evidence to us.

First, perhaps you would like to introduce yourselves for the record, and then we will start with the questions, starting with Professor Winfield.

Professor Alan Winfield: I am an old-fashioned chartered engineer. I worked on safety-critical systems in industry before moving back into academia 25 years ago. I co-founded what is now the Bristol Robotics Laboratory, and, as you said, I am a professor of robot ethics in the University of the West of England, Bristol. I am probably the only professor of robot ethics in the world, which is kind of cool, I think.

The Chairman: I suppose you wrote your own job description.

Professor Alan Winfield: Something like that, yes. I am also a visiting professor at the University of York and associate fellow at the Leverhulme Centre for the Future of Intelligence in Cambridge.

To give you a brief summary of what I have been doing in robot ethics for the last 10 years or so, first, I co-organised a joint EPSRC/AHRC working group that drafted the ethical principles of robotics. Those were published in 2011⁴. That led me to become a member of the British Standards Institution working group that drafted BS 8611, a guide to the ethical design of robots and robotic systems. We believe it is the world's first published ethical standard in robots. To come right up to date, I am involved with Konstantinos in the IEEE Standards Association's global initiative for ethical considerations in artificial intelligence and autonomous systems, and I am a member of its executive. I also co-chair its general principles committee and, for my sins, I am chair of the Working Group P7001, which is developing a new standard on transparency in autonomous systems⁵.

The Chairman: Thank you very much indeed. Professor Karachalios.

Dr Ing Konstantinos Karachalios: First, thank you for pronouncing my name correctly; it is not so easy. Secondly, I am also an engineer. I studied in Germany and have a doctorate in nuclear reactor safety, so I understand a bit about safety issues. I worked for the public interest at the great European organisation, the European Patent Office, in many functions. My path led me to my current position as managing director of a global organisation that makes standards in every aspect of life you can imagine. You may know the wi-fi family of standards, for instance; it is produced through my organisation. Every time you connect to the internet you go through the protocols that we created. This is quite influential. That is my management role. There is another part of my personality, and Alan mentioned it: I launched the initiative on ethical aspects. There is a reason for this and I hope this will come out in this discussion. It is the reason why I am here in front of you today.

Q19 **The Chairman:** Thank you very much indeed. I am going to start with a broad question; you have seen the way they develop. Does the development and use of artificial intelligence give rise to new and distinctive ethical issues, or are they variations or pre-existing ethical questions? The question is: will established ethical principles suffice, or do we need new ethical principles? Obviously, we may go on to talk about the kinds of mechanisms that may be needed as well, but perhaps you would start by answering that broad opening question.

Dr Ing Konstantinos Karachalios: This is a very big question. I watched the previous sessions you have had here and they have tried to answer part of it. I was very intrigued by the session with the journalists, I must say. They posed some very critical and interesting questions. I would like to start from there and not repeat them. There is a series of issues which is part of the work we are doing in the initiative. We have produced a book. The second version is about to come out and I would be pleased to send this to you because it has chapters with all possible

⁴ Note by witness: Engineering and Physical Sciences Research Council, *Principles of robotics* <https://www.epsrc.ac.uk/research/ourportfolio/themes/engineering/activities/principlesofrobotics/> [accessed 5 January 2018]

⁵ Note by witness: IEEE Standards Association, IEEE Project 7001, *Transparency of Autonomous Systems* <https://standards.ieee.org/develop/project/7001.html> [accessed 5 January 2018]

questions, such as on transparency, accountability and legal issues, where the best people in the world have come together, framed them in a very succinct way and proposed recommendations and general principles to address them. I would be pleased to send to you the second version, which will come out in two months. I am sure that Alan will go into more detail if there are specific issues.

I would like to use my presence here to address first the positive side of these technologies that I see coming and that can accelerate, to put it very briefly, the promise of technology to satisfy the material needs of humanity. This is a big promise. It is the promise of Buckminster Fuller in the 1970s—in the last century—that war will be obsolete as long as we satisfy our basic needs. The question is: why is it not happening? We must put a lot of energy in to make it happen, because the potential is there already. This is very positive, and all these technologies can accelerate this. There is nothing new per se, but they can massively accelerate it.

Secondly, there are critical aspects, three of which are beyond technical expertise. The first is that code is above law. It has already become a reality. Where decades ago Lawrence Lessig said that code is law, code is now above law, and I do not believe that the policymakers should accept this. You should reclaim the territory that you are losing.

There are examples that I can give you if you would like, but I will go to the second point, which is that the possibilities that are given by these technologies, together with sensing ubiquitous data-gathering and so on, may lead to an erosion of democracy. This is very serious, and it is already happening. It means that if our private life cannot be distinguished from public life, nobody can be an autonomous political actor. This is serious and happening already. We must not kill democracy in order to protect it. These technologies play a massive role in this context.

The third metaproblem is what I would call the Stanislav Petrov theorem. I do not know if you know this guy. He died impoverished a few months ago in Moscow. On the night of 25 to 26 September 1983, he was in a bunker surveying the computer systems and saw that the computers were telling him that the US had attacked the Soviet Union. He ignored them. He said two sentences, and this is the theorem: “We know better. We made you”, and he ignored them. This is a very important action. If we lose it, we have no future. We must always be able to say, like Petrov, “We know you. We are better than you. We made you”. If not, there is not much future for us. Those are the three levels of risk that I see.

The Chairman: Thank you. I may come back to you, because clearly you are talking about new impacts, so in a sense you are arguing that we need new ethical concepts.

Dr Ing Konstantinos Karachalios: Yes.

The Chairman: We need to unpack that a bit further. Professor Winfield, what is your take on that?

Professor Alan Winfield: Like Konstantinos, I think it is a super-interesting question. It is a difficult question, because one of the

problems with artificial intelligence is that we do not know what natural intelligence is. There is a fundamental definitional problem. We are trying to build something to emulate natural intelligence, yet we do not have a scientific theory of natural intelligence. Let me be a little more specific.

The first thing I would say as an engineer is that AIs—let us call them AIs, because those are the words that we have in front of us—are just engineered systems like washing machines and word processors. They should therefore be subject to the same engineering best practices, and if those AIs are safety critical, such as driverless car autopilots or medical-diagnosis AIs, they should be held to the same high standards of provable safety and reliability that we would expect from aircraft autopilots or medical equipment.

To come to your question, at the same time, all social AIs, such as chatbots, personal assistant AIs, care robots, robot pets and robot toys, have the potential to deceive. In other words, the human might come to believe that the AI cares about them, and in turn the human forms an attachment with it. An AI is an engineered system and, of course, it cannot care about someone any more than a washing machine can. This potential for deception is a special property of AIs and something that requires special ethical consideration. It is not the only example, but it is the one that I would like to present to you as evidence.

The Chairman: The way you talk about it is very interesting. It is almost like stripping away the mystique, really.

Professor Alan Winfield: Exactly.

The Chairman: Just treat them as machines, basically.

Professor Alan Winfield: That is exactly right, yes.

The Chairman: I am going to bring in Lord Swinfen in a minute, but I want to ask you about the mechanisms. You have both accepted the need for new principles, in a sense, although you have not been specific. But, let us face it, through partnership on AI, from what you and the IEEE have put together there are principles that can be adopted. How do you put those into practice? What mechanisms are required to ensure that any principles you adopt are going to be applied?

Dr Ing Konstantinos Karachalios: There are several elements. I think every institution should do its best. No one player can face it and tackle it all. It is your responsibility to regulate; it is our responsibility to self-regulate. I can tell you what we can do and how we can contribute. We, the technical community, have a duty. We have a duty to stop obfuscating the dialogue around these things and put them under the spotlight, which is what Alan is doing right here, and to do a better job from the beginning. These are engineering systems, and we cannot let out software systems that have an impact on the physical environment without them being thoroughly tested, as beta versions or whatever.

We must force the software engineers to do a job like the ones the civil engineers do. You build a bridge and you are sure it will work, or almost sure. We must take into account our safety, dignity and privacy in the design and not as an afterthought. This is what we are trying to do at the IEEE; we try to design and develop systems that will help the engineers

and architects to do a better job from the beginning, otherwise we can only say, “Okay, we did what we did and others will deal with the consequences and find solutions for the problems we have created”. You must stop accepting this from us and we must do a better job.

Lord Swinfen: Do we need a definition of natural intelligence so that we can define artificial intelligence?

Professor Alan Winfield: Probably, yes, but it is very hard to give one. Here is a very simple definition: natural intelligence means doing the right thing at the right time. However, that is not a scientific definition or a measurable definition. The truth is that there is no adequate scientific definition of natural intelligence. In fact, there is no general theory of intelligence.

Dr Ing Konstantinos Karachalios: May I add something? We can assume that our brains are computer machines, Turing machines—and there are people who believe this—or we can say that our intelligence has a source that we cannot define; it is deeper. People call it the soul, or the spirit, or whatever, but there can be different sources feeding our intelligence. I would not exclude it. To believe that we are machines is a very reductionist view of human beings that has not only social but political implications. This is an interesting discussion and worth getting out. Unfortunately, it does not take place within the scientific community.

The Chairman: I am going to bring in Lord Levene, but, really, we cannot rely on the Turing test as being anything useful?

Dr Ing Konstantinos Karachalios: We must leave it open that our intelligence has other sources deeper than this.

Q20 **Lord Levene of Portsoken:** Professor Winfield, one comment you made rang an immediate bell with me when you were talking about washing machines. We have a mews house. The young couple who were living in the mews house went out on a Sunday morning and left the washing machine running. The next thing was that it caught fire and virtually gutted the house. Of course, you get the insurance company involved, which is also going to have a big interest in this. That is a minuscule problem. The major one, if what we have read is correct, is that the disaster at Grenfell Tower was apparently caused by a fridge. What, if anything, at the moment is being done to try to bring into a clearer light who is responsible for it?

Professor Alan Winfield: It is certainly true that AI systems complicate the question of responsibility, but in many ways they should not. It is the designers, manufacturers, owners and operators who should be held responsible, in exactly the same way that we attribute responsibility for failure of a motor car, for instance. If there turns out to be a serious problem, generally speaking the responsibility is the manufacturers’. That would typically be the case, but going back to the point we have just made about not granting AI some special status, we need to treat AI as an engineered system that is held to very high standards of provable safety. If that AI happens to be in a washing machine, we need standards for that AI.

Lord Levene of Portsoken: I have one quick supplementary question.

You say the manufacturer, but which one? Is it the manufacturer of the washing machine or the one that made the printed circuit board in it that failed?

Professor Alan Winfield: Typically, if it is an aircraft, we would regard the final manufacturer—Airbus, Boeing or whoever—as responsible, irrespective of the fact that there is a huge supply chain. Typically—forgive me if I am wrong; you probably know better than I do—it is the final assembler of all the subsystems that assumes responsibility for the safety and reliability of those subsystems.

The Chairman: I am sure we will come on to that with our next evidence session on legal accountability and so on, but in principle you are right.

Lord Giddens: The issue of what intelligence is is so central to this whole thing, I find the elusive and vague way in which people use it amazing. I would ask you to develop your point a bit more. I taught psychology in Cambridge for quite a while, and you always had to teach the evolution of the debates on intelligence. There is a huge literature. There are disagreements about whether Spearman's *g* exists. Nevertheless, it seems to me there is one really distinctive thing and it is that to be a human being you have to know semantics, and you have to know the meaning of things. You cannot know the meaning of things unless you act in the world, it seems to me, and you are the possessor of a society and a culture. That would never, as far as I can see, apply to machines.

Therefore, I want to know what implications one draws from that. I do not know if you know Searle's famous philosophical attack on the idea of AI: that it could never rival human intelligence precisely because, no matter how clever you are, as a computer speaking Chinese you would not know the meaning of the words you are saying. That is what Searle argues.

Professor Alan Winfield: The Chinese room argument.

Lord Giddens: That is a really substantial point, because what intelligence is is the core of it all, and I have not found many people before, listening to you today, who have made that connection and explored that relationship. It seems to be so crucial to me.

The Chairman: We will come back to that, if we may.

Viscount Ridley: As a postscript to Lord Giddens, with whom I nearly always agree on everything, I am not sure I agree with him on this.

Lord Giddens: I agree that we agree on everything.

Viscount Ridley: Was it not Justice Potter Stewart of the US Supreme Court who said about pornography, "You can't define it but you know it when you see it"? Is that not good enough for intelligence?

The Chairman: Can we come back to that, maybe as an adjunct to answering Baroness Bakewell's question? Otherwise, we will run out of time, and we certainly do not want to do that.

Q21 **Baroness Bakewell:** We are getting in really deep here. Where is the responsibility to lie for ethical developments even if we cannot define natural intelligence? Someone has to hold responsibility. Should it be

within organisations? You speak about your institution being about bringing about consensus and why it takes so long. Can there be a consensus about something that is so amorphous? We use words such as “duty” and “obligation”, but all these are personally governed by a whole background of information, so where are we to lodge our trust in the decisions to be made?

Dr Ing Konstantinos Karachalios: We are not going to solve all the problems, but I think we can get better. Although it is vast and difficult, we can get more concrete. This is what we are trying to do. Also, morals and ethical values are not the same everywhere; they are very contextualised. What we are trying to do with the first standardisation project we have started, is to enable the software architects and engineers to take at all ethical considerations into account—they are not doing that now—to think and to be more self-reflective. This is the problem. It is not a problem for which we can find a perfect solution. You need to reflect upon what you are doing, and if you are not doing that you are a danger to everybody else; you are just a war machine. The engineering community must stop being a war machine against society. We must assume our responsibility and be self-reflective. This is what it is all about: it is an education of ourselves. We owe this to society.

Baroness Bakewell: But I must ask: who are “we”?

Dr Ing Konstantinos Karachalios: In our case it is the technical and scientific communities that have a duty. It is your duty to support us or to encourage us or to pressure us to do this, because nobody can do our job. We have to do our own job.

The Chairman: Do we need an identifiable cadre, therefore, of AI engineers? Is that what we are arguing for?

Professor Alan Winfield: I would say that more generally we need every AI engineer to be trained in ethics, and in responsible research and innovation.

Baroness Bakewell: Should we also have specialists as we do in hospitals? We have ethics committees, and the people sitting on those ethics committee have specialist knowledge.

Professor Alan Winfield: That is absolutely right, yes.

Baroness Bakewell: We need specialist ethical engineers.

Professor Alan Winfield: I would certainly advocate—and I have said this for several years—that we have reached the point when, for instance, AI research projects should be subject to ethical approval as medical research has been for many years, for the reasons that we are discussing.

Baroness Bakewell: Where are you to find a consensus on who they should be?

Professor Alan Winfield: That is a great question. There are two sides of the coin, one side is education, which we are talking about here, but the other side is standards, because it is in standards that you effectively express, articulate and formalise ethical principles, and these are precisely the routes that we are taking here in the (IEEE) initiative.

The Chairman: We have to have people like you setting those standards.

Dr Ing Konstantinos Karachalios: I found what Lord Giddens said very interesting, and I would like to say a few words on this; I have been thinking a lot about what Lord Giddens has mentioned. The conclusion I came to is that human intelligence is not the capacity to choose between options that you are given; it is the capacity to pose a dilemma that is there but which people do not want to see. It is there in the conditions, and it is the courage of your heart to pose a dilemma, not to solve it. The dilemma is very painful and you pay a price. Only humans have this capacity. A chicken has the capacity to choose different options—which seed to pick up—but a human has the capacity to pose a dilemma. This is something no computer can do.

The Chairman: Offline there may well be further discussion, but you make a very interesting point.

Q22 **Viscount Ridley:** In relation to the last conversation, I remember having a conversation with Ian Wilmot shortly after he cloned Dolly the sheep. In a similar session to this, he was asked to describe the science and then asked to leave the room while the grander people discussed the ethics. He was very cross about that and said, “No, we biotechnologists have views on ethics, too”. It is worth remembering that point as one goes forward, and I am sure you do.

My question is about transparency. We cannot be far from the point when artificial intelligence diagnoses a disease or offers a legal opinion without being able to explain how it reached that conclusion. That is already true of human beings in some sense, but is it a particular problem with AI, with robots, and how transparent should artificial intelligence systems be?

Professor Alan Winfield: I take a very hard line on this. I think it should always—always—be possible to find out why an AI made a particular decision. That is very easy to say, of course, but we know that in practice it can be very difficult. It seems to me absolutely unacceptable that one might accept the decision of a medical-diagnosis AI or a mortgage application recommender system without understanding why it made that decision. Many members of the AI community will get very cross with me—I am sure they are right now if they are watching—because, of course, what I am effectively saying is that we should not be using systems that are not transparent, such as deep learning systems. I would simply apply an engineering approach whereby we need to be able to understand why the system makes the decisions it does; otherwise, if a system goes wrong and causes harm, we simply cannot find out what went wrong.

I would like to offer the analogy of aircraft autopilots. We all understand very well the standards that we set for the engineering of those systems, and an AI autopilot and driverless cars, for instance, should be held to no less a standard of safety, and explainability or transparency.

Viscount Ridley: Is not AlphaGo already failing that threshold?

Professor Alan Winfield: It is, yes.

The Chairman: By which you mean deep learning?

Viscount Ridley: By which I mean move 37 in game two, which we were told about and nobody knows why it did it.

Professor Alan Winfield: My challenge to the AI community—and they are smart guys, smart men and women—is to invent AI systems that are explainable. I do not believe that it is technically impossible.

The Chairman: Has the horse not already bolted, though?

Professor Alan Winfield: It may well have done, except that we can still regulate—and I believe that we should—and say that it is simply not acceptable to have, for instance, a medical-diagnosis AI that cannot be explained.

Viscount Ridley: I have a supplementary question on that that I was supposed to ask. Does the degree of acceptable transparency differ depending on the situation? In other words, some of our evidence has suggested that it will be fine in some circumstances but not in others, so in a game it might be all right and in a medical diagnosis it might not be. Would you accept that?

Professor Alan Winfield: I would indeed. I am focused primarily, as you can probably tell, on safety-critical systems. That is where if the AI goes wrong, physical, financial or psychological harm could result—in other words, where harms result from a failing AI.

The Chairman: Dr Karachalios, do you agree with Professor Winfield's hard line?

Dr Ing Konstantinos Karachalios: Yes, and I believe that there are different aspects of it. Transparency is not explainability. It is different. A system that can explain what it does is different from a fully transparent system. There are efforts to make decisions more transparent. There are elements of the European Commission's GDPR that are forcing this transparency. This is the right way to go—and it can be done—to give direction to the researchers and scientists.

In addition, it is very difficult to understand why dataset A has been transformed to dataset B, because we do not remember the way computers remember. We do not have this capacity. Even if we see it, we do not understand how it came about. This is a problem that we need to understand, because if you do not get a job or you are refused medical treatment, nobody can tell you why. We need systems that can explain this, I agree with you.

Baroness Bakewell: Would your view that there should be knowledge of how these things work, governed by your ethical concerns, be shared beyond western culture and western thinking? Would they be shared with other more elusive cultures that may have a different concept of natural intelligence, conscience, duty and so on?

Professor Alan Winfield: The question of interoperability applies, for instance. If a manufacturer somewhere in the world wants to sell us a driverless car, that driverless car should be fitted with the equivalent of

an aircraft flight data recorder, and just like an aircraft flight data recorder the data contained in that should be available by law to an accident investigator, regardless of the culture from where the car comes. The point is if that culture wants to sell us driverless cars, we should insist that the logged data should be publicly accessible.

Baroness Bakewell: Can you imagine down the decades, with the development of different and divergent cultures, these values not being universal?

Professor Alan Winfield: I can, but in a global economy you are not going to help yourself economically.

Baroness Bakewell: So the driver would be economic, would it?

Professor Alan Winfield: I guess, yes.

The Chairman: You seem to be making a distinction between transparency and explainability. Would either of you like to unpack that?

Dr Ing Konstantinos Karachalios: Full transparency is having the source code to see how it is written. It is fully transparent but it does not explain anything. You do not understand it.

