HOUSE OF LORDS

Liaison Committee

Corrected oral evidence: Artificial Intelligence Committee: follow-up

Wednesday 14 October 2020

10.15 am

Watch the meeting

Members present: Lord McFall of Alcluith (Chair); Lord Bradley; Baroness Hayter of Kentish Town; Earl Howe; Lord Low of Dalston; Lord Tyler; Baroness Walmsley; together with Lord Clement-Jones; Lord Hollick; The Lord Bishop of Oxford; Baroness Rock; Lord St John of Bletso.

Evidence Session No. 2 Heard in Public Questions 7 - 13

Witnesses

I: Simon McDougall, Deputy Information Commissioner; Carly Kind, Director, Ada Lovelace Institute; Roger Taylor, Chair, Centre for Data Ethics and Innovation.

Examination of witnesses

Simon McDougall, Carly Kind and Roger Taylor.

Q7 **The Chair:** I welcome Simon McDougall, deputy information commissioner; Carly Kind, director, Ada Lovelace Institute; and Roger Taylor, chair of the Centre for Data Ethics and Innovation. I am Lord McFall, Chair of the Liaison Select Committee. I am joined by fellow Liaison Committee members. I am delighted to be assisted by Lord Clement-Jones, chair of the former Select Committee on Artificial Intelligence, and some former members of that committee.

This is the second of a series of follow-up sessions focusing on the key recommendations of former special inquiry committees. Thank you for participating in this new procedure, which was one of the many recommendations of the 18-month review of House of Lords committees, which reported last year.

I draw attention to the declarations of interest of Members who are participating in the meeting. I hand over to Lord Clement-Jones for the question session.

Q8 **Lord Clement-Jones:** I repeat my interests as declared in the first session this morning. I warmly welcome Simon McDougall, Carly Kind and Roger Taylor, all of whom have made a huge contribution to the ethics, governance and regulation discussion and policy-making over the past few years. They are ideal witnesses for our purposes. Thank you very much for coming.

I will start off with the first question, to do with the ethical development of AI. What is the best way to ensure the ethical development and use of artificial intelligence by government, industry and academia? You may want to unpack the difference between ethical principles, corporate governance and regulation as part of your answer.

Simon McDougall: It is a great question. If we unpacked it fully, we would be here for the full 45 minutes on the first question. In looking at how we ensure the ethical development of AI, we have great material to work with. Two and a bit years since the original report came out, which was ground-breaking and an excellent summary of where we were, we have a lot of material to move from the theoretical to the applied. I would echo what the witnesses in the first session said about how this is now a real-world technology. We are dealing with real issues, rather than abstract philosophical questions.

As you would expect, being a regulator I would highlight what we have in data protection regulation as material we can use. There is a great onus in the UK data protection regime on transparency and accountability, which is applicable across the board. Within our current regulatory framework, we have tools such as data protection impact assessments, principles of data protection by design and default, and rules on automated decision-making. All these are very applicable to modern-day AI.

At the ICO, in the last couple of years since the original report, we have sought to build on that with different pieces of guidance. In the last year, we have published our work on AI explainability, jointly with the Alan Turing Institute. As part of our AI auditing framework, we have produced very practical AI guidance for engineers and data protection practitioners. We will be producing imminently our data-sharing code of practice, which is also very relevant to the world of AI.

This is a very fast-evolving area and we cannot rest on our laurels. We have to keep going with what is a fast-moving challenge, but we have a body of work to apply. That is critical if we are to maintain public trust and confidence in how we are using AI.

Lord Clement-Jones: Do you need a broader remit to explicitly include AI?

Simon McDougall: That would be challenging right now. Defining what that broader remit actually means is quite difficult. If we look at the wicked questions, use cases and challenges we are seeing now, most of the time they include the processing of personal data, in which case data protection is already coming into play. We find that we are very actively engaged in questions of AI. I am sure we will come to Covid and our experience over the last six months in the next 45 minutes. There, we have found ourselves being very actively involved in questions of personal data processing. We are not feeling that there is a really big gap to address in how AI is being used currently.

Lord Clement-Jones: Carly, you have been very prominent in the black box debate. Where do you sit on this?

Carly Kind: Thank you for having me. There is no one answer to this. No country out there has yet solved this question about how we ensure ethical development of AI. It is a process that we are involved in. The UK is very well placed among other countries to be developing this process.