Professor Alan Winfield: An example might help to clarify. If you have a driverless car and it is involved in an accident, the accident investigators need to understand what happened to cause the accident. That requires transparency, not necessarily explainability. Explainability for the user of a care robot—say it is an elderly person with a care robot in their home—means that that elderly person should be able to ask the robot in some fashion, “Why did you just do that?”

Lord Giddens: Following up what Lord Ridley said, this would imply that you need analogous structures to those that exist in the professions, would it not? In medicine, you have organisations that oversee the decisions that are made, and obviously you cannot investigate all the decisions that are made, but in all these cases, including airlines, you have supervisory institutions, and the argument would seem to me you must have those also wherever IT is deployed.

Professor Alan Winfield: I think that is right, yes.

Lord Giddens: In the case of driverless cars, it is possible that it will never take off at all because you will not be able to solve the question of responsibility in law.

Professor Alan Winfield: Exactly.

The Chairman: Do you want to roll that into your question, Lord Giddens?

Q23 **Lord Giddens:** What system of accountability should there be for artificial intelligence systems? How far should they be held accountable compared to other institutions? In a way, I tried to suggest how I would answer that.

Professor Alan Winfield: The first thing to say, and I am sure that I speak for Konstantinos as well, is that an AI can never be responsible, so it is humans and not robots or AIs who are the responsible entity here.

The question then, of course, is figuring out who exactly is responsible in a given situation, and that is where the transparency is so important. If we do not have transparency, we cannot properly have accountability. That is probably not a very satisfactory answer.

Lord Giddens: Yes, except the most important thing, it seems to me, is to be able to enforce transparency where it is needed.

Professor Alan Winfield: Exactly.

Lord Giddens: You could never have a system in human medicine where every decision is transparent. You simply have a system whereby decisions can be questioned and a code of professional ethics.

Professor Alan Winfield: I absolutely agree. In January, the Alan Turing Institute published a report in which they recommended what they call an IA watchdog. I am not clear in my own mind whether you need a single body or a body that is specific to particular domains. Care robots might be one body, driverless cars might be another. I think we absolutely need to have those kind of watchdogs.

Lord Giddens: If I might interject, your example is good, because in airline systems you already have that.

Professor Alan Winfield: Exactly, and I think that is a great model of ethical governance.

The Chairman: You are not saying that the use of AI should be totally risk free, because there is bound to be an element of risk in autonomous vehicles or whatever it may be.

Professor Alan Winfield: It cannot be.

The Chairman: When we talk about accountability, it is accountability in so far as you do not accept risk yourself?

Professor Alan Winfield: I am sorry, I am not sure I understand the question.

The Chairman: There is an element of risk that you accept.

Professor Alan Winfield: That is absolutely true.

The Chairman: You cannot expect AI to be totally risk free, so you are not 100% accountable for the consequences in every case for the use of AI.

Professor Alan Winfield: Indeed, yes, and there is no question that there will be accidents for which there is not.

The Chairman: Otherwise we would never adopt AI in any particular respect.

Professor Alan Winfield: That is exactly right. Sorry, I misunderstood.

The Chairman: I am just testing the proposition, basically.

Professor Alan Winfield: Yes, I agree.

The Chairman: You have a very robust approach to the adoption of AI, but the question is: does it go all the way to saying that you are legally 100% liable for every consequence that occurs?

Professor Alan Winfield: It cannot, I agree with you.

Q24 **Baroness Grender:** I am going to ask about the issue of bias. It was interesting that you used the phrase “potential to deceive” earlier on. We all believe we lack bias, but all research suggests that we are wrong. One of the questions that we want to ask is: how can artificial intelligence systems be developed so that they are not discriminatory or there is no bias in the algorithms built in by the individuals who have written it?

Dr Ing Konstantinos Karachalios: Of course, if we replace "bias" with our "human preferences", this is what makes us human. You can never be neutral; it is us. This is projected in what we do. It is projected in our engineering systems and algorithms and the data that we are producing. The question is how these preferences can become explicit, because if it can become explicit it is accountable and you can deal with it. If it is presented as a fact, it is dangerous; it is a bias and it is hidden under the table and you do not see it. It is the difficulty of making implicit things explicit. This is one of the most difficult things in life, and in politics.

We have a specific project that we started precisely on how we can make our algorithms, data and the interface explicit, to make clear the preferences and the bias. This is the best we can do. We cannot have bias-free anything, but at least it can be explicit. Our project, P7003, aims to define precisely what an unacceptable bias would be. The words must be explicit: how you can treat data, the collection of data and the quality of the data, and so on. This is very important, because otherwise people will be faced with decisions that will affect their lives and they will not know why they have affected their lives in a negative way and so on. It is time to do it, and we have started this.

Professor Alan Winfield: I would say, again as an engineer, that a biased AI is a badly designed AI. It is an AI that has been designed with uncurated data. More than that, as Konstantinos said, AIs reflect the unconscious biases of their human designers. It is critical, therefore, that design teams reflect the gender, age and ethnic mix of the societies that they hope to serve; otherwise, there will be inevitable unconscious biases.

Baroness Grender: Children cannot be part of the design team, for instance, and the Children’s Commissioner for England noted that the risk of children being excluded from the development of AI is where bias can creep in. Yet they are potentially the most exploited group, because they are so susceptible to everything. Let us talk very specifically about children and the bias issue. Just as an example, what do you do to overcome it?

Professor Alan Winfield: There is a very troubling example of one of these conversational AIs—a loudspeaker that you can speak to. In the US it was discovered that very young children were using this and were becoming bad mannered. The problem is that an AI system does not require a child to say, “Please can you tell me what’s on television?”, and, “Thank you”, afterwards. To me that suggests that the design team did not even have parents let alone child psychologists or teachers on the team. It would have been ever so easy technically to build into the chatbot a requirement for politeness—“please” and “thank you”,

essentially. You are quite right about young children. These were pre-school children, and they are too young to be part of that design team. Surely adult proxies should and must be part of the design team.

Baroness Grender: How do you bring that about? What do you need to do to ensure that happens?

Professor Alan Winfield: It is part of well-established frameworks, and responsible research and innovation would require this, such as the 2014 Rome declaration⁶ and the EPSRC AREA framework⁷.

The Chairman: You used the interesting expression “uncurated data”. We are going to look at the strength of data protection and so on in the context of AI applications. It is interesting that in the current Bill and the GDPR there is that point about automated decision-taking, but it is only automated decision-taking; it is not partly automated decision-taking. Do you think that makes any difference in the current context?

Dr Ing Konstantinos Karachalios: Of course. Fully automated decision-taking is very dangerous. There should always be supervision by humans as to the final outcome, if it is of critical importance to other humans. It is not only about the safety systems that guarantee the safety of physical systems. I do not know if you have heard what happened after the Las Vegas massacre, but there were news feeds on Google and Facebook with terrible news that was fake. They were feeding the public with fake news, because there was no journalist there to supervise them.

Professor Alan Winfield: Machine-generated news.

Dr Ing Konstantinos Karachalios: It can cause a lot of harm to society and to people. Just because there is a computer in the loop, they let it be. We should not accept it any more. We should say, “Stop it. You are unqualified for this”. This is very evident to me. We should not have any lenience. Just because the computers are in the loop, they do not deserve any asylum for this.

Q25 **Lord Hollick:** We come to the ownership and exploitation of data in the public domain. When the Royal Free NHS Foundation Trust entered into an arrangement with DeepMind to hand over 1.6 million patient records, it obviously sparked controversy about consent and privacy, but it also sparked a lively discussion about how that data should be exploited and who should get the benefit of that exploitation and over what period. Can you both please tell us the ideal structure of a deal to protect the public interest, and to ensure that the public interest is suitably rewarded in terms of data, information, apps and financially and in a way that would still encourage the digital companies to enter into those arrangements? What is the ideal deal?

⁶ Note by witness: European Commission, Rome Declaration on Responsible Research and Innovation in Europe, https://ec.europa.eu/research/swafs/pdf/rome_declaration_RRI_final_21_November.pdf [accessed 5 January 2018]

⁷ Note by witness: Engineering and Physical Sciences Research Council (EPSRC), Framework for Responsible Innovation, *Anticipate, reflect, engage and act* (AREA), <https://www.epsrc.ac.uk/research/framework/area/> [accessed 5 January 2018]

Dr Ing Konstantinos Karachalios: This is a huge question and probably one of the most important ones. I would say, though, before we go there, that we must reverse the current situation. The current situation is we have no agency over our personal data any more. This is a global problem. In the way we connect with the networks, we have lost any agency over digital identity and our data. This is a huge failure and a danger to democracy. This is not about money; it is very political. What can be done? There are practical solutions in how we structure our access to the networks so that we keep agency and we reveal and conceal what we want about ourselves and get some control over it. This is the first thing.

The second concerns the data that we have generated anyway. A lot of it is in the public interest and we should not privatise it. Data generated in the public domain can be useful for many purposes, such as helping clinicians to understand the behaviour of diseases. This can be of benefit to humanity and it should not be privatised. Of course, it must be done in a way that does not expose personal information.

This goes back to the first question I mentioned. This is an extremely complex aspect and has different levels of consequences. The highest level is the danger to democracy, because with no control of your identity you are a slave. You may be in a golden cage with a lot of high-tech gadgets and so on, but you are a slave. This is not the future; we are there, in my opinion.

Professor Alan Winfield: The problem is that we have all made a Faustian pact with the big AI companies. They give us really cool systems such as search engines, machine translation or social networks for free in exchange for our personal data. It transpires that data is extraordinarily valuable. What worries me, in addition to the question of what to do about it, is that many people simply do not understand the extent to which their own data is being used in this way.

Lord Hollick: Let us stick to medical records, because nobody who goes to see a doctor or to hospital expects their data to be made available. They do not give consent, but they may do when they use a search engine, so your argument does not apply there. How are the interests of the public to be protected?

Professor Alan Winfield: By having very strong privacy protections on personal medical data.

Lord Hollick: How do the Royal Free NHS Foundation Trust and similar bodies in the public domain get access to the skills in order to argue their side of the case to protect the public interest?

The Chairman: Are data trusts a solution?

Professor Alan Winfield: They may well be. To be perfectly honest, this is outside my area as an engineer, but I am on the ethics advisory board for the big EU Human Brain Project, and we take extraordinarily seriously the privacy of brain data from brain scans. I understand from colleagues that there are really great technologies for anonymising that data so that science can still get the benefit from that data without it being possible, as it were, to trace it back to an individual who was the source of that

data. I think we need to have data trusts, as the Lord Chairman suggested.

Lord Levene of Portsoken: Clearly, people give up the data, but have they given permission for it to be used? We are all faced with using something online and you have an eight-page document that you are supposed to read and click “yes” at the bottom of if you want it to happen. I think the number of people who read those is tiny, yet by doing that you are giving permission for you know not what. Has any thought been given as to how that can better be regularised?

Dr Ing Konstantinos Karachalios: I call it “broken consent”, and I alluded to it in my previous explanation here. We have started a whole series of standardisation projects in this effort that we are making, and it is no coincidence that more than half of them are about data—children’s data, student data, data gathered in the workplace and general privacy considerations. There are no easy answers to this, but I can tell you that hundreds of people right now are working on this. They do not have an agenda and they are not paid by any multinational or whatever. We are trying to reclaim this lost territory and do whatever we can. I think it is extremely important and I am ensuring the entire organisation is putting a lot of emphasis to try to address these issues.

Lord Hollick: There is a stark contrast between the public sector and the private sector, because in the private sector there are battles royal being fought about who owns the data and you are not allowed to use that data. I fail to see why the horse has bolted in the public sector but in the private sector this remains a very live issue of keeping the horse in the stable.

The Chairman: I am sure we are going return to the whole data issue later in our inquiry, but I would like to come on to Lord Swinfen now.

Q26 **Lord Swinfen:** Many respondents to our call for evidence believed that the growing use of artificial intelligence was likely to exacerbate existing economic disparities. Do we need to address the potential economic disparities that could be caused by widespread use of this new system, and, if so, how?

Professor Alan Winfield: I absolutely believe that we do, yes. It is vital that not only the benefits of AI should be shared by all in society but the wealth created by AI. That is really important. I often tell people not to forget, of course, that the basic technologies were all funded by taxpayers, so in a sense there should be a premium back to the taxpayer. It would not have happened without that initial investment in military research, if it is the internet, or civilian research. The world wide web was funded by European taxpayers and basic AI technologies were all developed in universities, funded by taxpayers. It seems to me perfectly right and proper that we should all share the wealth. How do you do that? First, of course, the big AI companies should pay their taxes, but we also need to have innovative wealth distribution systems. We should be looking at things such as a universal basic income or a universal conditional income or a negative income tax. I am not an economist, but one thing is clear to me: something has gone very wrong in the advanced economies when, for some of decades, we have had increased

productivity but wage stagnation. This can only be because of automation.

Dr Ing Konstantinos Karachalios: I agree. I would start with the famous curve where productivity has grown and middle incomes have stagnated since the 1970s, since the onset of ICT technologies. They are not fulfilling their promise for more equality or more progress in society. More and more wealth is accumulated by fewer and fewer people, and not only wealth but power. These technologies that we are talking about may even accelerate the pace, and this is not a place we want to be. A British artist, James Bridle, made a nice exposition called “car trap”. He says that this technological progress may take us to places we do not want to go to as a society. We may lose skills, jobs and personal autonomy. We may lose—this is also very important and goes back to our intelligence and cognisance—the capacity to make sense of the world. We do not want to go there, which means that we cannot just leave the technology to rage. We must put a framework around this. It is not the regulation of technology production itself but the outcome and the distribution that needs to be regulated, and I think this should be done.

The Chairman: You have certainly stiffened our backbone generally in the course of this afternoon.

Q27 **Lord St John of Bletso:** We have read the recent report by Professor Dame Wendy Hall and Jérôme Pesenti on growing the artificial intelligence industry in the UK and the general recommendation that there is no need for regulation. Dr Karachalios, you spoke about IEEE P7000 and IEEE P7001 and a range of international standards on ethical developments. My question to you both is: do you believe that the ethical development and use of AI requires regulation? If so, what type of regulation?

Professor Alan Winfield: First, we need standards, and I think we need standards more than we need regulation. Standards are the hidden infrastructure of the modern world. Standards are as important as roads and telephones. They are part of the infrastructure of modern life. There is no doubt in my mind that some standards will need teeth, particularly if, and I am sorry to keep going back to the phrase I used, they are safety critical. If the AI is part of a safety-critical system that has the potential to cause serious harm, those standards need to be mandated, so there needs to be regulation. I am not saying that we need to regulate everything, but we need to be selectively regulating particularly the safety-critical systems that we have already mentioned.

The Chairman: But not the ethical principles?

Professor Alan Winfield: The ethical principles underpin the standards, some of which are mandated. I see that there is a flow. In fact, I submitted written evidence to the Commons Select Committee last year and I made exactly this point that ethics underpin standards, some of

which become mandated in regulation⁸. I think that is the way that ethics flows through to regulation.

Dr Ing Konstantinos Karachalios: I would make three very brief points. The first is that we should mandate human oversight of critical systems and of systems that influence public opinion and so on. I gave an example before. These should be mandated. We cannot say that just because there is a computer in the loop you can do whatever you want. This is what happens now and it is apparent that this does not work.

The second point is self-regulation. This is what we are doing. This should happen at the origin of system design and construction. The architects of the systems must start taking the contextual aspects of technology into account and not just speed to the market and functionality. The technology must protect our dignity by design, not as an afterthought. I think the time of innocence is over. We should say “enough is enough”.

The third point is that if you combine this, regulation would promote rather than hinder innovation, because a lot of these technologies would not take off because we would not trust them, and we would have reason not to trust them. If you regulate in such a way that you infuse trust and make it trustworthy, the technology will take off. It is to the benefit of everybody to have a kind of regulation that will help technology take off because people will trust it.

Lord St John of Bletso: I am a bit concerned that I might run short of time. I just looked at the ethics and legal and data capital and there were two approaches: the soft approach and the hard approach. Obviously we have the hard approach with the GDPR, and I was going to go on to the soft approach.

Dr Ing Konstantinos Karachalios: We are making the soft approach before the hard approach comes.

Lord Swinfen: Who is to draw up these standards, bearing in mind that the Civil Service is probably not in the forefront of artificial intelligence?

Professor Alan Winfield: I think it is primarily the role of professional bodies, of which the IEEE is one, and there are other standards bodies, such as the International Organization for Standardization and of course the British Standards Institution in the UK. These are the bodies that should be responsible for developing these standards, but, essentially, it is within the community, as Konstantinos said.

Dr Ing Konstantinos Karachalios: Everybody should do a job at a different level. You have a responsibility to mandate something. We have a responsibility to do a better job at the beginning. These are different types of standards and we should not confuse them. Ours are bottom up and voluntary, but they can be very powerful. Wi-Fi is voluntary, but everybody uses it because it is useful. Our ambition is to produce standards that will be used by the technical community, because they make sense and this will make your life much easier.

Viscount Ridley: Dr Karachalios, you said that good regulation can

⁸ Note by witness: Written evidence from Professor Alan Winfield ([ROB0070](#)) received by the House of Commons Science and Technology Committee inquiry on Robotics and Artificial Intelligence

stimulate intervention, but you would surely concede that bad regulation can do the opposite. There is a moral hazard and an opportunity cost if you get it wrong and it could end up stifling innovation and preventing good things from happening.

Dr Ing Konstantinos Karachalios: Of course.

Q28 **The Chairman:** One final sentence from each of you, and this is probably a very unfair question: if there was one recommendation you would like to see the Committee make at the end of this inquiry, what would it be?

Dr Ing Konstantinos Karachalios: I speak in the country of Lord Byron and Winston Churchill and so on and coming here is very inspiring. I would say: give the kids digital literacy but do not confuse it with literacy per se. Give your kids literacy and give them the means to make sense of the world. Without this, digital literacy means nothing.

Professor Alan Winfield: I would go back to my plea for ethical governance. We need a body, a kind of AI watchdog that essentially is responsible for finding that balance, as you said, which I agree is not easy, but we need a body that is responsible for making those recommendations to government.

The Chairman: Do you mean to build on Dame Wendy Hall's AI council or something of that sort?

Professor Alan Winfield: Indeed, yes.

The Chairman: Thank you very much indeed. We have had a very stimulating session. Clearly the IEEE is putting itself front and centre in this debate and I think that is extremely welcome. Thank you very much indeed.

Dr Ing Konstantinos Karachalios: Either as a Committee or personally, if you want to continue this discussion—and there are many things that we have opened up but we have not been able to explore in depth—we are available.

The Chairman: Thank you very much.

Professor Rosemary Luckin, Miles Berry and Graham Brown-Martin – Oral evidence (QQ 181–189)

Professor Rosemary Luckin, Miles Berry and Graham Brown-Martin – Oral evidence (QQ 181–189)

[Transcript to be found under Miles Berry](#)

Major Kitty McKendrick, Professor Noel Sharkey, Mike Stone and Thales Group – Oral evidence (QQ 153–162)

Major Kitty McKendrick, Professor Noel Sharkey, Mike Stone and Thales Group – Oral evidence (QQ 153–162)

Evidence Session No. 16

Heard in Public

Questions 153–162

Tuesday 28 November 2017

[Watch the meeting](#)

Members present: Lord Clement-Jones (The Chairman); Baroness Bakewell; Lord Giddens; Baroness Grender; Lord Hollick; Lord Holmes of Richmond; Lord Levene of Portsoken; The Lord Bishop of Oxford; Lord Puttnam; Viscount Ridley; Baroness Rock; Lord St John of Bletso; Lord Swinfen.

Examination of witnesses

Professor Noel Sharkey, Major Kitty McKendrick, Dr Alvin Wilby and Mike Stone.

Q153 **The Chairman:** Good afternoon and welcome to our witnesses for our 16th formal session of the inquiry. This session is intended to help the Committee discuss the military use of artificial intelligence. We have a very strong panel this afternoon. We have four panel members and I can see they will be fighting to get a word in edgewise. We have Professor Noel Sharkey, emeritus professor of artificial intelligence and robotics and public engagement at the University of Sheffield. Fresh from the moral maze, we have Major Kitty McKendrick, who is a visiting fellow at Chatham House at the Royal Institute of International Affairs. We have Mike Stone, who is the former chief information officer at the Ministry of Defence, and we have Dr Alvin Wilby of the Thales Group: vice-president for research, technical and innovation. A very warm welcome to you.

As always with these things, I have a little rubric to read out first before asking you to introduce yourselves. The session is open to the public. A webcast of the session goes out live, as is, and is subsequently accessible via the parliamentary website. A verbatim transcript will be taken of your evidence. This will be put on the parliamentary website. A few days after this evidence session, you will be sent a copy of the transcript to check for accuracy and we would be grateful if you could advise us of any corrections as quickly as possible. If, after the session, you wish to clarify or amplify any points made during your evidence or have any additional points to make, you are very welcome to submit supplementary or written evidence to us. Perhaps, first, you would like to introduce yourselves for the record. Mr Stone.

Mike Stone: I have embarked on a third career now as an adviser, but I spent 28 years in the military first and retired as a one-star around the turn of the century, and went into business. Subsequently, I was

seconded back into the MoD, initially as the chief executive of the Defence Business Services—all the back office of defence—and then as the chief digital and information officer for the last three years. I left that post earlier on in the year, at the end of March.

Major Kitty McKendrick: I am a visiting fellow at Chatham House where I am on placement from the Army where I am still a full-time serving officer. I have served 12 years as an officer in the Royal Electrical and Mechanical Engineers, deployed on operational tours to Iraq and Afghanistan, and I am a chartered engineer. At Chatham House my research is on artificial intelligence generally and the use of predictive algorithms in support of public safety. It is probably worth noting that I have not been directed to study that particular area by defence but have chosen it and, although I am here and obviously still a serving member of the Ministry of Defence, I am not representing it in an official capacity today.

Dr Alvin Wilby: I am the vice-president for research, technology and innovation at Thales UK. I am a systems engineer by background with a maths and physics education. I have worked in defence and aerospace for about 35 years now. I am a fellow of the Royal Aeronautical Society, the Institution of Engineering and Technology and the Royal Institute of Navigation. I am not an AI expert but I have experience with defence systems that need faster-than-human response times and have worked on safety-critical software systems for aviation. My company is a leading supplier of autonomous systems for both civil and military applications. If you came here on the Jubilee line, it would have been our signalling system that was controlling the train. For the Army we provide the Watchkeeper remotely piloted drone.

Professor Noel Sharkey: I am an emeritus professor of robotics and artificial intelligence at the University of Sheffield. I have been working in this and related areas for almost 40 years, particularly in the area of machine learning which has use for many things from robot control to the diagnosis of diesel engines and safety-critical systems. In 2007 I found out about autonomous weapons systems being developed in the United States and I have been working to get new regulations for these ever since, particularly at the UN now, with the CCW, a committee that prohibits weapons systems. We have been working there for five years now to get a new, international, legally-binding treaty to prevent these weapons from being used. Finally, I am co-director of the Foundation for Responsible Robotics, a group based in The Hague. We look at every other application apart from military, so all ethical, responsible issues in AI and robotics.

Q154 **The Chairman:** Thank you very much. We obviously have quite a large panel today and I am very anxious to hear the answers to the questions that you are particularly expert in answering as individuals. It is not necessary to answer at length on every question, so a nod or a wink will be sufficient as we go through. Obviously, we have quite a bit of ground to cover. I am going to ask a very general question, which I am sure you will all want to answer as we go through. To what extent is AI already being used in military applications? Where is its use currently most

prevalent? Which areas are likely to see the most growth in the near future, and how might this change the nature of warfare? You can make the distinction between the logistical use of AI and the military use of AI. That might be quite useful in considering the way forward.

Mike Stone: In the first instance, it is most prevalent in the back office, at the moment, rather than in the military arena. Intelligent process automation is beginning to come into its own. There is a degree of AI within the cyber arena. I was looking at how I could gamify all the networks in the MoD and create decision-support tools, essentially. A big part of AI is about improving evidence-based decision-making. There is obviously a use in terms of safety-critical. Then we have use in unmanned vehicles. We need to make a distinction between autonomous and unmanned. A drone is unmanned. It is driven by a person. The decision to fire is made by a person. It is an unmanned vehicle as opposed to autonomous or semi-autonomous. The only application I am alive to that is entirely autonomous is the Type 45 defensive ship-to-air missiles, where they need to be able to operate at speeds at which humans could not think. Of course, they have to be switched on by humans and switched off subsequently by humans, but if you had a human in the loop on that, as you were going, it would put the ship at peril. Otherwise, I am not alive to the use of autonomous vehicles.

The Chairman: The use of vocabulary is rather important in this sense, so that we know exactly what we are talking about.

Major Kitty McKendrick: It is a very broad question. There certainly is a tendency to look straight to autonomy, but the use of direct and indirect artificial intelligence in military applications is probably a lot broader than that. Apart from skills-based functions that are automated for the reasons of speed of human control, there is also optimisation of logistics which uses artificial intelligence and more knowledge-based functions as well as data analytics for intelligence. I would say the most sophisticated artificial intelligence in use in direct military applications is probably for modelling a simulation for disaster response, staff training and wargaming military plans. It is worth noting that civilian development in artificial intelligence, particularly the use of machine learning and deep neural networks, from everything I have seen from my research, which is predominantly open-source, outstrips its use by the military at this stage.

The Chairman: Is there any useful distinction between autonomy of systems and intelligence of systems?

Major Kitty McKendrick: A distinction has to be made. While the level of autonomy of a system might be linked to the exact type of artificial intelligence that is underpinning it, that is not an adequate definition of autonomy and one that will not withstand technical development. It is important to separate the terms “artificial intelligence use” and “autonomy”.

Dr Alvin Wilby: There are some very good points made. One of the challenges is that we do not have precise definitions of what we mean by “AI” or “autonomy”. That tends to get us tied up in knots when we talk about this. To illustrate, some of the research areas we are looking at that would fall under the definition of “AI”, if I follow the Engineering and

Physical Sciences Research Council definition, are things such as self-protection systems. As Mike says, you need to react faster than a human, so self-protection systems for helicopters and fixed-wing systems would be in that category. More than self-protection systems, they are capable of being dangerous. If you fire flares on the runway, that is extremely hazardous, so there is a safety implication in that, but it is definitely an autonomous system.

We are looking at an autonomous mine-hunting system, using autonomous surface and underwater vessels, for example, but also using convolutional neural networks, which Kitty mentioned, to support armoured vehicles, for example. You can use the camera images and spot threats more quickly. There is a wide range of potential applications and I think that is the challenge.

In terms of systems that are capable of autonomously selecting targets, I am aware of only a few. Again, I have no specific insight, just looking on the open web, but the Israelis have a system called Harpy, which is an anti-radar system. It is capable of loitering over an area and deciding which target to go for. The Russians have some anti-submarine mines that, again, can do autonomous target detection. At that end of the AI spectrum there are very few systems. As far as I am aware, the UK is not developing systems like that at the moment.

The Chairman: Do you find all these distinctions useful or are your concerns spread right across the spectrum?

Professor Noel Sharkey: My concerns, in terms of military use, are just two functions, in fact: autonomy of target selection and the application of violent force. Everything else, as far as I am concerned, is fine; I have no problem with it. They have covered most of the things. I am also seeing quite a lot of research on strategic research, so what to do in the battlefield—risk assessment—and some very good work on trying to work out where insurgents will strike next: picking up patterns of activity, feeding them into neural networks and seeing if they can make a good prediction. Some of it seems to be coming along quite well. I am seeing things that have not mentioned, which are military uses for cleaners for the hulls of ships when they pull away from the Navy. The United States is starting to use this kind of thing.

What I am not seeing yet, unfortunately, is AI used in bomb disposal. We do not have autonomous bomb disposal, which is one of the things that kills most of our young soldiers. We could do with a few billion pounds being invested in that and less in offensive weapons, as far as I am concerned. We could save a lot of lives. They could go ahead of the convoy and look for explosives. It is not an easy task because the place is full of explosive fragments. That would be quite easy to see.

The other comment is: what are you calling AI? Most of what I see are computational systems. I would not call the Harpy AI. It detects a radar signal, looks up a database and says, "Is that our signal"? If it is not, it dive-bombs the target. The term AI is running out of control, at the minute. You are seeing it everywhere because there is massive commercialisation. Despite the fact that it looks as if there is an exponential increase, there is not. You are seeing the same techniques

that I was using in the 1980s, but you have very fast machines—that is creating a difference—and very large amounts of data. The thing is spreading sideways through the commercial world.

The Chairman: We should not assume that, in a sense, things are more developed than they are

Professor Noel Sharkey: Exactly.

Lord St John of Bletso: You have partly answered my question. My understanding is that in GCHQ they organise themselves into the attack group and the defence group. How effective has AI modelling been in the defence group? Major McKendrick, you have answered that in part. On the other side, to what degree will robotics and drones play a part in the future of the attack group?

Major Kitty McKendrick: I am not sure what the answer is to the first part of your question. I am also not even sure how I answered it when I spoke—robotics in the function of an attack?

Lord St John of Bletso: I am talking about the defence side of things, initially.

Mike Stone: I think it was me who said that. I was seeking to use gaming techniques to model my networks, to therefore be able to play through thousands of different courses of action as a result of a particular incident, and decide the optimal route.

The Chairman: That is in real time. That is not in training.

Mike Stone: Very, very near real time. If you feed it real time information about the infrastructure and the network paths and you either simulate an incident or have an incident take place, you can work out the optimal solution to that. That is a highly computationally intensive approach though. What can be done from an offensive perspective, we have always had jamming capabilities. An extension of that, clearly, is an opportunity. I am not equipped to talk about the offensive side.

Q155 **Lord Hollick:** Can we come to the issue of human control over these systems? To what extent should direct human control over the deployment and targeting of weapon systems be required? When restrictions are placed on the development and use of AI-enabled weaponry which can target autonomously, how is it practical to put a human decision-maker in that process? You have been clear that there is not much weaponry yet deployed which has artificial intelligence. As you said, once missile systems have been turned on, they have to act very quickly. If we look ahead a little, autonomous identification of targets is just round the corner. How does the human element in the chain stop that? How do we prevent those systems from taking leave of the human control?

The Chairman: Professor Sharkey, we will start with you. You have quite strong views in this area, I believe.

Professor Noel Sharkey: I think I have reasonably logical views in this area rather than strong. This is the key question going on at the UN at the moment. The ICRC—International Committee for the Red Cross—has

taken up this idea about human control, and what is effective human control under international humanitarian law and international human rights law. If there is a human properly in the loop, it is not an autonomous weapon system. If there is autonomous targeting, there is no human there.

It is more than that. You have to press, for instance, the idea that the UK will say there will always be human oversight. The UK say they will never use an autonomous weapon system. That would make me proud except for one thing, which is what they define as an autonomous weapon system. They set the bar so high compared to the rest of the planet. They are out of step with the rest of the planet. I find this embarrassing. They say an autonomous weapon system will have to be aware and show intention. Then they say, "This is unlikely to happen, so we will never have autonomous weapon systems, and the fact that we will never have them means that we will never support a ban for them as well". I do not get the logic of this, but that is our Ministry of Defence for you.

When pressed on what human oversight is, I cannot get an answer, so we have "on the loop", "in the loop", "around the loop" and "wider loop". In the United States we have, "appropriate levels of human judgment". At the moment, we are pressing; we have two weeks next year at the UN where they are going to have to come and explain what they mean by this.

Lord Hollick: In your view, how should that be defined?

Professor Noel Sharkey: I am an academic and have written several papers on human control. There is a scientific research area called human supervisory control, which has been going for about 30 years, and you look at different ways of control. The lowest level would be pushing a button and letting the weapon go, as you would with a missile. Do not forget, you have already selected the target. More or less, our Brimstone can veer a bit and look for the ship it was sent for. Generally speaking, that is one method. Another method is the Patriot missile system where you have six seconds to veto. That ended up with a British Tornado shot down by two US fighter jets because when you have six seconds to veto, you veto and your fellow soldiers are killed, you are not very popular—put it that way. It is a panic station. Another way is that the machine gives you the target and you decide whether to press the button or not. One of the problems there from supervisory control is that that leads to an automation bias, and there have been lots of experiments even with Tomahawk missiles where, after a while, when it gets it right a number of times, you start obeying what the machine tells you. You have to develop an interface that stops that.

Another method, and this is almost the last method, is that the machine gives you a list of targets and you choose one from that list. That is not good if you feel, "I must choose one from that list". It is also not good if the machine rank-orders the target, because experiments have shown that you will take the top one. I am sorry, I am using some psychology here. In psychology, we separate between automatic reasoning and deliberative reasoning. Automatic reasoning is what we do when we are playing tennis, riding a bicycle and—most of the time—driving our car.

Deliberative reasoning is when you engage with what you are doing and think about it. That is what we need. The final level is proper, full engagement about what the target is. Make sure it is a legitimate target, it is a military target, it is not a surrendering soldier—any of those things—and then, for every attack, initiate the attack. The last one is a little idealistic, but in the last two or three we can find a way of doing it.

Major Kitty McKendrick: I hope I can elucidate slightly. Compliance with international humanitarian law—the principles of distinction, proportionality and military necessity or precautions in attack—is the absolute minimum standard for the use of a weapon system. From a military point of view, all weapon systems are simply tools we use and are the responsibility of the operators and commanders. Irrespective of the level of automation or if that tool is autonomous, depending on how you define it, it is always the responsibility of the operator and the commander that that weapon system is used in accordance with the relevant international humanitarian law. Perhaps confusion arises externally because whether you require a human “in the loop”, “on the loop” or “out of the loop” entirely depends on the type of weapon and the context. For anti-ballistic missile defences you would not consider putting a human in the loop. It would be reasonable and in accordance with international humanitarian law to use those weapon systems with the human out of the loop. Certainly, for other weapon systems, such as remote weapon systems that identify and perhaps prioritise targets, a human in the loop is required to make a firing decision for them to be used in accordance with international humanitarian law.

Mike Stone: I absolutely agree with everything that Kitty has said. The question demonstrates the need for a very clear lexicon. AI is such a broad area. When I was studying years ago, we were talking about AI. As the professor said, rules-based was considered to be AI. That has been in military applications for a long time. That is a simple, “If this, then that occurs”. We have had neural networks over the years. These days we are talking about machine learning, peer to peer, the insights from big data—all terms we need to have clarity over. Then we have this issue of autonomous, semi-autonomous and whether a drone, even if it is delivering a missile, is semi-autonomous. I would say it is an unmanned plane, in the same way as a subsea vehicle can be driven from onshore and a human has to make a decision about it. That is very different. It is a bit like a plane. These days you do not need a pilot; the moment they are looking at the screen it is, effectively, flying itself. In terms of drones, we are talking about unmanned rather than autonomous vehicles.

Dr Alvin Wilby: I want to pick up one of the points, which was this issue of when you use this and when you design and develop it, which I think I heard in your question. I fully agree with everything that Kitty has said about how you should view the law relating to armed conflict, and so on. We are talking here about systems that will, in a sense, make judgments about the circumstances they are in. I want to be clear when I say “judgment”, because one of the problems is we talk in language that imbues these systems with sentience; we think they are thinking and making decisions for themselves. That is a long off. We are a very, very

long way from anything that would pass a Turing test, even in the narrowest way.

Some of the judgment, therefore, is what we build into it when we develop it. If we tell it that “a tank looks like this and a civilian looks like that”, that is something we have built in during the development process. The issue for me is not about whether it then decides to do something rogue; the issue is whether we have to guard against what I call artificial stupidity, which is when the algorithm comes up with an unexpected result that would cause a problem. There is a need to think about how you bound the behaviour of these learning systems so you can have adequate confidence, if you are making a target decision that might be based on it, that you will get the outcome that you expect. In a sense, that is no different from designing any safety-critical system, but it is a bit harder than for run-of-the-mill safety-critical systems.

Viscount Ridley: Last week we were talking to the medical profession and there was quite a lot of emphasis on the idea that artificial intelligence is there to augment, not replace, human beings; that it is there to help the radiologist make better use of his time, or be better at his job. Is it possible to maintain this distinction in this field? Can you argue that is still what AI is likely to do in defence, or not?

Professor Noel Sharkey: I would say so, yes. AI could be of great assistance. Let me say something about compliance with IHL. This is mainly done through Article 36 weapons review, and this would be very, very difficult with autonomous weapons because of their predictability. We are going to have to write something in to Article 36; we are getting a moving consensus that there should be something in there about human control.

Mostly we hear about smart weapons. I talk to the military a lot and the manufacturers. Frankly, smart weapons save civilian lives, but that is not the main goal of them. The main goal is to save ammunition and hit the target accurately. If we move more towards a world where we are thinking, “Let’s get zero civilian casualties; let’s use artificial intelligence for that purpose”, it could be of great assistance, not in control but bringing information to us about exactly what is happening on the battlefield, assuring us that someone is not surrounded by civilians. That would be very good. Yes, I am for that.

The Chairman: Is there something the others want to add to that?

Major Kitty McKendrick: I would second that view and take the discussion away slightly from weapons specifically. In my opinion, the low-hanging fruit for military use of artificial intelligence is collaboration in the planning process. It can make the military more able to make good, evidence-based decisions in acceptable timescales, deal with complexity and plan iteratively to operate more effectively.

Professor Noel Sharkey: I should have finished with: on the proviso of the other things I said about human control, that you are not going for automation bias. The same goes for the medical profession—not that we accept things and do it.

The Lord Bishop of Oxford: It has been an interesting exchange. I

want to register some concern that the UK is so out of step with other nations. I would like to draw you out a little more on the ethical foundations for the arguments you are putting forward about autonomy. You have made a lot of distinctions between different degrees of autonomy but you have not said much about why it is important to make that distinction and why this may be a critical and bad thing. That may seem obvious to you, but it will be helpful to know on what ethical basis you are basing those positions.

The Chairman: Lord Bishop, we have the question about ethical principles later on. I am wondering whether we should not subsume and park your supplementary for later.

The Lord Bishop of Oxford: By all means.

Q156 **Viscount Ridley:** To your knowledge, is the UK currently developing autonomous or semi-autonomous weapons? Are there any formal restrictions on developing this weaponry in the UK? Should the UK be considering restricting the purchase or procurement of particular kinds of AI-enabled weaponry? Should we consider restricting the sale of such things? Against the background of this, will any of this matter if there is an international arms race developing in this area?

Mike Stone: I think I answered it, in part, in my response to the first question.

The Chairman: We are very good at repeating ourselves.

Mike Stone: As far as I am aware, the only time where we have autonomous weapons is with these defensive elements, such as the Type 45, where the speed of response required is far greater than humans can provide. Perhaps I could come back on it after the others have responded.

Major Kitty McKendrick: On the question raised of the possibility of an arms race, with regards to specific autonomous technologies, and whether or not prohibitions on development might stand to stop that, in preventing that escalation of technological development, a lot of the tendency towards that escalation will be on the military utility of the weapon. If a weapon is not in control and not subject to appropriate human judgment, I fail to see how it could be militarily or operationally useful. For that reason, the limits on utility of weapons that are out of control might prevent escalation of their development, certainly by defence sectors.

Viscount Ridley: An example we have been given is unrestricted submarine warfare. The Americans said they would never do that, but within six hours of Pearl Harbor happening they changed their minds. You can see why, but they did. Do these prohibitions ever matter once things start getting out of control?

Mike Stone: The genie is out of the bottle, particularly in the civilian world. Controlling that will be very difficult. In the past, a lot of the innovation and investments came from the military world; today, much of the innovation is coming from the civilian world and being brought into the military world. The likelihood is that even if we were to act in a

different way others would not. Therefore, there is going to be a degree of a race to get to a point where there are capabilities.