Since the House of Lords began its inquiry in 2017, 100 different ethical codes and guidelines have been published. As the previous witnesses and Simon said, we are a stage now where we have general consensus on the ethical principles that apply to AI. We need to work out how to apply them in practice and shore up public trust and confidence along the way.

That is quite difficult, because many of these ethical principles change in different cultural contexts and have no one universal meaning. We have seen this in recent months with the A-level results. There are many different ways to define fairness, for example, or to conceive of justice. We need to involve the whole of society in the conversation.

We have learned some lessons about applying ethics and ensuring ethical development. First, discussions about ethics in abstract are not helpful. It is more helpful to look at use cases and applied examples of AI and understand the ethical issues they raise. We know a lot more now than we

did two years ago about facial recognition technology and the ethical issues that arise, because we are looking at the use case specifically.

We know that self-regulation and internal ethics processes have not kept up and have not proved to be sufficient to ensure accountability and public trust. In the Ada Lovelace Institute's work on public deliberation, we hear time and time again from members of the public that their trust in technologies is contingent on external oversight of those technologies. In both private sector and public sector tech development, they want to see external oversight, not solely internal self-regulation. Regulators will be playing an increasingly important role. Independent oversight mechanisms, ombudsmen and other external scrutiny measures are really important.

We have learned that the ethical development of AI is not just a technical question, but a sociotechnical question. Trust is not just about trust in the tech itself, but the system the tech is embedded in. We need to think about not only making the technology comport with ethical principles, but the humans using the technology and the system as a whole.

We have learned that we need to involve members of the public in these ethical processes. One of the most successful examples of ethics committees in action is the West Midlands Police ethics committee, which has proven to be really active in engaging with issues of technology. They have included lay members of the public in their ethics committee to deliberate on these ethical challenges. That has proved to be a really good model.

I am sure we will talk shortly about other process to ingrain ethical development. Wendy Hall referred in the previous session to developing tools and methodologies for ensuring ethical development, such as audits and impact assessments. The Ada Lovelace Institute is also very intent on that. It is a really rich area of research, exploration and application to be pursued. I can speak more about that, but that is the way forward on ethical development.

Lord Clement-Jones: Roger, the CDEI published its barometer not that long ago and you are getting into that sort of risk evaluation side now, are you not?

Roger Taylor: That is exactly right. To build on what Carly said, this absolutely is a process. In the way we look at it, we see four key things that we need to achieve. The first is to monitor how AI is being used and identify the barriers to its adoption and where there are gaps in regulation that can lead to misuse. As you rightly pointed out, the AI barometer, which was produced with civil society organisations, industry and regulators across different sectors, addresses exactly that.

From that, we can pick out three particular areas that get in the way of the ethical adoption of AI. The first is the need to educate and engage with the public to understand what is acceptable in the way these technologies are used. This is particularly important in areas of government use of these

technologies. A principle that is often referred to in this area is human centricity. There is significant work to do in this area to understand how we make these systems so that people feel comfortable with the way they operate. We did a lot of work, for example, on microtargeting. We brought together groups of people, explained how it worked and then got their sense of what was okay and not okay in that area.

Next, I would point to gaps in regulation. There are areas where the law needs to be adapted. Social media is one and we have pointed out that there is a gap there. Much more often, we are talking about the situation where we have existing laws and regulations. The problem is that nobody quite knows how to apply them in this new context.

A good example of that would be the work we are doing on bias. That is helping organisations to identify how to conceive of bias, where there are many different definitions you can use, and how to measure it, in a situation where you will definitely find your algorithm will fail a certain measure of bias. As a responsible organisation, you have to be transparent about how you are making those decisions and be confident that you are making them appropriately. It is a real barrier to adoption if people do not feel they know how to operate responsibly in this area.

If we want ethical and responsible development, we have to ensure that organisations with appropriately skilled and diverse workforces can access high-quality data. There is a significant barrier and uncertainty as to how people can safely use data in a way that respects and protects privacy but still allows it to be used for the public good. We have quite a full agenda, but these are all issues that the ICO, the Ada Lovelace Institute and the Centre for Data Ethics and Innovation are working on and making some progress on.

Q9 **Baroness Rock:** Each of the witnesses has talked about regulation. I would like to dig a little deeper into this and ask you whether the regulatory framework and the extent of governance are fit for purpose for the pace of development of AI—we all know it is very fast—as well as addressing the potential implications of the use of technology. Some of you have touched on whether more or less regulation is needed. How do we decide where regulation is appropriate?