Viscount Ridley: Should we be much stricter on, for example, the sale of something that looks like an autonomous or intelligent weapon to any old regime around the world?

Mike Stone: I think that would be very appropriate, yes.

Dr Alvin Wilby: As Mike said, and it features in some of the quite impressive videos that your organisation has developed around the use of killer drones, and so on, that technology is in the civil world. It did not get there from defence companies or contractors; it did not get there as a result of being exported by us or being stolen by terrorists, which you see in some other scenarios. It is out there and it has come from the civil world. The issue is how we protect ourselves against that. In terms of the legal framework in the UK, if the UK were deliberately developing systems that were intended to act indiscriminately, that would clearly be illegal under current international law.

Professor Noel Sharkey: As has been pointed out, we have a range of SARMO weapons in the UK—that is sense and react to military objects. For shooting down mortar shells or missiles, that is all fine. Some people might call them automatic; it is a funny distinction. We have been developing other things as well. Our main defence contractor, BAE Systems, has a fully autonomous Taranis fighter jet that is weapons based, but there are no weapons in there yet. That has been widely tested in Australia—finding targets in teams even, so that they can co-ordinate for the targets. BAE Systems is currently also making the Ironclad, a small tank, and autonomous defensive drones coming out of it as well. We are moving there. We are a long way behind the United States, which has fighter jets, submarines and tanks, et cetera. Russia is pushing very hard on the Armata T-14 super tank, which is 10 years ahead of anybody else in the world. They are trying to make it autonomous as soon as possible.

This technology is out there and we are already seeing the beginnings of an arms race, which is why it is urgent to try to prevent this happening. Once the genie is out of the bottle, you have had it. I have spoken to people in the United States a lot. Operationally, they have big hopes for this because they have lost their edge on missile technology. They are using terms such as “mission completion”. If you are using drones and are fighting a sophisticated enemy, the first thing they are going to do is jam all your signals. These things can complete missions. The big area is swarming. Everybody talks about swarms because you need force multiplication: one soldier, many weapons. That is another operational one. I have done some work with NATO, and their thing was “area denial”. That is what they wanted to use them for. You block off an area. There has been a lot of talk about them happening strategically.

Q157 **Lord Hollick:** You have talked about an arms race, and there is no international agreement to prevent such an arms race in this area. Surely it is incumbent upon countries such as the UK and companies such as BAE Systems and Thales to develop defensive systems. Those defensive systems would no doubt be able to replicate the same abilities and

functionality of the attack systems. In those circumstances, they would be able, at the flick of a switch, to behave autonomously. Is not the corollary of the arms race you have described, and all the potential which bad guys will want to pursue, going to push us into a position where we have to develop a satisfactory defence, so we cannot be caught napping, either by nation states or by rogue terrorist organisations?

The Chairman: I am going to ask Lord Levene to ask his question and we will put both questions to the panel, if you could make a note of the first one.

Lord Levene of Portsoken: Interestingly, we have heard this afternoon that civilian applications are leading the military rather than the other way round. That being the case, as there is so much work being done now on autonomous cars, which presumably are fed with certain criteria as to how they should manoeuvre themselves, presumably it is not a very big leap to go to autonomous weapon systems, which are also fed with certain criteria and, therefore, make up their own minds whether they are going to activate themselves. Is that right?

The Chairman: Who would like to kick off with Lord Hollick's question?

Dr Alvin Wilby: I will make a couple of—

Professor Noel Sharkey: Could you repeat the last part of that question? I want to be very clear.

Lord Levene of Portsoken: As autonomous cars are being rapidly developed now, which presumably are fed with criteria as to how they should drive, presumably it is not a big leap from there to autonomous weapon systems that sense a certain scenario and set themselves off according to the way they have been programmed. Is that right?

The Chairman: Let us hear from Dr Wilby. He was just waxing lyrical in answer to Lord Hollick.

Dr Alvin Wilby: I am not sure that I can guarantee waxing lyrical, but I will try. I was going to pick up on the autonomous car scenario and make another point which I think is relevant to the cyber challenge here. If an autonomous car is reprogrammed to kill pedestrians, it has become a lethal autonomous weapon system. That is a credible terrorist threat. We should be taking cyber defences, many of which will need AI to be effective against those sorts of things, quite seriously.

Mike Stone: To take Lord Hollick's point first, we have to respond. If we did not respond anyway there would be developments in the civilian world that we would be able to take hold of. The genie is truly out of the bottle in terms of the development of capabilities that are driven by what is now referred to as AI. I say again that we need to have a tight lexicon on the meaning. Going back to what Professor Sharkey said, linking into what Lord Levene was asking, most of the things he described I do not believe are autonomous; they are unmanned. So long as they are directed or piloted remotely and decisions are taken, they are not autonomous. They are only autonomous when they are released and told to go off and find their own way in an area, identify targets and destroy them, and there is no human in the loop whatsoever. I absolutely take

Alvin's point that we need to be tackling the cyber aspects of all this for civilian aspects of life, let alone anything else.

Major Kitty McKendrick: I would agree we are still suffering somewhat from a clear definition of "autonomy". Certainly some of those capabilities mentioned, some weapon systems we have in service now, are described as autonomous but are not genuinely so, as in the example of the self-driving car that decides it wants to do something. Even the critical functions lexicon, which we say can select and engage a target, is problematic because quite often we use that to describe machines that are detecting targets in accordance with a predictable program. They have been told to look for certain features and they have identified those targets on that basis, which again I would not consider to be a genuinely autonomous system.

The Chairman: Did you want to add anything?

Professor Noel Sharkey: Yes, I certainly do. The challenge to what I said were not autonomous is ridiculous, if you do not mind me saying so. It is, perhaps, a bit disrespectful to say that and I should say that it was not as accurate as it might have been. The weapons I talked about, for instance the X-47B, which is the American fighter jet, will take off—

The Chairman: We like to provoke disagreement among our panel, by the way.

Professor Noel Sharkey: It will take off and land on aircraft carriers completely on its own. The Russian Kalashnikov neural network targeting system is certainly fully autonomous, so is what BAE Systems is calling an autonomous aircraft—the Taranis—and if they were here they would say that themselves. All these things are being touted as autonomous. Mind you, they said the Guardian in Israel was fully autonomous and had an autonomous targeting system until our campaign happened to pop up and take it to the UN and now, of course, there are humans in control of the weapons on board. There is a lot of slime in there, in a sense, if you see what I mean, and back stepping.

In terms of the beginning of an arms race, which you were asking about, yes, we could respond by developing further weapons and they could respond by developing counterweapons, and so on. One of the big reasons why people want autonomy is they say that the battlefield is getting too fast for humans. I would say there is no rush to kill people; let us slow down a bit. The idea is that these things can move faster and faster and faster. I do not want to live in a world where a war can happen in a few seconds accidentally and a lot of people die before anybody stops it. The best way to stop this arms race is to put new international restraints on it, to have a new international legal document. It will not be completely satisfactory. We went through all this on aerial bombardment all the way through the 1920s and 1930s and now the norms have shifted and it is perfectly all right to use.

The Chairman: I am going to stop you there because the next question is highly relevant.

Q158 **Lord Swinfen:** How realistic is it that lethal autonomous or semi-autonomous AI applications will be used by non-state actors or rogue

states? What might be done to prevent this? How easy is it to convert civilian AI applications to military use? How might AI be used to conduct cyberwarfare? Is the alleged use of chatbots to sway public opinion in a number of western democracies a sign that this is already happening?

The Chairman: They are all a bit unrelated so do not feel the urge to answer every question.

Professor Noel Sharkey: What was the end of the chatbots bit? I did not quite catch the end of the chatbots.

Lord Swinfen: Is the chatbot being used to sway public opinion in a number of western democracies?

The Chairman: There is a sign that it is already happening with AI applications by rogue actors.

Dr Alvin Wilby: I will not take all the questions. Your first question was how likely it is that things we would label as having AI would be used against us by rogue states or non-state actors. I would say that was an absolute certainty in the very near future. An example of a drone attack in 2015 was in Japan, where someone used a drone to try to deliver radioactive sand to the Prime Minister's office, which I think we would agree is an undesirable use of civilian drones. As we have said a number of times, the genie is out of the bottle and this capability is out there. The technological challenge of scaling it up to swarms and such things, again, does not need any huge inventive step; it is just a question of resources, time and scale. That is an absolute certainty that we should worry about.

You mentioned cyberattack. It has not come up in the conversation so far but the Stuxnet example is one that is perhaps worth thinking about. Again, that was a cyberattack that showed some elements of what you might call artificial intelligence. It was an example of something that was released into the wild. It was given the targets it was after—the centrifuges—and it was designed to cause real physical damage, not just to disrupt ISs. The degree of control exercised over it after it was released, I would argue, is perhaps not meeting the threshold we might wish. Again, those sorts of things are out there.

Mike Stone: It is absolutely inevitable that this is going to get into the hands of non-state actors and certainly rogue states. North Korea and Iran are, perhaps, top of the list in most people's thoughts around that. We have to be very careful about offensive cyber because you end up reaping a reward for the use of indiscriminately offensive cyber, and end up having to do quite a bit of defending. One has to be very careful about the application of that. As I said, I am not in a position to talk about offensive cyber. Chatbots are little programs which talk to each other. We are allowing them to be virtual agents that we can communicate with, but they do not have to be talking in the sense of chatting; they can chat to each other rather than chatting to humans. In life, everything is now mobile and social media driven; we expect to have service on demand. Much of that is driven by chatbots. I do not see that as having any impact on the way people view the military world.

Professor Noel Sharkey: I will try to be very brief on this. Rogue states, ISIS, IS or Daesh—whatever you want to call them—have been

using drones for some time now. They are not autonomous drones; they drive them with remote control. They put a grenade on them, take them in and drop it. They have also been using them in a way that is almost autonomous, by loading them up with explosives, letting them go and when the battery runs out they fall. They are already starting to do that. One of the problems is the more we develop this technology we will see very bad copies coming back that have made no attempt at discrimination or proportionality. That concerns me quite a lot—never mind authoritarian dictators getting hold of these, who will not be held back by their soldiers not wanting to kill the population.

Do not worry too much about an AI cyberwar just now. DARPA, the defence research wing of the Pentagon, has had a big competition going on for a couple of years on AIs for cyberattacks. They did quite well, but they took the very top one to the hackers' conference, DEF CON, and they ran it as part of the competition. It came last by a long way. Human hackers still have the power there.

When you said "chatbots", I think you were referring to the whole idea of bots, such as Tweetbots, which Russia has been clearly using for spreading propaganda in both the United States and here during Brexit. These are very difficult to stop, but we have to make an extreme effort to do so, because I believe they are destroying our democracy.

Major Kitty McKendrick: Drawing on that, the use of Tweetbots, for example, opens a bigger question about how artificial intelligence can be used to gain strategic advantage outside war, and the implications nationally.

The Chairman: Lord Bishop. In some senses, we have probably covered quite a bit of this ground. You may want to tweak it in an appropriate way.

Q159 **The Lord Bishop of Oxford:** I would say that a short answer is fine on this. Should the UK support an international ban on the development and deployment of fully or semi-autonomous weapons? Is an international ban on such technology a feasible proposition? You might want to give quite short answers because we have covered it.

Mike Stone: A short answer would be no.

Major Kitty McKendrick: I am not in a position to make a definite policy recommendation but the most important thing is that our current obligations in international humanitarian law are upheld and that new weapons reviews are used to address as thoroughly as possible emerging technologies, including autonomy. More information should be forthcoming to develop a common understanding of things such as human control so that that can be used in weapons reviews, which should be undertaken at earlier stages of the concept of weapon design and development.

Dr Alvin Wilby: May I pick up on that? We touched several times on this challenge of defining what we mean by all these terms. It is a real issue. One of the exemplars that I know people campaigning for the ban are using is the convention on laser weapons. That is a slightly flawed comparison, in some ways. The laser weapon ban codified a certain area

in terms of what we mean by “inhumane behaviour”. You could have argued before that we would probably have thought that was inhumane and, therefore, would not have done it, but it makes it absolutely clear that that is viewed as inhumane. In that sense, it is positive. The difficulty with using that as an exemplar is if I am given a laser weapon I can tell you with 100 per cent certainty whether it has a laser in, what type of laser it has and whether it is powerful enough to cause blinding. With AI, I have no ability to do those things in a meaningful way, at the moment. The ban, in that sense, is misguided.

Also, to pick up a point that Professor Sharkey made earlier, I think he said the UK was out of step. The UK is probably in step with the majority of the nations that have been discussing this, at the moment, in thinking that we are not in a position to put such a ban in place today.

The Chairman: Professor Sharkey, you might want to respond.

Professor Noel Sharkey: Yes, I will. When I said “out of step”, I was talking about the UK’s definition, which is clearly out of step with everyone else’s. It is not remotely like anybody else’s; it is science fiction. The point about an international ban is you can say that we will never do anything wrong and always comply with the laws of war, yet we also hear if there is an arms race going on we had better get them together so we can use them ourselves. Generally, I have been trying to encourage nations to take their blinkers off and stop thinking about themselves and to think about global security and the disruption of global security by introducing a brand-new autonomous weapon that is unlike anything else we have ever seen before that can free-roam and destroy targets without human intervention. This is a real problem and I am not sure where it is going to go. In the UK I would like to see, at the very least, a national policy, not one line from government Ministers saying, “We’re never going to use them so we’re not going to ban them”. I would like to see a national policy on what they find as acceptable human control, not just a single phrase such as “human oversight”. Let us have the UK, at least, step up to the UN with a proper definition of what they mean by human control. We will work from there.

On your moral question, obviously there are the ethics of the laws of war. That is a compromise between moral realists and ethicists. There are all those issues. One that has come out, strangely, is from American generals. They say the most undignified thing would be having a machine delegated with a decision to kill you; that it would be beneath human dignity. Certainly, this is Germany’s view as well. They have not come out for a ban yet, but number one on their constitution is human dignity.

The Chairman: We will carry that forward into the next question. Thank you very much. Did you want to come in?

Mike Stone: Only to say that the answer is quite clearly no, in my view. I support what Professor Sharkey said about understanding clearly what we mean by human engagement. We have to be very careful here in loosely banning AI if we are talking about truly autonomous things that are set off and can, willy-nilly, go off and select a target. That is very different from most of what we have been talking about.

The Chairman: In a sense, learn for themselves.

Q160 **Baroness Bakewell:** With the phrase of the afternoon being “the genie is out of the bottle”, I feel this is a rather out-of-date question. However, I am eager to ask it because there are so many issues arising from it. What ethical principles—an old-fashioned phrase—should companies developing AI systems for military and security applications use? I am interested in this concept of ethical values because they are, by no means, universal, as I believe Professor Sharkey would agree with. If one company makes one set of ethical values and another, another, we still have the arms race, have we not? Dr Wilby, this is for you.

Dr Alvin Wilby: I had a feeling that was coming in my direction. I have a fairly simple approach to this, and we have touched on it already. There is a body of international law—the Law of Armed Conflict—that sets out generally accepted views of what is and is not reasonable in warfare. We all wish we did not have warfare but we have to be pragmatic; it exists. My starting points are that body of international law and the views of our parliamentary democracy, which I also believe in. We would do things that we expected to be entirely compliant with those frameworks. We would make sure that we followed the spirit and intent of those laws, as far as we could. If I take my example of laser weapons, I could very easily think of a way of using LEDs rather than lasers, thinking that might circumvent the legislation, but we would recognise that the intent here is to prevent blinding, and therefore we would make sure that we were not doing that either. We would also make sure that other systems which are quite legitimately using lasers for other military purposes—laser range finding, for example—were designed to minimise the risk of causing blinding. I would use that type of framework in my thinking. If anybody has a better way of anchoring the industrial perspective, I would be glad to hear it.

Baroness Bakewell: You speak about international law. Does it need radical updating to meet all these changing situations?

Dr Alvin Wilby: You clearly need to keep it under review. The fundamental principles in it are about proportionality—do not use more force than is necessary to achieve the effect—and discrimination, so that you can tell the difference between combatants and non-combatants. It sets a framework that says, “Use the minimum force necessary to achieve your objective while creating minimum risk of unintended consequences”—civilian losses or whatever. I am glad I do not have to make the judgments, but those are the judgments the military has to make in any use of lethal force. Those international laws, which are codified for the military in JSP 383, which is 800-plus pages of the manual of armed conflict, produce rules of engagement that say, “In this particular scenario, this is a behaviour that we think would be appropriate”. For me, that has to be the starting point of any ethical position that I would take, as an industrialist.

Baroness Bakewell: Do you agree, Major Kitty?

Major Kitty McKendrick: It is valuable that industry is cognisant of our obligations under international humanitarian law and produces weapon systems that can comply with them. I would augment it to say that in terms of data analytic software, for example, it also needs to be

cognisant of our obligations under international human rights law, which in co-operation with the defence industry is something that would be considered at the development of any of these systems. Another useful ethical principle is about knowing your customer. A close relationship between the military and defence industry is helpful in that regard, taking into consideration non-proliferation at that stage as well in the design.

The Chairman: I am going to bring the Bishop's supplementary in before we come to the other two, who may want to answer.

The Lord Bishop of Oxford: I think it has been answered, thank you. That has been helpful, particularly the reference to the international body of law and the principles there.

Baroness Bakewell: I used this phrase "ethical principles". Have the new developments which you have elucidated so exactly put the philosophical concept of that phrase under strain? Does it still have validity—that there should be an agreed, international ethical principle on which all these developments can move forward? The eagerness with which people are striving to get ahead of everyone else is apocalyptic, is it not?

Professor Noel Sharkey: I would not like to bring moral philosophy into weapons manufacture, frankly. I do not see that it has much of a place. I think you answered it very well, as to what your job is and what you have to do under international law. There is a moral principle in international humanitarian law, which you did not mention, called the Martens Clause, which is a get-out-of-jail free clause. The Martens Clause is that no weapon should be developed that is against the dictates of the public conscience. How to work that out is a little difficult. None the less, at the UN delegations can sit down and say, "I think this weapon is immoral. I don't think we should go for this". They are the people who can decide it. We have that principle in there, just in case.

The Chairman: Did you want to add anything?

Mike Stone: All I would say is that I welcome the debate. I think it is absolutely right that we should be testing this. The law of armed conflict is absolutely on the minds of everybody in the Ministry of Defence and they intend to fulfil that. We need to be in this debate to shape it, frankly.

Q161 **Lord Levene of Portsoken:** This takes it a little further. Should there be requirements for intelligibility or auditability in AI systems used for military or security applications, particularly those which have the capacity to kill or injure human beings? If so, how might this be enforced?

Major Kitty McKendrick: I will answer as succinctly as possible. In terms of intelligibility, as a military commander or operator that is absolutely essential to having a level of trust in the weapon system you are using. Irrespective of the level of autonomy or the complexity in that system, the operator, the commander, must be trained in its use. That has to be based on some kind of understanding. In terms of auditability, autonomous weapon systems give us an opportunity to have electronically generated audit trails of their use including trusted user

logs, for example, so we know exactly who used a weapon at a certain time. That prevents their misuse, perhaps. Beyond those standards of intelligibility and auditability, there should be other standards in development of validation and verification of software. In fact, Dr Wilby will probably be able to expand on that.

Dr Alvin Wilby: I would fully agree. There is a lot of discussion about white box versus black box AI systems, by which we mean if it is a black box you have no idea how it came to the decision or the identification that it came to. With white box you have more insight into it; it followed this rule and therefore you had that outcome. To make this concrete, we have talked a bit about neural networks, the image recognition systems that people will be familiar with. If you catalogue your home photos you will probably have one of these systems that tries to suggest face names, and such like. Often it gets it spectacularly wrong and will try to label the corner of a building or a tree, which is an example of what I referred to earlier as artificial stupidity; it has gone badly wrong and we do not quite know why. With those systems it can sometimes be hard to understand precisely why it came to the classification it came to.

As Kitty says, you can take a log of what the system's state was at that point. You could investigate it afterwards, you could show it other images and see what classifications it came up with to get some insight into why it got it wrong on that occasion, but you may not be able to get a complete understanding of why it got to the point it got to.

Mike Stone: It is fair to say that there is greater auditability with these approaches than there is when humans are directly engaged in it.

The Chairman: Do you mean when humans are solely engaged?

Mike Stone: Correct. The intelligibility piece is around decision support. Decision support is vital for the military to be able to get inside what is called the opponent's OODA loop—their observe, orient, decide, and act loop—so that we can act, in an informed way, much faster than the enemy.

The Chairman: You have an audit trail, so to speak.

Mike Stone: It is absolutely vital that we have that audit trail, I would agree. One of the key issues we will have, in due course, is we will have to deal with how immutable that data is and how long we keep it. If everything has to be kept for ever, we are going to come up against serious storage challenges, with judicial inquiries, and such like. At the moment, that has entailed literally keeping everything for ever.

Professor Noel Sharkey: Memory is getting cheaper all the time, so we should be able to keep it for ever. Intelligibility and auditability are absolutely essential for any system that impacts on our lives. I do not want to go into AI generally, but we are finding there is a lot of bias—gender and race bias—with predictive policing, et cetera. We need a clear chain of intelligibility. Quite frankly, when you are using neural networks, you have a large matrix of numbers and you cannot tell what feature it has selected. It might be that it is a little hair coming out of the end of my nose. I have no idea what feature it has selected. That makes it very difficult.

In terms of auditability, we need that, and we have Article 36 weapons reviews that assess intelligibility, but with auditability I am not as sure as they are about keeping a nice trail. Machines can malfunction in ways we can never predict. You have all had computers crash regularly. A computer can be hacked, can have a bullet go into it or can be blown up. There could be some unanticipated action which the programmer did not expect to occur and it might do something really strange. If I wanted to commit a war crime, the first thing I would do is destroy the robot that I used to commit the war crime. I am not sure that auditability will be that much easier, but it is essential.

Dr Alvin Wilby: I have one quick follow-up on that. This touches on areas such as safety-critical system design and safety-critical software. If you are designing a safety-critical system that is on your civil aircraft, or a train system, for example, you follow very precise methodologies and rules on how you develop the software and how you test it. You are required to do very extensive fault analysis. You look at everything that could possibly go wrong so that you can guarantee, in principle, that a system is acceptably safe. I accept our view of the world sometimes is not quite as good as it could be, but there is quite a good body of well-developed work which has some relevance to this. In developing systems that are using AI, we have to be sure we understand how we know how well they will perform in the environments into which we want to put them.

Professor Noel Sharkey: I have to come back on that, if I may. We have no formal method at all, at the moment, for computer systems. We can make very good judgments and safety-critical systems by having very good engineers. When you are talking about warfare, you have a lot of unanticipated circumstances. We can formally verify some programs if they are not too long, but if they are long we cannot do it at all. When it comes to a neural network, there is no known method we can use to formally verify it. All we can do is empirically test it. If you are going to do a weapons review, how do you empirically test every possible circumstance, since there is an infinite number of unanticipated circumstances that can happen in a conflict zone?

Lord Levene of Portsoken: Is there not a danger, though, that we are talking about shutting the stable door after the horse has bolted? If the system has killed humans and we say afterwards, "That is all right; we know now exactly why it did it", it is too late.

Lord Swinfen: I have, possibly, an unusual question. We have, at the moment, anti-personnel mines and anti-tank mines that can be left in the ground for months. Would you consider those to be autonomous weapons?

The Chairman: A one-word answer would be absolutely fine.

Major Kitty McKendrick: No, but I see where you are coming from in terms of "human out of the loop".

Lord Swinfen: I did not hear what you said.

Major Kitty McKendrick: No, I would not consider them autonomous weapons.

Major Kitty McKendrick, Professor Noel Sharkey, Mike Stone and Thales Group – Oral evidence (QQ 153–162)

The Chairman: Because of the human agency in laying them, and so on.

Major Kitty McKendrick: Absolutely.

Professor Noel Sharkey: I would call them victim-activated explosives.

Q162 **Baroness Rock:** Thank you all for an extremely thought-provoking session.

The Chairman: You are going to be very strict with our witnesses.

Baroness Rock: I am going to be very strict, I am afraid. If there was one recommendation that you would like to see this Committee make at the end of this inquiry, what would it be? I would be grateful if you could stick to one. Mr Stone, perhaps I could start with you and go along the panel.

Mike Stone: I would like to leave you with the thought that if you want to consider restricting anything you need to be very careful about the way in which you define things. Autonomy and AI are very different. AI enables autonomy but we do not want in any way to be labelling everything as AI and therefore throwing out all sorts of fantastic things. Apart from anything else, as I said, the civilian world is driving this now, not the military world.

Major Kitty McKendrick: Building on that, I would ask the Committee to consider the positive benefits of artificial intelligence for human security, specifically in conflict prevention and mitigating the effects of conflict through resilience. There is huge potential there.

Dr Alvin Wilby: We have talked about the discussions that are going on now, which seem to me to be polarised into “ban everything” and “do nothing”, if I can put it that way. It is too important a subject not to be studying. For me, it would be making sure that the MoD, together with industry and academia, does a piece of work that says, “How would we use these types of military systems in the scenarios we realistically might expect? What would that look like, if we were to do it properly in accordance with the law of armed conflict in a sensible way?” and see if we can generate some useful guidance.

Professor Noel Sharkey: Mine would be quite clear. Some other European nations are going to step forward and say they want to join the 22 countries which have already called for a ban and the 100 other countries asking for a new legally binding instrument to be discussed. The UK has a great opportunity to take a leadership role at the UN and say that they will stand up and stop these obnoxious weapons. If they will not do that, could we, at least, have the development of a sound national policy that defines “human control”? We need to take away the ridiculous definition and have somebody with some expertise look at it, other than the MoD.

The Chairman: We are back to the definitions point, very helpfully. Thank you all very much indeed. It has been a very interesting discussion. Thank you for your self-discipline, if I may say so. It is very easy to answer everything at length, but we got through with some useful points. Thank you very much.

Professor Peter McOwan, Professor Sir David Spiegelhalter and Professor David Edgerton – Oral evidence (QQ213–223)

Professor Peter McOwan, Professor Sir David Spiegelhalter and Professor David Edgerton – Oral evidence (QQ213–223)

[Transcript to be found under Professor David Edgerton](#)

MMC Ventures, Project Juno and Eileen Burbidge MBE – Oral evidence (QQ 46–54)

MMC Ventures, Project Juno and Eileen Burbidge MBE – Oral evidence (QQ 46–54)

[Transcript to be found under Eileen Burbidge MBE](#)

Professor Dame Henrietta Moore, Professor Richard Susskind OBE and Future Advocacy – Oral evidence (QQ 95–104)

Professor Dame Henrietta Moore, Professor Richard Susskind OBE and Future Advocacy – Oral evidence (QQ 95–104)

[Transcript to be found under Future Advocacy](#)

National Data Guardian for Health and Care Dame Fiona Caldicott, NHS Digital and Dr Hugh Harvey – Oral evidence (QQ 128 – 142)

National Data Guardian for Health and Care Dame Fiona Caldicott, NHS Digital and Dr Hugh Harvey – Oral evidence (QQ 128 – 142)

[Transcript to be found under Dr Hugh Harvey](#)

NHS Digital, Dr Hugh Harvey and National Data Guardian for Health and Care
Dame Fiona Caldicott – Oral evidence (QQ 128 – 142)

**NHS Digital, Dr Hugh Harvey and National Data Guardian
for Health and Care Dame Fiona Caldicott – Oral
evidence (QQ 128 – 142)**

[Transcript to be found under Dr Hugh Harvey](#)

NVIDIA, Dr David Barber and Digital Catapult – Oral evidence (QQ 38–45)

NVIDIA, Dr David Barber and Digital Catapult – Oral evidence (QQ 38–45)

[Transcript to be found under Dr David Barber](#)

Ocado, Dyson and Fujitsu – Oral evidence (QQ 105–115)

Ocado, Dyson and Fujitsu – Oral evidence (QQ 105–115)

[Transcript to be found under Dyson](#)

Sarah O'Connor, Andrew Orlowski and Rory Cellan-Jones – Oral evidence (QQ 9–17)

**Sarah O'Connor, Andrew Orlowski and Rory Cellan-Jones
– Oral evidence (QQ 9–17)**

[Transcript to be found under Rory Cellan-Jones](#)

Open Data Institute, Open Rights Group and Privacy International – Oral evidence (QQ 65–75)

Evidence Session No. 8

Heard in Public

Questions 65–75

Tuesday 31 October 2017

[Watch the meeting](#)

Members present: Lord Clement-Jones (The Chairman); Baroness Bakewell; Lord Giddens; Baroness Grender; Lord Hollick; Lord Holmes of Richmond; Lord Levene of Portsoken; The Lord Bishop of Oxford; Lord Puttnam; Viscount Ridley; Baroness Rock; Lord St John of Bletso; Lord Swinfen.

Examination of witnesses

Frederike Kaltheuner, Olivier Thereaux, Javier Ruiz Diaz.

Q65 The Chairman: A very warm welcome to our witnesses today: Frederike Kaltheuner, policy officer from Privacy International; Olivier Thereaux, head of technology, the Open Data Institute; and Javier Ruiz Diaz, the policy director of the Open Rights Group. I hope I have pronounced all your names correctly, given my poor linguistic skills. It is very good to see you. We are expecting a vote in the House very shortly, but I will go through the rubric and ask you to introduce yourselves. Hopefully, the vote will be thereafter and we can resume after it.

The evidence session is open to the public, and a webcast of it goes out live and is subsequently accessible via the parliamentary website. A verbatim transcript will be taken of your evidence and will be put on to the parliamentary website. A few days after this session, you will be sent a copy of the transcript to check for accuracy, and we would be grateful if you could advise us of any corrections as quickly as possible. If, after this session, you wish to clarify or amplify any points made during your evidence, or you have any additional points to make, you are very welcome to submit supplementary written evidence to us. If there is a vote we will adjourn for 10 minutes or so. It might be slightly longer, depending on how long it takes us to vote. We will adjourn as soon as the Division bell goes, so that it does not inconvenience you in trying to speak over it. Would you like to introduce yourselves for the record? Then we will begin with questions.

Javier Ruiz Diaz: Good afternoon. I am the policy director of the Open Rights Group. The Open Rights Group is a grass-roots organisation that works on all aspects of digital technology. We work on privacy and on access to information, intellectual property, and generally anything where fairness and the rights of individuals may be affected.

Olivier Thereaux: Good afternoon. I am a technologist. I am the head of technology at the Open Data Institute or the ODI. We are global, non-partisan, independent and non-profit, and we are working to build a better data economy. We work with companies and Governments globally to help, inform and enable them to work better with data.

Frederike Kalthener: I am a policy officer at Privacy International where I lead our strategic work on data exploitation in the private sector. We are an international charity and we have partners around the world, particularly in the global south, which gives us a unique perspective on what happens when emergent technology is applied without appropriate legal or technical safeguards.

Q66 **The Chairman:** Thank you very much. Since the Division bell has not gone off, I will ask the very general opening question that we asked of our previous witnesses. Who should own data and why? In that context, is personal ownership of all data generated by an individual feasible, and if so, how?

Javier Ruiz Diaz: The Open Rights Group would also be against the idea of ownership. It is worth recalling that the Court of Appeal declared in the *Datateam* case in 2014 that ownership of data in the UK was not the way to go. So that was established in law. We believe that a framework of rights applies that can be quite complex. There are also commercial confidentiality issues, but the idea of ownership as such is not useful.

Having said that, when it comes to trying to protect the public and establishing the public interest in data in any other context we would normally look at intellectual property, and that is the challenge that we have here. For example, we will be looking at open source, with publicly-funded software or open access journals with publicly-funded research. The challenge is how we do that in artificial intelligence with personal information when we cannot use the same framework. We agree that that is a difficult problem.

I would go even further and question not just the idea of the property of data but the idea of applying intellectual property to other aspects of artificial intelligence. Understanding how artificial intelligence systems work is quite complicated, but it is becoming quite clear that the separation of data and software is increasingly blurred. Before, we could say, "Let's just make it open source", and have a data framework to go with it.

The Committee suspended for a Division in the House.

The Chairman: Let us resume. One of the ironies is that here we are, looking at a future technology, and it takes a rather long time to get through the Division lobbies because we do it the old way. There we are. You were just concluding, Javier.

Javier Ruiz Diaz: The use of an intellectual property framework to provide certain protections and defend the public interest in the realm of artificial intelligence is quite tempting, but it may be problematic, particularly because we are not sure what type of intellectual property, if any, will apply to an artificial intelligence system. If you look at the contracts that DeepMind has been signing with NHS trusts, it lays claim

to exclusive ownership of any and all developed IP, as it calls it. That includes designs, works, inventions, software, data, techniques, algorithms, know-how or other materials. The problem is that nobody really knows what these outputs look like. As we said, a combination of data and evolving software is being produced.

The Chairman: I am going to move on in a minute. Just to cut you short, you object to the wide-ranging terms of that particular type of contract.

Javier Ruiz Diaz: I certainly do, yes.

The Chairman: Thank you. We might unpick whether or not there should be a statutory override, and so on.

Olivier Thereaux: I agree with Javier that the notion of ownership of data is really problematic and generally unhelpful. Data has a few qualities that makes it incompatible with notions of ownership. I can hold it, you can hold it, and my holding of it does not impact the value you derive from it. Some types of data are facts and should be in a data commons, such as a bus timetable or a list of UK addresses, which would go towards ownership. Take the data about a phone call. I make a phone call to a friend. The data about that call has me as the data subject, but it could not easily be owned just by me because there are other data subjects. The person to whom I made that phone call is a data subject, the companies through which I made that phone call are another kind of data subject, a secondary data subject, and so on. There is quite a lot of complexity once we look at a single data point, which makes it quite difficult to apply the notion of ownership. However, once we talk about control, rights and responsibilities, even within that complexity we can apply those things.

I would add a small caveat to that. My believing that the notion of ownership is not useful does not mean—that this goes to the second part of your question—that individuals should not have control over it. On the contrary, saying that they do not own the data about themselves does not mean that they should not have control. They should have control, some of it directly and some of it through democratic means.

The Chairman: You would agree with the contractual override point.

Frederike Kalthener: I agree, and I think there is a reason why in civil society the very idea of data as property or data ownership is highly controversial. One point is feasibility. The other is: is that really what we want? I do not want to live in a world where privacy becomes a luxury. If something is property, that means you can sell it away. Data governance, be it data protection or other forms of regime, essentially governs imbalances of power between those whose data is being processed and those who are controlling it. This is one reason why we also like to think of data as a form of control that comes with rights.

As we stand, and this is important when we look at artificial intelligence today, individuals have lost complete control over their data. This is confirmed by a study conducted by the ICO in 2016, which shows that only one in four adults in the UK trust businesses with their data. That is quite concerning. There is a profound lack of confidence in companies'

handling of data. In the same study, 55% of individuals did not think that businesses are transparent about their processing. We read this to mean that people are very concerned.

The context of this concern is that people's understanding about advanced processing is extremely low. Things such as profiling, tracking and automated decision-making require a degree of technical expertise that I do not think we can assume in the general public. We still think about privacy as what we share online, whereas we are being tracked increasingly in our homes and in public spaces. The privacy implications of AI applications are in their power to recognise patterns and to infer very sensitive details from data that we might not think is very personal or sensitive.

The Chairman: Taking this a bit further, because you are all agreeing on the rights questions as opposed to the ownership question, if data is generated by multiple data subjects, how do you work out who is entitled to control it, share it and so on? That is quite tricky, is it not? Would you like to follow that up?

Olivier Thereaux: We are doing some work on that at the moment, so my answer will be mostly preliminary. The idea is to identify the data subjects and then give them control. I am sure we will have a chance to talk more about that later in our discussion, but being able, as a data subject, to access data that is about me when it is held by a data broker or a company, through means that are easy to use—

The Chairman: It might be generated with another data subject, i.e. it could be somebody you are on the phone with, or somebody you sent an email to. How do you regulate that, so to speak?

Olivier Thereaux: That is problematic, because we cannot necessarily work under the assumption that the secondary data subject will be identified and identifiable and will know about this.

The Chairman: I only throw out these ideas to confuse the situation. Let us move on.

Q67 **The Lord Bishop of Oxford:** The Government's review of AI was published two weeks ago. It addresses key issues relating to data and AI. Do you think it goes far enough? If not, where would you go further?

Javier Ruiz Diaz: We think the review has many positive points. However, it is quite focused on growing the AI sector, and possibly not sufficiently dedicated to seeing in which cases we perhaps should not do AI and in which cases we should not have access to data. That is one point. We would agree with many of the recommendations on data and access to data. However, we have lots of questions about the idea of data trusts. We think that individuals affected would need to understand each individual processing of the data, not just that their data has been put into a data trust, and lots of different users are going to have access to the data once it is there.

The Chairman: You would be rather suspicious of that, would you?

Javier Ruiz Diaz: We would be suspicious of the idea of simply creating an environment and then removing the data to the data trust itself and

saying, “Now you have to trust that the data is going to be kept safe”. Individuals will want to know what happens in each case. If an SME is going to come in, they will want to know what the SME is going to do with the data rather than just have it and they will want to know how it is going to put it in a safe place.

On the other recommendations, the recommendation for a framework to be created jointly by the ICO and the Alan Turing Institute is also potentially problematic for one simple reason: the ICO’s remit is based on data protection, and, as was discussed before, wider issues of privacy, fairness and discrimination may not be covered by the ICO. We think that you might need input from other regulators and take a much broader approach.

The recommendations on skills are positive. However, we think they should include broad interdisciplinary skills, such as data ethics, and not just focus on computer science.

Olivier Thereaux: I will add to that. First, the report truly and rightly establishes access to data as one of the cornerstones of the development of AI. Data trusts are an interesting idea. I am not sure whether they go far enough. At the moment, we have market failure in access to data. Data held by the private sector is not made as open as possible to enable AI to take off.

The Lord Bishop of Oxford: If not data trusts, then what?

Olivier Thereaux: I will answer with anecdotal evidence. We have been working with the pharma sector, which is quite private and secret by default, but in the case of antimicrobial resistance it has the incentive to collaborate to a greater extent than it typically would. We have been working with seven big pharmaceutical companies to share data. It is important that you match these incentives. In this case the incentive is fairly existential, but it will not always be the case; sometimes there will be a business incentive and sometimes a regulatory incentive. Once you have the incentive, you must have mechanisms in place. Digital trusts are one such mechanism, but they might not be the only one, to safely enable not just contractual access to data but equitable access to data as much as possible. Contractual access between two parties is good, but making sure that access is made equitable, especially when we are talking about data available to or held by government for instance, is a recommendation that I would make.

Frederike Kaltheuner: The report had a very specific mission, which is to interrogate how the AI industry can grow in the UK specifically. That focus inevitably meant that there were areas that were underexplored. It is vital to guarantee access to data. The question is what kind of data, when, at the same time, individuals feel that they have lost complete control over it. We cannot talk about one without also addressing the other. It is important that AI and innovation in AI does not become a pretext to further undermine individuals’ trust. It is also important, because with AI privacy can be a time-shifting risk; information that is not a problem for you today can be very harmful in the future.

There is a recommendation to improve the skills of computer scientists, and now there have been discussions to teach ethics to computer

scientists. It is also important to improve the skills of those who hold complex AI systems to account. Let me give you an example. There are a lot of vital civil society organisations that work on housing rights, welfare rights, policing, justice and LGBT rights. All these organisations are now faced with new logics of discrimination that require very technical skills to challenge them. We need to talk about this in terms of civic skills in a vital society. I can elaborate on this in answer to the next question, but we think this is one of the reasons why it is important also to have such a thing as collective redress in data protection law.