Roger Taylor: This technology is developing very rapidly. It is no surprise that regulators and regulation have to move at speed to catch up. That is inevitable when a very new technology arrives on the scene. There are areas where we can see very clearly that regulation needs to be updated. I would point to social media as a very new thing that requires our media regulation to be adapted to deal with it.

As I was saying before, the more significant gap is understanding how to make sense of our existing laws, regulations and ethical standards. Knowing what is acceptable and how you tell whether your data-driven systems are behaving acceptably is a challenge for organisations. It is a real barrier to innovation and development.

We are doing a lot of work on AI assurance, about how an organisation can go through a process. As Daniel was saying in the previous session, we need to be really careful to avoid a situation where, if you pass a test, you can rely on your system being good to go. This requires careful monitoring; sophisticated analysis of the behaviour of these systems; and boards and accountable organisations having the skills to understand how their algorithms are behaving and whether they are behaving acceptably.

Regulators also need to upskill. There is variation between regulators. Some of them are moving at pace and addressing this. Simon has talked about the work of the ICO; there is the Financial Conduct Authority and others. In other areas, there is much more to do. It varies between industry associations. Some are doing a lot of work in this space and some are doing less.

We have very good foundations in this country. We have strong regulators and a culture of anticipatory regulation that organisations such as the ICO have led on. We have a good legal framework underpinning how we address this, so we are in good shape, but there is a lot of work to do.

Baroness Rock: Simon, as Roger has touched on the ICO's work, perhaps I could ask you to expand on what the ICO is doing in this area.

Simon McDougall: It is in the vein that we have discussed so far. The regulatory framework itself, while there is always room for improvement, is broadly applicable to the challenges we are facing with AI. It is about pushing on to help organisations apply this in practice. I would draw the Committee's attention to our AI auditing guidance, which came out earlier this year. It recognised that, in terms of good practice, organisations need to take a range of measures throughout the development life cycle of building an AI system, concerning how the business case is visualised, how the coding works, the implementation and, importantly, the measuring of outcomes afterwards.

Going back to a point Professor Hall made in the first session, it recognises that one aspect of a lot of machine learning is that the algorithm will continue to evolve after the initial launch. Any system of governance and controls has to recognise that and be an ongoing process, rather than one stop in time. Our guidance was intended to enable organisations to work through all those stages and have practical measures to address that. That is important going forward.

In our view at the ICO, AI may end up being a little like cybersecurity. While we get very interested in the bleeding edge of these technologies, and it is important to be up to speed with those, in most of the security breaches we deal with there are obvious flaws that could have been dealt with a long time ago. They are unpatched servers or social engineering. These are basic errors. Over time with AI, we may see some of the worst situations arise from missing basic steps along the way, rather than more theoretical, interesting, cutting edge questions. We need to get the basics right.

Baroness Rock: That is incredibly helpful. Thank you.

Carly Kind: I have a couple of thoughts. There are a number of limitations with the existing regulatory framework, which is not to say that we need to rush to legislate now. We are still understanding where regulation would be useful and appropriate. I have a few observations. There are some very specific areas in which the regulatory framework is clearly not fit for purpose. The Court of Appeal has recently ruled that facial recognition technology is one of those areas where there is not a sufficient legal framework.

Roger alluded to the use of AI on online platforms and social media. That is clearly an area in which the use of AI is causing harm and creating ethical risks. Self-regulation has not sufficiently kept up with that. Some of those risks are being dealt with within the domain of data protection by the ICO. Some of them are dealt with by the CMA. Those regulators are showing a real willingness to work together, but there is a question about whether we need some type of framework to enable a more overarching look at online platforms and the use of AI there.

More broadly, I have a question, which is what role regulation could play in shoring up social licence for these new technologies and creating a sense of social responsibility on the part of both public and private sector entities when deploying them. This is not about an enforcement regime. It is about saying that AI development and research should be allowed to flourish unimpeded in an innovation environment. Before it is deployed on the public at large there needs to be some kind of quality assurance, circuit breaker or mechanism to validate that a piece of technology is ready for public deployment.

These technologies are being beta tested on the public, without any sense that they are causing harm in the process. I am not advocating that we create a regulatory bottleneck for new technologies, but there needs to be some process for giving them the go-ahead. Think about self-driving cars, facial recognition or AI in healthcare. In order for these to make a difference to our society, they need to enjoy indelible public trust. Unless they get a stamp of approval through a regulatory mechanism, I worry that that will not happen and their benefits will not be realised. That is something to think about when it comes to regulatory interventions.