Finally, whenever there are positive reports—and AI is exciting—we also need to talk about what AI cannot do. I have seen from the submissions to this Committee that the police are very interested in using AI for all sorts of applications. AI creates knowledge that is inherently probabilistic, and we need to have a siloed discussion about areas where we feel comfortable with using and relying on such kind of knowledge, and areas where we would rather rely on other forms of evidence.

Q68 **Baroness Rock:** That leads very nicely on to my question on legislation and data protection. We heard in the previous session about the importance of accountability, and, indeed, the Information Commissioner talked about the new tools of law that did not exist before. I have two questions. Is the Data Protection Bill currently being considered in Parliament fit for future challenges? Is there a risk that new data protection legislation could hinder progress in AI research and, indeed, in the work that is happening particularly with SMEs and start-ups?

Javier Ruiz Diaz: We are doing quite a lot of work on the Data Protection Bill right now, and of course we think there are lots of things that could be improved. We do not believe that, as it stands now, it would cause a competitive disadvantage to the UK. I think that was discussed before; the whole world is moving in this direction. Even in China, many people are saying that they are doing their homework to be able to process data from Europe, too, in the future.

Having said that, we do not believe that the law itself will provide all the assurances. We need individuals to feel that their rights are protected and can be enforced. In that sense, as Frederike was saying, there is a very important clause in the Data Protection Bill, Clause 173, under which we are asking for collective redress to be expanded to situations where individuals might not be aware that their data is being processed, which we think is particularly relevant to artificial intelligence systems. If we can have strong legislation, which is there, strong civil society helping the regulators to enforce it, and individuals who feel empowered and confident that their rights are protected—and if they are not they can complain and something will happen—we will have a better situation.

Frederike Kaltheuner: I should say from the outset that we cannot subsume our privacy concerns of AI under data protection. That said, the GDPR was the fruit of many compromises. It reforms without revolutionising data protection and it gives more rights to data subjects. That is a good thing. However, it also gives Governments a lot of discretion. To frame it in economic terms, what will definitely harm the AI sector in the UK is if the UK does not obtain adequacy post Brexit. It is

important to look at the Government's discretion in this scenario. There are things that we are concerned about, such as the very broad set of grounds for public interest. We like public interest, it is a good exemption, but it needs to be clearly defined.

This also has relevance for processing of data revealing political opinions, as was mentioned in the previous session. This is very important politically, not just for privacy. The Bill should not be seen as in a vacuum. The ePrivacy regulation, which has been adopted by the European Parliament, also has implications that are quite relevant for AI. If this text is adopted, and ideally strengthened, it will provide privacy by design. It also regulates machine-to-machine communication, and it is crucial that this regulation is adopted in UK law. Finally, data protection privacy is a fundamental right. There is also the point to be made that these rights should be enshrined in a more abstract and fundamental way that goes beyond the Data Protection Bill.

The Chairman: You are going to draft even more amendments than you already have, by the look of it.

Olivier Thereaux: Speaking here as more of a technology expert than a policy expert, those two pieces of legislation are a good start, but there are two particularly interesting things, especially in relation to AI, that are worthy of more discussion and scrutiny. The first is the right to explanation. We touched on that a little. It is important to understand that the hardest thing in AI is explaining the AI. You can publish the weights of a neural network, but that tells you absolutely nothing about what it does and why; you will have to look at it as a black box. Unless we have access to training data, the right to explanation cannot be done in a technical manner, so we need to rely on people to do the work and to take responsibility for the right to explanation. I know that the proposed legislation has no right to explanation in it per se, or certainly not at a technical level, but the proposal in the AI review to have a framework for explaining processes, services and decisions delivered by AI is a good one. Pushing for more research so that we can eventually have better technical solutions for explanations of AI is good, but in the meantime we do not have this ability when we are using a black box.

The second point that I wanted to make, less as a policy expert than as someone who has been building digital systems for a long time, is about the notion of informed consent. At the moment, the user experience of informed consent is, if you will allow me a technical term, a joke. Presenting the individual with 100 pages of legalese that they could not possibly comprehend and a single box between them and the service that they want to be using puts the onus on them to take the responsibility to say, "Yes, I have made informed consent. I have ticked that box". The providers of services could do much more to make that consent informed and unambiguous. We are not pushing them hard enough towards that.

Q69 **Lord St John of Bletso:** Do you believe that we need a specific AI watchdog or regulator to deal with the issues surrounding data and privacy? As you know, our previous witness was the UK Information Commissioner.

Javier Ruiz Diaz: We do not think there is a need for a specific regulatory body, because the issues raised by AI—the right to privacy, fairness, discrimination, competition—go beyond data protection, as we said before. However, it would be useful and helpful to have a trusted independent body that could provide expert advice and support, that could issue opinions and recommendations, and potentially help individuals if they had a complaint or a query about a system, with the caveat, as Olivier said, that explanation and explainability may be limited in some cases, but we strive for it.

No, we do not think a regulator as such would help. We need to get many regulators to be AI informed and to be able to incorporate AI into their work. We think of equality, human rights, for example, and discrimination. Public procurement in particular has no specific regulator, but incorporating AI into how public procurement works would be quite important. Yes, we think information and advice is very useful, but a new regulator may end up overlapping with many other regulators, and in the end they may also have many gaps and cracks for things to fall into.

Olivier Thereaux: I agree with Javier that ultimately it feels premature to have a regulator for AI. It is probably more useful right now to recognise that AI is going to be used across many sectors, many of which already have regulators. Those regulators need help to figure out and understand the impact of AI and how they can deal with AI in their specific context.

The Chairman: That is really interesting: integrate it with existing regulation as opposed to erecting a new edifice.

Olivier Thereaux: Yes. For the moment, that would probably be more impactful.

Frederike Kaltheuner: I agree. We use AI to mean many different things at very different levels of complexity and abstraction. The implications, not just for privacy, can be radically different in different contexts. The trend, at least in the context of data protection, is to recognise the ICO as the key body to ensure respect for data protection rules. The ICO will have more powers in the coming years, and it is important to make sure that the body is able to exercise its mandate. I agree that not all risks fall under data protection, and there is a need for discussions about defining some of the terms suggested primarily by industry bodies at the moment; “fairness” and “transparency” are not defined and we need workable definitions.

Lord St John of Bletso: Would a limitation to an AI watchdog be getting the right staff with the relevant experience and expertise in this sector? There appears to be a bit of a shortage of that level of expertise and experience. I am referring to a watchdog rather than a regulator.

Javier Ruiz Diaz: In our experience, there seems to be a general lack of skills. We were joking before that most AI meetings start by someone declaring a conflict of interest in that they are going to start working for DeepMind, which is taking most scientists in the country. We think there is a shortage of skills. However, that is something to build in the programme. We do not see it as a limitation.

Q70 **Lord Levene of Portsoken:** How can data be managed so that it is used for the public good?

Javier Ruiz Diaz: We listened to the previous discussion. It can get quite tricky to define the public good. As you say, there are conflicting public goods. Taking an ethical approach, as was discussed before, rather than just a strict application of the letter of the law, particularly when it comes to data protection, which can sometimes become a bit narrow, is important. As we said before, applying the law in the strongest form possible is also important, as is having a strong civil society, as discussed. I am very critical of the idea of consent and control, with the caveat that of course we support the stronger consent models in GDPR. We need to understand some of the limitations to the idea that we can just push consent as the solution, and we need to start building systems that simply do not do bad things, rather than relying on people to say, “No, I don’t want you to do bad things with my data”. If you do not do it, you do not need to ask. We think that is the solution long term. Of course, that requires taking a much broader approach than looking at the short term.

The Chairman: Are you suggesting some kind of ethical screening?

Javier Ruiz Diaz: Yes, and not just relying on saying, “We are going to do all these things”, when it is a bit unclear what they are and it may or may not be good. You need to be a lot more up front about what you are doing, and if you identify something that you think people may not like, either do not do it or build an alternative and you explain why this needs to happen—as is the case with public interest legislation or the public sector—how it may have detrimental effects for some people but how it will also benefit or provide some public interest. You can construct an argument, particularly for the public sector, for why you are doing something for the public good and explain at the same time that it will be detrimental to certain people. That happens all the time.

The Chairman: Just connecting the answers to the last question to this one, with apologies to Lord Levene, are we saying that in a regulatory approach in different sectors you might require ethical screening for certain types of data before it can be used or stored? Is that the logic I am hearing here?

Javier Ruiz Diaz: I think that would be logical. Even within the strict application of data protection, if we look at how data protection should be applied, there are many grey areas, such as what implementing data protection by design means exactly. You can do it in a minimal way or you can do it properly. Even within the application of the law, you can do it right and ethically or you can do it to get away with as much as you can. We agree that we need to build an ethical imperative in many of these areas.

Q71 **Lord Levene of Portsoken:** What technical approaches might help to preserve privacy while also ensuring that the benefits of AI can be realised?

Frederike Kalthener: In answer to the previous question, if that is not rude, when it comes to really complex systems good intentions can still

result in harm. That is why it is so important to have appropriate safeguards and ask people whether they are consenting. We also need to look at the fact that sometimes the risks and opportunities of a project with the best intentions are not equally distributed. That is another aspect to keep in mind.

When it comes to technical approaches that help, this is a very broad question, because AI can mean different things. I am happy to provide more evidence in writing. In general, privacy-enhancing technologies are the gold standard. The way this question is framed suggests there is a potential trade-off between privacy and ensuring the benefit of AI. This is fundamentally misguided. We can only have progress in AI if privacy is protected. Especially when it comes to consumer-driven AI devices—smart home appliances—companies have a very long way to go. There is an information asymmetry between what consumers know what their devices are currently doing, what they are capable of doing, whether they have been compromised, whether they are secure, and the knowledge of companies. It is a problem that it takes organisations such as the Consumer Council of Norway, to uncover the fact that smart toys for children are fundamentally insecure and privacy-invasive.

Lord Levene of Portsoken: They are fundamentally what?

Frederike Kaltheuner: The Consumer Council of Norway frequently does studies on consumer products that rely on AI technology. They did a project on a doll called Cayla.

The Chairman: You said it was insecure and privacy-invading.

Frederike Kaltheuner: Yes. The doll was recording children’s voices and conversations and transmitting those back to the company in an insecure way. The threshold is quite low for companies to make products privacy-protective. That said, there are ways of making smart devices more secure and more privacy-respecting. We also have to keep in mind that there are business models based on exploiting people’s data. The advertisement-driven internet is based around tracking and targeting.

The Chairman: Do you mean things such as RFID and that kind of tracker?

Frederike Kaltheuner: It means we are being tracked beyond cookies and not just on every website you visit. When you go to advertising conferences, the next frontier of this tracking and profiling is the home. We have to keep in mind there is a potential misalignment of interests. We used to say that whenever you do not pay for a service you are the product. Now we are paying for devices and we are still the product. This is a fundamental problem.

Viscount Ridley: On this specific point, you make the point that AI can help privacy, or need not hinder privacy, as it develops. Thinking about it, I am a lot more comfortable with a bot knowing a lot about my habits through my internet usage than I am with a person knowing that. That person might tell the *Daily Mail*—not that my habits are worth telling the *Daily Mail* about, I hasten to add. I think there is a degree of comfort that when some ridiculous ad pops up that can only have popped up because it knows I was looking for dog leads on Amazon a week before, so it is

trying to sell me more dog leads, as happened to me the other day, the fact that a machine has done that is to some extent reassuring to me because I feel no one is going to care. Am I being unreasonable here?

Frederike Kalthener: I am glad you asked this question. The Equifax breach and the excellent hearing in the US Senate on this issue was very informative, because individuals are being tracked, both on websites and in public spaces. Behind all these targeted advertisements that sound mundane and seem as if they do not have a significant effect is an entire ecosystem of thousands of companies that are in the business of amassing profiles on people. Equifax has 800 million customers around the world with data brokers who categorise your romantic interests, your neuroticism, et cetera, and this ecosystem that has been built for advertisement, is increasingly tapped into by unrelated institutions. It is interesting to policing, intelligence, credit scoring, employment hiring, and this is what data protection governs.

The Chairman: I will bring in Lord Hollick, but I would like Olivier and Javier to come in. Some of us may be less relaxed than Viscount Ridley about these matters.

Olivier Thereaux: On the question of privacy, AI is both the good cop and the bad cop, depending on how it is used. There is a way to use AI that is more privacy-preserving than a lot of other computer systems. AI in particular is relatively computer-unintensive compared to a number of other systems, which means that you can train it somewhere, give the AI model to a device that is relatively underpowered—most of our consumer electronics are underpowered—that could run the model in your device without ever having to call home. So there are ways in which you can use AI that are very nicely privacy-preserving, because they do not need to call home, unlike most other systems. You can get a personal recommendation and a personalised service based on your personal data without personal data ever leaving your device. That is one end of the spectrum.

At the other end of the spectrum, AI is particularly problematic, because it can be used extremely efficiently to re-identify people. I stress that, because we have pretty good de-identification methods. I also stress that I am not saying “anonymisation”, because that is a misnomer. There is no such thing as anonymisation; it does not work. Anonymisation is transient. You may be anonymous through some technical mechanisms, for example differential privacy, but according to experts, in some circumstances differential privacy, only works when you put so much noise in the data that it becomes useless. Therefore, de-anonymisation is something that happens and AI is responsible for making that easy to do.

The Chairman: I am going to get Lord Hollick to ask his question and, Javier, you might want to respond. I have Lord Giddens looking excited over there, too. It depends on the time.

Q72 **Lord Hollick:** We heard in the last session that the Information Commissioner will have oversight of the exploitation of public data and presumably will be able to decide whether it is an appropriate thing to do and what the terms would be, whether they are financial terms or public benefit. Do you agree?

Javier Ruiz Diaz: I am not completely sure that the Information Commissioner could set financial terms.

Lord Hollick: Let us say public benefit terms.

Javier Ruiz Diaz: There is public benefit, yes. That is one of the things that needs to be developed in the context of GDPR. At the moment, we have very good ideas, as was discussed before, about public benefit and public interest in the context of freedom of information. We do not have a very well-developed corpus on data protection yet, but hopefully that will happen. However, whatever is defined there will be strictly framed under the remit of data protection, which is fairness in information processing. Things such as setting up financial terms may be outside the remit of the ICO, which is one reason why we have said before that we may need to look a bit beyond the ICO in some of these areas.

Lord Hollick: Is there not a danger that very valuable public sector information will be handed over to people for a few pennies and they will make a tremendous amount money out of it?

Javier Ruiz Diaz: One of our main concerns here is that public sector bodies do not seem to understand the value of their data and seem almost to be going round begging the IT companies.

Lord Hollick: Who is going to help them to appreciate the value of their data?

Javier Ruiz Diaz: We think the ICO will have a role to play in ascertaining the fairness of the transaction from the point of view of individuals affected. From the point of view of the institution or defending the value for the public sector, or the Government as a whole, we think we may need some extra instruments at the ICO that it is simply not set up to provide yet. We absolutely agree on the need to answer this question. I mentioned public procurement earlier. We need to start looking in that direction.

Q73 **Lord Giddens:** Could you comment very briefly on what impact you think the increasing sophistication of face recognition technology makes to all this? I was daunted to read of the number of restaurants and public environments that have introduced digital cameras. They are often not visible, and they are used to track people, record their conversations and now potentially identify them. This is an awesome thing. In China, the state is using them directly, of course. Here, they are used mainly by companies, but we have never seen anything quite like that. How would you assess that?

The Chairman: Who is best for answering that one? Frederike, you are in the hot seat, clearly.

Frederike Kaltheuner: This is clearly a concern. The study you refer to found that with AI technology there might be the potential to lip-read people's conversations from CCTV.

Lord Giddens: They quite often record them as well, apparently, in restaurants.

Frederike Kalthener: There is such a thing as anonymity in public space, which we need to explore as a value in and of itself. At the same time, there are links to profiling; facial recognition is used to profile individuals to learn, for example, whether they have a problem with alcoholism.

The Chairman: It is covered by the GDPR.

Frederike Kalthener: It depends. The way forward is to say that we can mitigate some of these risks by limiting the resolution at which, for example, CCTV operates or at which facial recognition can operate.

The Chairman: We will not go into the thickets too much. Clearly, that is a very interesting area that we ought to be thinking about.

Q74 **Lord Puttnam:** How can we mitigate against reinforcing unintended prejudices in the use and misuse of data? I do not want to lead you, because you will be asked by the Chairman what you would have us do, but we have a Data Protection Bill going through Parliament at the moment, which strikes me—and I have been away for a while—as extraordinarily weak in the area of education. It strikes me that the only serious answer to the question I have asked you is a generation of re-educated young people who understand the dangers and are able to combat them. What do you feel that we as a Committee, and in fact Parliament, could do in this area to strengthen democracy?

Javier Ruiz Diaz: I am glad you mention democracy. The first thing we need to understand is that we cannot delegate all our responsibility for making decisions to computer systems. We cannot expect fairness and non-discrimination from systems. There will be an ongoing battle to make sure that these things are fair. We also need to understand that the decision has to be made at the policy level and then for us both, if humans and computers are involved.

The first things here are discrimination and the use of data, but this goes to the very idea of fairness. Much of data protection is based on the idea of the liberal self that has free will and autonomy to operate. We need to understand that artificial intelligence is taking place in combination with behavioural economics and with manipulation techniques. It is bringing into question not just whether there is so much a privacy infringement to an individual, as described before, but whether there is systemic manipulation of society, bringing into question the very idea of free will, our role and democracy. In that sense, eternal vigilance is the short answer.

I agree with you about education. A big part of our work is to raise awareness of these issues in civil society, and it would be good to see it done in a more systematic way. I am afraid there are no simple solutions. The problem is even bigger than we are seeing.

Olivier Thereaux: I am reminded of the words of another technologist called Maciej Cegłowski, who said that machine learning, which is one of the methods of artificial intelligence, is like money laundering for bias. We take bias, which in certain forms is what we call “culture”, put it in a black box and crystallise it for ever. That is where we have a problem. We have even more of a problem when we think that that black box has

the truth and we follow it blindly, and we say, “The computer says no”. That is the problem.

I want to mention a model that somewhat poetically comes from the world of chess, which, as you can imagine, was one of the first to be disrupted not so much by AI but by very powerful computing. The best chess players today are not humans or computers but what we call centaurs; the mythical creature that is half-horse, half-human. The idea is that if you have the body and the strength of a horse and the head of a man you get something that is better than the sum of its parts. Follow me for a second through that analogy.

The Chairman: The mind boggles. What is AlphaGo then?

Olivier Thereaux: AlphaGo would probably benefit from that centaur model, even though the problem is that AlphaGo is weird and people do not understand what it is doing. There is a move by AlphaGo that nobody can explain. Something we could explore, especially in the usage of AI for decisions that have significant impact on people’s lives, is having the model whereby the AI is used to inform the decision but the ultimate responsibility is with a person. It is not an operator who is going to say, “The computer says no”, but someone who understands how the AI has been trained and who possibly understands some of the inherent bias in the AI and can overrule the AI by saying, “In this case, this is wrong and I take responsibility for overruling it”.

Frederike Kaltheuner: It is important from the perspective of a consumer. When I sign up to a system I would like to know whether it is racist. I want to know whether the operator of the system has tested for such biases, and I want to know retrospectively if I have been wronged. These are different things. On the one hand, there is mitigation, but there is also redress and oversight. I very much liked Sandra’s comment in the previous session on the right to explanation. That is a first step. We have to keep in mind that the GDPR is very narrow in Article 22; it covers a very narrow set of decisions. We are concerned. I agree with what you said. However, it also works the other way round. If you have a human decision-maker who cannot critically assess the decision-making process and the decision-making process is biased, we have the same result as if the computer made the decision. This is relevant, let us say, in predictive policing with propriety software. If the guidance is followed, we have biased decisions.

The Chairman: There is also nothing at all about transparency by design anywhere that would lead somebody to think of what they are constructing and its implications.

Javier Ruiz Diaz: Having said that, the accountability principles in GDPR could be interpreted as requiring quite a high degree of transparency in data processing.

The Chairman: In advance, in design?

Javier Ruiz Diaz: Yes, and being able to explain what data you use, where it comes from and what you do with it. As I was saying before, how hard you want to implement it is a grey area. You can have a de minimis interpretation or you can have a really strong interpretation

where you lay out all your algorithms and your systems, and there is your accountability.

The Chairman: We reach a crescendo with Viscount Ridley.

Q75 **Viscount Ridley:** Is there one suggestion that each of you would like to make on what this Committee recommends? Be as specific as you can, because we are hearing a lot of problems that need discussing rather than solutions to those problems.

Javier Ruiz Diaz: I thought about that quite a bit when I got the email.

The Chairman: Do not forget that you are allowed only one recommendation.

Javier Ruiz Diaz: Yes. I decided to be restrained and realistic rather than asking for the world. Focusing on the area of public procurement may be beneficial, mainly because this is an area where there could be some direct impact in directing government. Trying to regulate private sector or the wider technology in the abstract sense may be harder.

Javier Ruiz Diaz: Focusing on public procurement of AI and ensuring there is fairness, that there is a public benefit and that individuals' rights are respected would be particularly fruitful.

Olivier Thereaux: My one recommendation, which is not perhaps massively practical, would be work towards making AI and data boring and safe, by which I mean: consider it as infrastructure rather than considering it as a raw material. If we consider it infrastructure we will work on skills, standards and access, and this will help.

The Chairman: As opposed to treating it as voodoo or something. What is the countervailing possibility?

Olivier Thereaux: The narrative that is often used, and which is extremely problematic, is considering data as oil, as something that has value. I did not react earlier, but assigning a financial value to data is problematic, especially when you are talking about public services. Public services should work towards delivering value, not necessarily bring about revenue through selling their data. Instead, if you think of that publicly-held data as infrastructure and think about how to make it accessible, how to have standards to make it safe to access and use, we are having a much more informed debate and something that is probably much safer.

The Chairman: A last word from you, Frederike.

Frederike Kaltheuner: We live in interesting times. Technology is embedded in the fabric of our society and at the core of a lot of political concerns that citizens have. We can absolutely not afford to turn AI into a race that is about winning, because this will mean that fundamental rights are undermined.

The Chairman: That has been a very stimulating session, and some very interesting ideas came through. Thank you, all three of you, very much indeed.

Open Rights Group, Privacy International and Open Data Institute – Oral evidence (QQ 65–75)

Open Rights Group, Privacy International and Open Data Institute – Oral evidence (QQ 65–75)

[Transcript to be found under Open Data Institute](#)

Andrew Orlowski, Rory Cellan-Jones and Sarah O'Connor – Oral evidence (QQ 9–17)

**Andrew Orlowski, Rory Cellan-Jones and Sarah O'Connor
– Oral evidence (QQ 9–17)**

[Transcript to be found under Rory Cellan-Jones](#)

Dr Jérôme Pesenti – Oral evidence (QQ 201–212)

Evidence Session No. 21

Heard in Public

Questions 201–212

Wednesday 13 December 2017

[Watch the meeting](#)

Members present: Lord Clement-Jones (The Chairman); Lord Giddens; Baroness Grender; Lord Hollick; Lord Holmes of Richmond; Lord Puttnam; Viscount Ridley; Baroness Rock; Lord Swinfen.

Examination of witness

Dr Jérôme Pesenti.

Q201 **The Chairman:** A very warm welcome to Jérôme Pesenti, who is the CEO of BenevolentTech and BenevolentAI Ltd, and co-chair of the review into artificial intelligence on behalf of the Government. This is the 21st formal evidence session for the inquiry and the session is intended to help the Committee discuss the review carried out by Professor Dame Wendy Hall and Jérôme Pesenti on how the artificial intelligence industry can be grown in the UK.

I have a bit of rubric that I need to read through, so we know where we are. The session is open to the public. A webcast of the session goes out live, as is, and is subsequently accessible via the parliamentary website. A verbatim transcript will be taken of your evidence and this will be put on the parliamentary website. A few days after the evidence session, you will be sent a copy of the transcript to check it for accuracy, and we would be very grateful if you could advise us of any corrections as quickly as possible. If, after this session, you wish to clarify or amplify any points made during your evidence, or have any additional points to make, you are very welcome to submit supplementary written evidence to us. Perhaps, to start with, you would like to introduce yourself for the record, and then we will begin with questions.

Dr Pesenti: I am Jérôme Pesenti. I have been in the AI world for the past 20 years. I am currently director and CEO at a company called BenevolentAI that uses AI to accelerate drug discovery. I have had the pleasure to work with Dame Wendy Hall and co-chair of the review on AI.

Q202 **The Chairman:** Thank you very much indeed and thank you for coming to this specially organised session of the Committee. I am going to start with a very general question and then, of course, you have seen most of the questions in advance, so I hope we will be able to move through in the time allotted. What were your motivations in co-authoring the review

of the AI sector, and why did you think the review was necessary? As an adjunct to that, were you happy with the way the review was conducted? Is there anything that you might have done differently in carrying out the review?

Dr Pesenti: First, I am passionate about AI. I really think it is something that will be good for the world in the future, but it also comes with lots of challenges. AI is transformational technology that will change, in some way, the order of the world and create winners and losers—more winners than losers, I hope. It is important for countries to be aware of what is going on and prepare. I am also working for a British company, so it is in my interests that the UK is well positioned in that AI race.

In terms of the process of the review, this was quite new for me. It was an interesting choice to combine me with Dame Wendy Hall. We came from different areas. She came from academia and she has much more expertise with this kind of process. I came very green to this kind of government process but I am an entrepreneur. I like to get things done quickly, and combining it was a good approach. Overall, I was quite satisfied with the process. We tried to keep it very small, very agile and very quick. The review was written in around five months and we came up with some clear recommendations.

As to what I would have done differently, it was a little difficult to gauge the level of ambition of the review. We had very good civil servants with us, and they were very helpful, but the message we got literally changed every other day: “No, tone it down”; “No, make it more ambitious”; “Tone it down”; “Make it more ambitious”.

The Chairman: It went up and down like that. One minute, it was a bit ambitious; the next minute, it was not quite as ambitious.

Dr Pesenti: That is right. It depends who they were talking to, because there were two ministries plus No. 10. Depending where the message was coming from, the signal was different.

The Chairman: Broadly, were you able to say what you wanted to say?

Dr Pesenti: Yes, absolutely. It was a good review. It is important to think of the review as a first step. This is why we did it the way we did, not making it super ambitious or talking about concepts. We said, “What are the pragmatic steps that it should take?” Do not believe that the review is the end game. Other countries are investing a tremendous amount. I tend to believe it is better to take some clear first steps to start with, and then elaborate on it and become more ambitious as we go on.

The Chairman: You saw this as putting a stake in the ground and a first step.

Dr Pesenti: That is right.

The Chairman: I have one general question. As we have moved through, we have been more sceptical about talking about an AI sector, yet the review is about an AI sector. Does AI not permeate a number of sectors as opposed to being a sector by itself?

Dr Pesenti: It does. AI will be a little like a technology such as the internet or mobile. It will be part of everybody's business. There are some major players right now, one of them being DeepMind. There are also companies such as mine, BenevolentAI. It is an emerging sector. It is true that the bigger impact will be outside this small group, in the broader industry. Most industries will be affected by AI.

The Chairman: Would you describe an AI sector essentially as those who, in a sense, create the AI networks? Then the users are not part of the AI sector, as such.

Dr Pesenti: That is correct. If you do not do that, the sector will be everybody, because everybody will be affected. I would say it is primarily the people who are interested in creating novel AI. Academia is a big part of that, as is government, to some extent.

Q203 **Lord Hollick:** Following on from that, what were the recommendations that you were asked to tone down?

Dr Pesenti: One of the recommendations was around the number of PhDs that would be funded. It is quite important that the number down the line becomes quite high. If you look at the demand right now, it needs to be counted in the thousands, quickly, in the next decade for sure. You cannot get there tomorrow because people are not able to be trained. You need to have faculty and fellows, which we also recommended in the review. There was this question: should you put the big number first or should you start with 300? There, we got a lot of back-and-forth.

There was also some discussion around education. We looked at it as a *fait accompli*—"This is the talent we have in the UK"—but high school education is a big deal. The fact that mathematics is not a requirement for the whole time until you are 18 is a big problem that needs to be solved. We were told, "You cannot really go there. Looking at the curriculum is usually very thorny", so we decided not to do that. I know that there are some other initiatives. Education in STEM, potentially in combination with other things, is very important to developing AI in this country.

Lord Hollick: Your report has generally had a very warm reception, and the Government have now put £75 million behind it. Are there other recommendations that you believe deserve early support or greater support? What else would you like the Government to pick up and run with?

Dr Pesenti: I would say that £75 million is a drop in the bucket, but it is okay. It is a bit of a mistake to say, "Give me a few billion dollars and I do not know how I am going to spend it". This is just to initiate it. Money is one aspect. It is critical to deliver on the recommendations in a short timeframe. With government, timeframes sometimes worry me. Even right now, I am a little impatient. We should have the council set up. We had a meeting this morning around data trusts. Things are moving and the signal is very good. I would always encourage the Government to go more quickly: small, fast, and then iterate, because this is the first step.

After that, you will need to invest way more to ensure the industry is growing quickly.

Lord Hollick: The industrial strategy has included £1 billion for building out the digital infrastructure, in particular 5G and one or two other aspects. When we saw the Minister yesterday, he was reluctant to share or did not have any targets about how many premises should be linked up to fibre. We are very low down in the pecking order internationally on that. Was that something that you considered when you were doing your report? How have other countries managed to get way ahead of us?

Dr Pesenti: I would not say that other countries are way ahead of the UK in AI right now. The UK is well positioned. In terms of the level of awareness of other countries, think of the US and China. Canada is also playing a role, as is France. It is very high. You hear numbers like, “This country is going to invest billions of dollars”. I am not saying that you should imitate that because sometimes they are empty promises, but you have to know that people are all focused on this, and it is transformational. Coming from the industry, I see that it is not just a buzz but it is really happening. It will change the nature of a lot of industries out there and it is important to be positioned. It is not a question of making a big announcement now; it is a question of keep iterating and moving fast. That could really give you a leg over the other countries.

Q204 **Lord Swinfen:** How would you measure the successful implementation of your recommendations? What metrics would be useful more widely when considering the progress of artificial intelligence?

Dr Pesenti: The review, again, is a first step. There are clear metrics. For example, in terms of education, we have given some numbers, which I referred to before, in terms of PhDs, masters—also, I believe, financed by the industry—and fellowships. These are pretty clearly defined. There is the concept of data trusts that need to be set up with an organisation, and there is the AI council. These are fairly clear in terms of delivery, but this is just the first step. If you asked me how to see whether the UK is ahead, I would give you three factors that are key to measure. A nice report was done by Stanford in the 100-year study, so there are other metrics there. It came up with 2,000 metrics. I would recommend that you look back at that report.

The first is whether the UK is considered a prime place for talent. You can measure it in terms of who the main people in AI are. Are they in the UK? How many high-quality PhDs and masters have been created in the country? I would track that over time. Secondly, there is an AI sector in terms of companies that are focused primarily on AI. You can measure the amount of investment in that sector, and the number of start-ups and jobs that have been created there. I would do that.

Thirdly—this is the first recommendation—is the UK one of the best places to have access to data? I do not know whether there are such rankings today but, if you look at the data economy and whether there is the right regulatory environment in the UK to enable people to get insight from data in the best and safest way, that would be a marker. If someone did a ranking, would the UK be at the top of it? These are the

three things: talent, access to data and the economic output of the AI sector.

The Chairman: You mentioned start-ups as being one of the measures, but would you include spin-outs and the academic-industry link as part of that?

Dr Pesenti: Yes. It would be nice to see some spin-outs. The established players are there. Through my company, we had many contacts with them. They need to recognise that it will be hard for them to make this kind of transformation inside their organisations. A spin-out may be a better thing. AI sounds like it is about technology, but it is a lot about culture, and these companies do not have the culture to make these changes. It is changing the game of how you do your business, so I would absolutely include spin-outs.

Q205 **Viscount Ridley:** On this question of international rankings, the Minister yesterday was at pains to draw our attention to a new survey that came out yesterday by Oxford Insights, which put us number one in AI preparedness, followed by the USA and Canada. Have you seen that?

Dr Pesenti: No, I did not see that. I guess that is good.

Viscount Ridley: I wanted to ask whether it is a good survey.

Dr Pesenti: Unfortunately, I do not know. I will look into it and follow up. I mentioned the ranking. I am a little worried sometimes by rankings, but it is important to have this perception that the UK is the place, from the perspective of both talent and data. Today, you are going to hear me talk about talent and data. That is the place where you want to do AI. People pay attention to these rankings.

Q206 **Baroness Rock:** One of the recommendations in your review is that, "Government should work with industry and experts to establish a UK AI council to help coordinate and grow AI in the UK". As we have seen, the industrial strategy has adopted this recommendation. Could I ask you a threefold question? How do you see the council operating? Who do you think should be represented on it? What should its remit be and how broad should its remit be?

Dr Pesenti: Let me answer number two and number three, and then number one. Number two is fairly simple: it should be a mix of industry, academia and government. For government, there are people who were already involved in the review. They are the natural players there. There is a strong academic presence here, so there are also some of these players there. In the industry, I would say there are major players. Even though DeepMind is owned by an American company, it is the major player here. It is a huge asset for the UK, so it needs to be there, as do start-ups in AI such as mine—there are many today—and representatives from industry who understand the necessity to change their business. I would see this mix. Do not necessarily make it too big. You may have one or two representatives from each.

The remit is to make sure that the recommendations are followed through. That does not mean that the council will be operating—I do not think that will be the case—but it will make sure that there is always

someone assigned to each of the recommendations so that they are moving forward. If it judges that one is not moving forward, it assigns it to someone else. So they are the overseers who make sure that people are following it through. They will also start making new recommendations as time goes on. This is just step number one. What we have done needs to continue. It was not a huge time commitment to do this. I have a full-time job in the UK and the US, and I was able to do this, so it is fairly easy to consult people and continuously do it, especially if they are part of the council.

Back to your first question, they will meet regularly. In my view, there will be two sides of the council: a face of the council, let us say the CEO of DeepMind; and the people who are more operational behind it, who meet more often, get it going and make sure that the recommendations are acted on. The council would usually not take on the operation of a recommendation, but find the right person to act on it.

Baroness Rock: On the representation point, particularly in industry, do you see it having UK-centric individuals and not having an international feel to it?

Dr Pesenti: No, I would keep it UK-centric but I would put international companies on it. You could have people with an interest in the UK who are part of these companies.

The Chairman: In the same area, do you see the council following up on some of the metrics that we talked about in the last question? Should that be part of its role?

Dr Pesenti: Yes, absolutely. Do not think that we thought about everything here. The council should follow through to make sure we are delivering, based on the metrics, and should start to measure. Again, it does not have to be the council that does that. I said there was this Stanford report. It could track how the UK is doing and give a quick report, every year, as to where we are and what progress has been made in the UK compared to other countries.

Q207 **Baroness Grender:** Could you please talk us through your data trusts proposal and how you anticipate them working in practice? When talking us through it, could you particularly give us a sense of how personal data might be overseen by one of these trusts? How would someone approach it as an individual and have a meaningful say in the data and the use? Is there a wider possibility of people in the UK having a stake in that?

Dr Pesenti: Let me walk you through it. Give me a few minutes to do that. The motivation, for AI and the new techniques using machine learning, is about learning from the data. To be able to have an AI business it is critical to have access to data. Today, data is siloed into some big organisations, which have disproportionate access to it, and it is difficult for other players to have that access. Government has access to a tremendous amount of data but does not really have the motivation or the medium to share it with UK companies that could really benefit from it. That is the motivation.

The challenge in trying to solve this is that there are many types of data, many ways to share it and many aspects to the agreement, from the

liability of sharing the data to the fact that the data is constantly changing. Some data is critical and other data could almost be open. One size definitely does not fit all when you share data.

We thought about this. Instead of trying to codify the way we share our data, we should say, "Let us have government set up a small organisation that will facilitate the sharing of data among multiple parties". Think of it as a government department, transport for example, or a city that wants to optimise its transportation system and share that with a few players, including some start-ups or newcomers. Today, to be able to do that, it needs to put in a legal agreement, and that is very difficult. Most people do that from scratch, and it is very wasteful to always restart from scratch. Let us have a few people within government with legal expertise who will facilitate these transactions and agreements around the data.

I feel very strongly about this. The only way you will learn how to codify and get this concept of data trusts out there is to practise it a few times. You want to practise creating an agreement around the data, so that you understand all the parameters and all the key questions. Hopefully, there are not 1,000 aspects but there may be 50 or 100 aspects that you see coming over and over. Then you can start to templatise this kind of agreement. You will have an organisation that starts being able to create these legal templates, so that, if you are department A and you want to share it with start-up B, when it comes to you, you can say, "Here it is. Here is the template. Here are a few questions. Fill it out". That would be phase 1.

We like this idea. We call it data trusts because this notion of trust has a very good connotation. Phase two is this idea that, at some point, when you give your data away, it does not necessarily go to a company that afterwards shares it with others, but it goes to a trust that defines how the data can be used. A civic trust has very much the same concept, where you set the rules around how the money or the assets are used, and you have a trustee to supervise that it is being used like this. I believe that, down the line, this is where we will end up. You, as a person, do not give your data to an organisation. Even when you go to a website, when they collect that data, they do not put it in their organisation. They do not own that data, but they put it in a trust, where it is very visible and clear how that data will be used.

We see that people are very willing to share their data, as long as it is used for a good use. Most people want to use it for good use. It is not 100%, because sometimes it is used to overwhelm you with silly advertising, but most use of data is pretty good. If people knew what data is collected, for what and who owns it, they would be much more willing to do it, which would trigger a virtuous circle of giving my data, insight being generated, output being generated, satisfying me and making my life better. Then I am more willing to give that data.

You have to start small. We are not the first ones who have thought about this; hundreds of people have thought about it. If they think the big end goal is saying, "This is the way you share data", they will never get there. The key is: let us practise it. Let us have a few people in

government—it will facilitate these transactions, especially when one of the parties is government; understand the practice of it for a year or two; come up with some templates that codify the way the data could be used; and take that and, at some point, put it into a trust. The codification will be: “These are the rules of usage of the data when you put the data into the trust”.

Baroness Greender: What about commoditisation of that data? If you are taking data out of the NHS, is there a way of putting money back in for it?

Dr Pesenti: That is a big question around data sharing. Where is the value going? What is the value of the data and what is the value of the derived insights? Who owns it? This is where I cannot give you a single answer but it will need to be solved. I suspect that there are many areas where ideally you do not take too much into account, because often the data is sitting out there and nobody is doing anything, so you are better off giving it to someone. Sometimes you could get people to pay for that, which could be money that comes back to the owner of that data. Sometimes, that could be the Government.

If you have a healthy—I do not want to say “market”. At some point, there will be a market around data; that could be happening, but these things are a few years down the line. That is why I do not want us to start there. We should start by getting the agreement. Sometimes there will be value and sometimes there will not, but the monetisation is definitely one of the parameters.

The Chairman: The Minister, Matt Hancock, yesterday suggested that the place to develop these data trusts would be within the data ethics commission body, whatever one might like to call it. Is that the right place to develop these trusts?

Dr Pesenti: Unfortunately, we created a bit of confusion. The trust, if it is an entity in itself, which is the ultimate goal I mentioned, will sit on its own. It will not be owned by government. The organisation to facilitate the exchange could be associated with that organisation. In the report, we deliberately did not look at the ethical aspects. It does not mean there is none; there are ethical aspects, but we did not want to do a bad job of thinking about them, so we did not include them. There is this aspect.