Q10 **The Lord Bishop of Oxford:** It is really good to see Simon, Carly and Roger. I have worked closely with Carly and Roger and learned a great deal from them over the last couple of years.

Do you think the balance is right in the UK at the moment for the use of data, between the privacy of the individual and the public good? If not, how should we address that? Do you think, as Dr Susskind was saying in the last panel, that balance is shifting during the Covid crisis? As a final question, can you comment briefly on the proposed national data strategy and how it addresses that question of private data and public good?

Carly Kind: I am sure you will not be surprised to hear me say that the idea that you can have either individual privacy or public good is a false

dichotomy. We have learned, particularly during the pandemic, that use of data for the public good can be done only with public trust and confidence, which comes from these data rights protections. Our own public deliberation research at the Ada Lovelace Institute shows that people do not generally feel there should be a trade-off between privacy and the public good. This idea of privacy versus the pandemic did not resonate with normal people. They were happy to see their data used in the pandemic response, but they also wanted guarantees of their data rights and privacy.

We also convened a round table with the Royal Society to look at lessons learned from Covid with respect to data sharing. The participants in that did not generally raise individual privacy protections as a barrier to data sharing or data access during Covid. They raised concerns about the absence of data, standardisation of data, interoperability of systems, a lack of political will, and cultural gaps that stopped data being shared, but not individual privacy concerns or the GDPR. The GDPR has stood up relatively well under the pandemic conditions.

I would reframe the question slightly away from individual privacy versus public good to the balance between individual privacy and public good on the one hand and private gain on the other. That is where the balance often does not get struck correctly. When we asked participants about their experiences during Covid, we saw that one of the barriers was the inability for the public sector to access private sector data.

At numerous points during the Covid response, researchers have argued in journals such as *Nature* that they need access to private sector data to do the type of research needed on the pandemic. For example, tackling misinformation on Covid online was stymied by the inability of researchers to access platform data. It remains the case that most data derived from the public's behaviour and the public's consumption is used for private gain and not for the public good. In fact, it is in the hands of a few companies, mostly American companies, not in the hands of the public sector.

This is a really important issue from the perspective of AI development. As you know, the future development of AI is forged by access to three things: computing power, skills and expertise, and data. In particular, at the moment, machine learning is very data intensive, so those who have the most data will build the best AI. That is why the leading AI research is coming out of companies such as Google, Facebook, Amazon and Apple: because they have access to the most data. Their monopolistic positions are reinforced by their advancements in AI; then they can attract more skills and more expertise.

This issue needs to be addressed. The development of AI in the UK is very connected to access to data and the monopolistic control of personal data by tech platforms. If I have any critique of the national data strategy, it is that it equates encouraging innovation with reducing regulatory burdens. To enable innovation, we need to use regulation in a targeted way to break up the monopolistic control of online platforms, and therefore data, and create a more even playing field, where new entrants, UK companies and the public sector can compete with these larger companies.

Roger Taylor: I agree with a lot of what Carly said, but there is a significant problem here. We are not using data as effectively as we could for the public good. Privacy is one of many reasons for that. The CDEI published a report on data sharing in the public sector. I agree with Carly's point that it is often cultural, but there is real fear and uncertainty about these issues. Carly is right that we do not need to see this as such a stark trade-off. One way of doing that is to make full use of privacy enhancing technologies, such as data trusts, that shift control over access to data in a way that is much more trustworthy and privacy preserving. That would greatly reduce the trade-off and we do not do nearly enough of that.

However, it is impossible to reduce this trade-off to zero. We must avoid suggesting that there is literally no trade-off here. The privacy lobby quite rightly points out that that is simply not true. This is where we run into problems about very fragile public trust. It is absolutely right that the public would like this data to be used for the public good. They are very clear, in research effort after research effort, that they want this to happen. They are more concerned about the public benefit than about privacy protection.

However, we often fail to communicate clearly and unambiguously how the data is being used and how privacy is being protected. That can often result in a confused public debate that undermines public trust and means that the people involved in developing these applications lack confidence or become risk averse, for fear of finding that what they are doing does not have public legitimacy.

This goes broader. We have strong regulation on data protection in this country, but many organisations are trying to navigate those laws, common law, confidentiality duties and a range of specific legal codes applying to particular bits of data they are dealing with. To be clear, the ICO does really great work in trying to clarify this, but it is enormously complicated. It is a serious barrier to organisations being able to derive public good from data and feeling confident as to how to navigate this space.