It could be the operational arm of that organisation or it could not be. It would not offend me if it was, but it needs to be very operational. That is the key. It needs to practise the exchange of data. I had a meeting this morning to speak about ideas. There is always a tendency for government to think, “Let us define the parameters of how we do it”. I was like, “No, let’s do it. You need to do it. You need to have people to do it”. It is a very important distinction: do not just pontificate about it; do it. You will not learn before you do it. I am okay if it is paired with something that will look at a more ethical and conceptual view, but you need people who are operational and who do these agreements.

The Chairman: That is very refreshing. Thank you.

Lord Giddens: Can I say, as an individual, how much I support those comments you just made? I am struggling slightly with where the

boundaries would be. It is perhaps relatively easy to see. I support the idea but, as applied to public institutions like the NHS, you have fairly clearly bounded parameters. The digital world is dominated by five or six giant corporations. Much of our data is there. This data is going to overlap with other data within public sector institutions. You have been involved with Watson and there have been quite a lot of controversies around that, so I would be interested to hear what your thinking is on how we should deal with that issue.

Dr Pesenti: First, you would develop a model of doing this data sharing well. Interestingly, when you look at the large organisations today—Watson is not really one of them, because Watson does not collect a lot of data versus the Googles or Amazons of the world—these companies self-police. If you look at the terms of data sharing that they do, they are usually very broad. I know for a fact that they will never use it to the full extent of what is allowed, because they are very conscious of the impact it would have. They police relatively well. Smaller organisations would do it in a much more haphazard way. There are smaller organisations you have never heard of that use your data in a much more abusive way than these big guys.

What sometimes disappoints me with the big guys is that they should make very clear how they police themselves. The concept of data trusts could provide this model. At some point, when we do not really accept that they are telling us, “Trust me with your data”, we would ask them, “I trust you but you have to tell me exactly how you do it”. Maybe it goes into a trust for that. This could become a model for how to do it well. If it does well, it will not impede people from doing useful things.

When you talk about patient data, that is a very clear thing. If you ask patients, they want you to use patient data to find new treatments, as my company is doing. They have absolutely no problem. Most people would want to use patient data to do that. In a way, both parties could have a win-win situation here. The only thing is that you need to make sure that there is no abuse at some point, but you have to be careful. When you look at what happened with DeepMind, it had the best intentions—I would not question that—but it was a bit cavalier about it. In the end, it puts us back a few years, for good things that it could have done. You do have to be careful because you need to keep the trust along the way. If you are seen as gaining access when you are not supposed to, you get a backlash. I do believe that most of the players, especially the ones you know, have good intentions, so it could get to a good area.

The data trusts would be an example of what should be done and potentially add some constraints. It is not just an agreement where I give you my life, but you tell me exactly what you are going to do, because the trust has defined these ways of sharing and using it properly.

Q208 **Lord Giddens:** Why did you think the Alan Turing Institute in particular will be best placed to act as a national institute for AI? What precisely do you think its new remit should include?

I wonder if I could add something that struck me very strongly. I am a social scientist, not a natural scientist, and there seems to be a huge gap between the teaching in the digital universe and in science, and social

scientists who study the consequences of this. The only institute—there may be others but by far the best known—is the Oxford Internet Institute, which produced the report that Lord Ridley just mentioned. It seems to me that we need a further expansion of those connections; otherwise, people make these wild generalisations. They know all about the technology but they do not really know much about the social and economic areas that they are making these generalisations about. It seems quite hard to bring these two sides together.

Dr Pesenti: You asked many questions. I am going to try to answer them. These are very good questions.

Lord Giddens: There was only one, really.

Dr Pesenti: The first one is around the Alan Turing Institute. We asked ourselves this question. You have an institute called the Alan Turing Institute. Turing is one of the most recognised names in AI and it is a great legacy to what has been done here. You cannot just go around it. You are not going to create a new institute that is Turing II. It was focused on data science initially, because AI was not where it was when it was created, so it felt pretty clear. I will tell you that there is scepticism in the industry that the institute is where it should be, in terms of efficacy and delivery, so it is really important for an institute to step up.

Lord Giddens: I completely support the Alan Turing Institute, which I have supported obliquely since its inception. It is just how you link up with the core knowledge of economics and other areas, to which you are then applying the analysis of the digital world.

Dr Pesenti: The recommendation is that the institute becomes the institute for AI, not just for data science, which it is today. Its remit says that it wants to be a leading institute there, which is corraling, more on the research and academic side, the forces in the UK. Canada has created the Vector Institute in a similar way, because it has a lot of very good people and wants to have a brand around this that says, “Here are all the strengths”. If you want to create a fellowship, for example, it would be great to have the Turing name associated with that. Prestige around AI could come from this institute. That is one aspect.

There are boundaries between social science and data science and AI, which are quite important. This has multiple aspects. Most basically, it is about diversity, for example. It is really important to have a diverse set of backgrounds, genders and races that come into that field. It is essential to data because, when you have algorithms from data, they get the bias from the data. Understanding of the concept of bias sometimes comes more from people who have lived it. I agree that people coming from social science are much more receptive to this kind of problem.

It also goes the other way. Social science should use a lot more data science. It is a tool that it should use. This is one of the recommendations. In my company today, I have people with multiple skills. They are not between social science and data science but between biology and AI, or between chemistry and AI, which are very different fields. People bridging these gaps are very valuable, in many areas.

It could be part of the Alan Turing Institute's remit to extend that and not be tone-deaf to these broader issues. It is very easy to introduce it in AI itself, because it has this concept of explainability, which is a problem that you can address. It is much more of a social science problem, in terms of how people perceive a decision made by AI and how you explain it. These are interesting problems that should be studied and addressed, through a combination of social science skills and hard skills.

The Chairman: We had Dr David Barber of the institute come and give evidence to us, and one of the points that he made was that it needed to have some restructuring, whether it has that broad remit that you were talking about or not. In a sense, he was saying that the secondment model was not one that would work in the future, if it was going to play the kind of role that you envisage. Do you agree with that?

Dr Pesenti: Yes, I agree with that. A lot of my information about the Alan Turing Institute comes from David Barber as well. I have talked to him a lot. He is pretty vocal and, in my view, he is right. From our perspective, coming to the industry, it will need clearer output of what it is doing. It could be around summer schools, producing some very good research or organising some symposia. I do not know. It will have to define what its clear outputs are. It is a young institute so it needs to grow.

Q209 **Lord Holmes of Richmond:** Given the coherence and clarity of AI-related strategies in countries such as Canada and Germany, is there a risk that the UK is being overly cautious in its approach to AI?

Dr Pesenti: There is a risk of being too cautious. It has not been the case yet. As I mentioned earlier, countries are investing a lot. I do not know too much about Germany. Canada is a really good example of someone batting above their weight. It is definitely delivering right now. The way it does that is by leveraging the talent that is there and the brand name. It is starting to get some brand recognition, which is being transferred to Canada as a whole. It is very clever.

There is a lot of talent here and you need to do the same thing. You have to observe what other countries are doing. This is why we are talking about the institute. Turing is a name. You have branding organisations such as Oxford and Cambridge, which could be a little problem, because then you have competing brands. Canada is taking the lead in that direction. People have talked a lot about the Vector Institute. You have to look at what other people are doing.

It is harder to compare to China and the US, because they just have such a mass, but there are clever countries out there. The UK needs to be nimble, clever and move quickly. If you do not move quickly, there is no chance.

The Chairman: Is that the distinctive role that we can play internationally, by being rather fleet of foot, nimbler and faster than others?

Dr Pesenti: Yes, in both talent and data. Data could be another area, because you could have a very good regulatory environment in the UK for this, which would give you a very interesting advantage.

The Chairman: You think we can play a distinctive role.

Dr Pesenti: Yes, absolutely.

Q210 **Viscount Ridley:** This is about getting the balance of regulation right to achieve speed and nimbleness. On the one hand, if we do regulation of data the right way, it will encourage the sector in this country—I am sure you would agree—but, if we do it in the wrong way, too heavily handed or badly, we could end up stifling the sector. Is that the balance we have to get right?

Dr Pesenti: Yes. This is why, in the review, we do not talk about regulation at all. It does not mean that you should not have it, because having clarity around data could be very helpful, but you are completely right. If you do it wrong, you can really shoot yourself in the foot there. This is why we are suggesting that government learns through this, so that the people doing these data trusts are very knowledgeable about it, and then lawmakers who are very knowledgeable about the field craft regulation.

Viscount Ridley: It is trial and error or natural selection.

Dr Pesenti: Initially, yes, before you regulate, in terms of understanding. That is what I would say. It is quick learning.

The Chairman: Do not play into Viscount Ridley's hands here.

Dr Pesenti: The other thing I wanted to say about other countries is that you have to be careful. It is not a zero-sum game. Another play is to partner with other countries. In January, I will have a session with France, for example. I am French, so it is good for me. There is a play in having very good partnerships with other countries that are also dynamic in that field. It would be silly to think it is a purely zero-sum game and that the UK can get it all. That is why you need to associate yourself with the major players out there, recognise who they are, learn from them and collaborate.

Q211 **Lord Hollick:** Absent a bottomless pit of cash, you have to focus. What impressed us about the gentleman from Germany was that they had a very clear focus on where they already have competitive advantage, in manufactured goods: "We are going to focus our attention on making AI make us even better at that". The UK does not have a bottomless pit of cash to invest in this, unlike China or the United States. Given the current state of development of AI, the start-up companies that we have had here, such as DeepMind, and the success we have had, if you were asked to make a suggestion about where we should focus, which sectors would you focus on? Where would you use the money most tactically?

Dr Pesenti: First, I agree with you that focus is key. Sometimes, buckets of money go against you and create a lot of noise. The UK, because it is smaller, is not going to the best at everything in AI. There are some sectors that could be winners. You have DeepMind, which is one of the best labs in the world, if not the best. It is quite remarkable that it is here. You have a play and a claim to say that theoretical research in AI and intelligence as a whole could happen here. You do not want to drop that piece because it is very prestigious. You definitely want to ride this

as much as you can, having DeepMind and having the universities tag along with it. Germany could not make that claim. Very few countries, maybe four or even three, could make it.

You need to do this key, generic, fundamental AI. Afterwards, it is a question of industrial sectors. Fintech comes to mind. Health comes to mind. The advantage there would be the fact that you have a level of centralisation here, which can help you have centralisation of data and allow you to be more ambitious. Sometimes in the US, it is all over the place, so you could play there. I would not be the most competent to pick the winners beyond this. We did not want to do that in the report because we thought it would be very complicated. You would have to analyse all the industries.

Also, it is not clear: when you say “money”, money for what? You have to be a little careful there. I would encourage the universities at some point, if we look at fellowships and things, to get a flavour of AI. This fundamental view and the fact that DeepMind is here can be played. It could be extremely prestigious, and that is where you need to go. With the amazing academic talent that you have here, plus DeepMind, you can play that.

Lord Hollick: One of our frustrations is that we are often very good at basic, early research and not terribly good at translating it into the D side of R&D.

Dr Pesenti: That is true. I am more on the D side—the application—so I understand the challenges. Some people on the R side do not see that. With AI, the R side at the moment is very prestigious, and you need to continue it. You need to have this smart organisation, because most start-ups should not do the R side; they should do the D side.

Q212 **Lord Puttnam:** I have one specific question, which is about your own desideratum. Do you have one thing that you would like to recommend to government? I will come back to that in a second, because one of the first things you said concerned this confusion over the number of PhDs. I have shown your report to people outside the UK, who are very impressed by the report, but right at the end they say, “You are not as serious as you are really suggesting. If you are suggesting that the opportunity is here and universities have a very limited number of PhDs, where does that fall?”

In terms of our report, as much as possible, government has to stop being ambiguous about its ambitions. Either this is a game-changer, in which case you would resource it and find other sources of money, or it is not. You probably wept a lot more than we did at the report. First, would you like to add a nought to that number?

Dr Pesenti: Honestly, I do not remember where it landed. We struggled a lot with that number, in terms of whether to put the 2,000 in there. We put it in the text afterwards, so I believe that it is in there, but it talks about multiple thousands. Sometimes, you put big numbers and then nothing happens. In business, one way to kill a project is to make it more ambitious, so it never sees the light of day. It is really important to go to that number very quickly—in the next two years—and it is not the end

game. I completely agree. The end game is not even in the thousands but in the tens of thousands.

The understanding is that AI will change the way that people do, for example, programming. It will change the way people approach problems. It means that a completely new workforce needs to be created that tackles problems in a very different way. For that, you are talking about the tens of thousands or the hundreds of thousands. I know the disappointment and I have heard this reaction. People say, "It is very good but are you really serious?"

Lord Puttnam: It is not just a question of disappointment. This is about messaging. You talked to vice-chancellors in this country and you have the number that ended up in the report. Yes, that is a nice-to-have. On the other hand, if the report is as adventurous as I think it is, you need vice-chancellors to say, "Wow, what a fantastic challenge that is. I have to get my whole faculty in and have a serious conversation about where we are going to put the resources". This messaging thing is incredibly important but, because government looks at it internally, it is like, "I am not going to have an argument with DfE, so I will not mention that". It fails to understand that it has just gutted its own report without any intention of doing so.

One other thing troubles me: as I understand it, the Alan Turing Institute has additional responsibility for fundraising. You are French, and there is a perfectly honourable history here of the Weizmann Institute, the Pasteur Institute and MIT. We have a cultural issue in this country. The institutes do not tend to be natural fundraisers. Can this work in the case of Turing? It is a dual responsibility.

Dr Pesenti: I am not sure I am competent to answer this. I am concerned. You raised two interesting points. You can be, on the one hand, very ambitious. Should you forget about fundraising and say, "We will give you the money. Establish the fellowships"? My view is that you will need to give some money to Turing to set up some very prestigious fellowships and get the top people in. You should give more money once the key people are in. If they are not, you rethink the strategy.

Lord Puttnam: It is a top-up.

Dr Pesenti: That is right: you top up. This is why I am a little afraid sometimes, when you put in these big numbers, because you need to track what you are doing and keep going at it. You are going to ask me what my recommendation is. My recommendation is about talent. You need to go to that 10,000 number that maybe was not in the report, but you need to make sure it happens in the next 10 years and to keep topping it up. Maybe you are right that it did not tell the vice-chancellors, "We need to do something here. Maybe we should privilege this. Let us make it happen this year and, next year, let us up the number".

Lord Puttnam: That is probably partly our job, which would be great. I would welcome it. I wanted to get the number off the table. If you had one thing you would love us to recommend, what would it be?

Dr Pesenti: Maybe that is it.

Lord Puttnam: That is what I was worried you would say. I was trying to

get a third bite of the cherry.

The Chairman: We have reached the end of our session. Thank you very much indeed, Jérôme. It has been a really informative and stimulating session, so we really appreciate your time.

Dr Pesenti: Thank you for your questions.

The Chairman: Thank you very much and congratulations on putting the first stake in the ground on AI officially for the Government. We hope to put at least another stake in the ground. Thank you.

Dr Pesenti: Thank you for the support.

PHG Foundation, Understanding Patient Data, Wellcome Trust and Dr Julian Huppert – Oral evidence (QQ 116–127)

PHG Foundation, Understanding Patient Data, Wellcome Trust and Dr Julian Huppert – Oral evidence (QQ 116–127)

[Transcript to be found under Dr Julian Huppert](#)

Privacy International, Open Data Institute and Open Rights Group – Oral evidence (QQ 65–75)

Privacy International, Open Data Institute and Open Rights Group – Oral evidence (QQ 65–75)

[Transcript to be found under Open Data Institute](#)

Project Juno, Eileen Burbidge MBE and MMC Ventures – Oral evidence (QQ 46–54)

Project Juno, Eileen Burbidge MBE and MMC Ventures – Oral evidence (QQ 46–54)

[Transcript to be found under Eileen Burbidge MBE](#)

Professor Chris Reed, Professor Karen Yeung and Jeremy Barnett – Oral evidence (QQ 29–37)

Professor Chris Reed, Professor Karen Yeung and Jeremy Barnett – Oral evidence (QQ 29–37)

[Transcript to be found under Jeremy Barnett](#)

Sage, SAP and IBM – Oral evidence (QQ 76–84)

Sage, SAP and IBM – Oral evidence (QQ 76–84)

[Transcript to be found under IBM](#)

SAP, IBM and Sage – Oral evidence (QQ 76–84)

SAP, IBM and Sage – Oral evidence (QQ 76–84)

[Transcript to be found under IBM](#)

Professor Noel Sharkey, Mike Stone, Thales Group and Major Kitty McKendrick – Oral evidence (QQ 153–162)

Professor Noel Sharkey, Mike Stone, Thales Group and Major Kitty McKendrick – Oral evidence (QQ 153–162)

[Transcript to be found under Major Kitty McKendrick](#)

Professor Sir David Spiegelhalter, Professor David Edgerton and Professor Peter McOwan – Oral evidence (QQ213–223)

Professor Sir David Spiegelhalter, Professor David Edgerton and Professor Peter McOwan – Oral evidence (QQ213–223)

[Transcript to be found under Professor David Edgerton](#)

Mike Stone, Thales Group, Major Kitty McKendrick and Professor Noel Sharkey –
Oral evidence (QQ 153–162)

**Mike Stone, Thales Group, Major Kitty McKendrick and
Professor Noel Sharkey – Oral evidence (QQ 153–162)**

[Transcript to be found under Major Kitty McKendrick](#)

Professor Richard Susskind OBE, Future Advocacy and Professor Dame Henrietta Moore – Oral evidence (QQ 95–104)

**Professor Richard Susskind OBE, Future Advocacy and
Professor Dame Henrietta Moore – Oral evidence (QQ
95–104)**

[Transcript to be found under Future Advocacy](#)

Thales Group, Major Kitty McKendrick, Professor Noel Sharkey and Mike Stone –
Oral evidence (QQ 153–162)

Thales Group, Major Kitty McKendrick, Professor Noel Sharkey and Mike Stone – Oral evidence (QQ 153–162)

[Transcript to be found under Major Kitty McKendrick](#)

UK Information Commissioner Elizabeth Denham, Dr Sandra Wachter and Dr Mercedes Bunz – Oral evidence (QQ55–64)

UK Information Commissioner Elizabeth Denham, Dr Sandra Wachter and Dr Mercedes Bunz – Oral evidence (QQ55–64)

[Transcript to be found under Dr Mercedes Bunz](#)

Understanding Patient Data, Wellcome Trust, Dr Julian Huppert and PHG Foundation – Oral evidence (QQ 116–127)

Understanding Patient Data, Wellcome Trust, Dr Julian Huppert and PHG Foundation – Oral evidence (QQ 116–127)

[Transcript to be found under Dr Julian Huppert](#)

Dr Sandra Wachter, Dr Mercedes Bunz and UK Information Commissioner Elizabeth Denham – Oral evidence (QQ55–64)

Dr Sandra Wachter, Dr Mercedes Bunz and UK Information Commissioner Elizabeth Denham – Oral evidence (QQ55–64)

[Transcript to be found under Dr Mercedes Bunz](#)

Professor Alan Winfield and IEEE-Standards Association – Oral evidence (QQ 18 – 28)

**Professor Alan Winfield and IEEE-Standards Association
– Oral evidence (QQ 18 – 28)**

[Transcript to be found under IEEE-Standards Association](#)

Professor Michael Wooldridge, Professor Nick Bostrom and Professor Dame Wendy Hall – Oral evidence (QQ1–8)

Professor Michael Wooldridge, Professor Nick Bostrom and Professor Dame Wendy Hall – Oral evidence (QQ1–8)

[Transcript to be found under Professor Nick Bostrom](#)

Professor Karen Yeung, Jeremy Barnett and Professor Chris Reed – Oral evidence (QQ 29–37)

Professor Karen Yeung, Jeremy Barnett and Professor Chris Reed – Oral evidence (QQ 29–37)

[Transcript to be found under Jeremy Barnett](#)