Simon McDougall: There are so many points to build on there, from both Roger and Carly. I agree with Roger that having public trust and confidence is fundamental to enabling innovation. People who are not comfortable sharing their data with new products, services or organisations will revert to traditional channels and established players. That is bad for competition, innovation and growth.

I wanted to drop in a tiny bit of good news on this, going back to the points on Covid. I lead the ICO's work on Covid as well as my day job in innovation technology. As we engaged with organisations supporting different efforts to fight Covid, we found a very high awareness of and engagement with privacy risk. People understood they needed to get privacy right and maintain public trust and confidence. They often needed help to do it, but they managed it well. We found that there was a refreshing change from the situation a few years ago. Awareness in organisations has improved. That said, there is still a long way to go. We need to see how we can apply

the crisis-driven lessons we have learned from the Covid pandemic to business as usual, as, in time, the pandemic recedes.

Q11 **Lord St John of Bletso:** I am glad to see you again. We have just discussed public trust and understanding. The public perception is that there is no unified approach on governance of AI across government departments. Do the departments, agencies and bodies involved in the development of AI policy in the UK work together effectively? If not, what can be done? We are aware of the Regulatory Horizons Council, but are the methodologies, tools and techniques ready for us to formulate regulations that can be rigorous, sound and effective?

Carly Kind: Roger and Simon might be slightly better placed to speak about collaboration within government. I would just make one observation from research we have done looking at the use of AI and algorithmic systems in local authorities. There is definitely a disconnect between central government adoption of principles and promulgation of standards for AI, and local government ability to put those into practice while using AI and algorithmic systems. There is a resource and capacity disjoint. Local authorities feel that they do not have the support and guidance they need to implement AI and algorithmic systems ethically. This disconnect between central and local government needs to be addressed.

Simon McDougall: It is a great question. I think so far, so good. A number of bodies have been formed relatively recently, since the original report came out. As the regulator, we find this an incredibly rich area to work in with these different organisations. Very recently, we had an excellent session, jointly with the Ada Lovelace Institute, on Covid, health certificates and immunity certificates. We managed to combine Ada's amazing network in the research world, the scientific, academic and medical communities, with people from government and regulatory bodies to get into the detail of that. We need more of that.

The situation will carry on evolving. This is a fast-moving area with a range of groups involved, including the recently formed Regulatory Horizons Council, with which we have had great initial conversations. We are looking forward to working with it more in the future.

Roger Taylor: I would answer your question on two levels. First, there is obviously a range of institutions: the ICO, the CDEI, the Office for AI, the AI Council, the Regulatory Horizons Council. Our communication and co-ordination efforts are going well. We are all talking to each other regularly, as well as to civil society, including Carly, to co-ordinate how we operate as far as possible. That is good. As this develops, we shall see over time what the long-term institutional landscape needs to look like.

You asked more broadly about the use of AI across government and Carly pointed to local government. I agree that organisations there are not operating to uniform standards. There is an issue about clarity as to the public sector approach to this. The data ethics guidance has been published, and it is moving in the right direction. The CDEI is working with

local councils. We are working with Bristol, but this is a very steep learning curve for organisations.

These are very basic concepts. How do you judge whether something is in the public interest? How do you ensure that it is doing what you said it was going to do and that the public accept it as a reasonable use of data? Organisations are developing approaches to this. We see our role at the Centre for Data Ethics and Innovation as trying to bring a consistent methodology to how organisations do this and a degree of consistency to how government operates in this space. You would be right to say that there is a significant variation in practice at the moment.

Q12 **Lord Clement-Jones:** You both mentioned the Regulatory Horizons Council, but you did not say where it fitted in the firmament. We would all be rather curious about that. Simon, would you like to unpack that?

Simon McDougall: It is early days for the RHC, but it is important to recognise that the digital landscape overall is evolving very rapidly. That is something we need to keep on top of. At the ICO, jointly with the Competition and Markets Authority and Ofcom, we set up the Digital Regulation Cooperation Forum. That is a recognition of the fact that data protection, competition and consumer protection, telecoms regulation and online harms overlap far more than they used to. We have borders and interplay that we did not have even a couple of years ago, which we need to work through.

We need to keep working on that, because the digital landscape is moving very quickly. One thing that Covid has done in the last six months is make real how reliant we are on our online world and online services, whether for getting goods or for connecting with people. The regulators in this space, as well as all the other bodies we have listed, have to keep working on this, because it is not a still picture. It is a very fast-moving environment.

Lord Clement-Jones: They are not duplicating your work, Roger.

Roger Taylor: No. We have had good initial discussions. I am sure everyone knows this, but, just to explain, the Regulatory Horizons Council is across all industries, whether it is 3D printing, genetic engineering or self-driving cars. It is looking at whether our current regulations need to change. The Centre for Data Ethics and Innovation—this is the way we have discussed it—is looking specifically at the use of AI. We are interested in the regulation of that, but we are also interested in how organisations govern it and how audit and third-party assurance are done. It is worth remembering that the majority of governance takes place within organisations, not in the organisation-regulatory interface. There are issues of public acceptability and education.

There is a very clear overlap between what the CDEI is doing and what the Regulatory Horizons Council is doing, which is where regulation and AI meet. We are co-ordinating our efforts to make sure that we neither

duplicate nor find ourselves pulling in completely opposite directions. I am confident that will work.

Lord Clement-Jones: You are not policing the boundary, so to speak.

Roger Taylor: We are working together on those issues; that is the right way to put it.

Q13 **Lord Hollick:** In November 2017, six months before we published our report, the Government published their industrial strategy, which placed AI and data as one of the four grand challenges. The mission was to embed AI across the UK, which will create thousands of good-quality jobs, drive economic growth, improve the functioning of government and improve healthcare. To what extent and how well have these objectives been achieved? How well organised are the Government to achieve these objectives?

Roger Taylor: There is a real benefit from greater co-ordination of our efforts in this space, through national leadership and the identification of clear objectives. Those may be sectoral, as with the industrial strategy. They could be in other sectors, for example the use of AI in education. Having clear, central government-set objectives in that regard is useful. There is also a real role under the industrial strategy for setting direction on cross-sector issues, such as data management, data trusts, ID, how digital identities work and how privacy enhancing technologies are adopted. That is another area where national leadership can work.

How well are we doing at delivering on this? We are making good progress, but I would refer to the comments made in the previous session. This is not a short-term game; it is a long-term game. We are in the foothills of getting this sorted and we have a huge amount to do. There is good progress, but the task ahead of us is probably much greater than what has been achieved so far.

Carly Kind: There has been less progress than ideal on the ambition, in the industrial strategy and the House of Lords report, for the UK to lead the world in safe and ethical use of data and AI. There has been a change in tone from government on that. We were hearing much more about that ambition a year ago than we are now. Perhaps that is due to the very changing circumstances. It also seems to be a shift in the rhetoric about where the UK wants to position itself in the global race on AI. There was this notion that we would have a third way of ethical AI and that the UK and Europe would pursue that. That does not feel like it is at the top of the agenda any more, which is a shame. It remains an area the UK can demonstrate leadership in.

If we look at other countries, Canada, for example, has now instituted a duty on public bodies to undertake an algorithmic impact assessment before deploying or procuring algorithms or AI in the public service. That is the type of intervention that makes a world-leading country in ethical AI. Going forward, we would like to see that prioritised more and that commitment in the industrial strategy returned to.

Lord Hollick: Simon, to what extent could the Government be using AI to help them make better decisions, for instance on how to manage the pandemic?

Simon McDougall: That is a fantastic question, because I do not think we have seen a huge use of AI in the first six months, or more now, of the pandemic. In hindsight—hindsight is a wonderful thing—that is not too surprising. We were dealing with a very novel situation, so it was not as if we had a huge amount of training data to utilise in the case of machine learning. In a crisis, everyone falls back on tried and trusted methods. We saw a lot of data use. We were involved in interesting discussions as to how best to share aspects of data with different parties to address the pandemic, whether it was shielded persons, the contact tracing app or customer logs. We did not see machine learning.

Now that we are gathering more data as to how Covid works and how our society works during Covid, AI may come to the fore over the next six months, as we have had time to regroup and use more innovative methods. I would echo Roger and Carly's point that, in doing that, it is critical to maintain public trust and confidence. I have had conversations with both private sector and public sector leaders in the last two weeks, where different pieces of research have pointed to the take-up rate of new technologies being driven by concerns about usage of the data, privacy and, again, trust and confidence. It is key that we have a regime where people feel they are protected and can therefore engage in new services. Otherwise, people will turn off the new AI-driven services we hope to see and we will be stuck where we are now.

Lord Clement-Jones: There is a common theme of public trust. Lord Hollick, when you put your question to Ministers, you might wrap up public understanding and public trust together. I think that would be very useful. Thank you very much indeed. That has been a really insightful session and we have picked up a lot.