HOUSE OF LORDS

Liaison Committee

Corrected oral evidence: Select Committee on Artificial Intelligence: follow-up

Wednesday 14 October 2020

9.30 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. 1 Heard in Public Questions 1 - 6

Witnesses

I: Professor Dame Wendy Hall DBE FRS FREng, Regius Professor of Computer Science, University of Southampton; Professor Michael Wooldridge, Professor of Computer Science, University of Oxford, and Programme Director for Artificial Intelligence, Alan Turing Institute, London; Dr Daniel Susskind, Fellow in Economics, Balliol College Oxford.

Examination of witnesses

Professor Dame Wendy Hall, Professor Michael Wooldridge and Dr Daniel Susskind.

Q1 **The Chair:** Welcome. I am Lord McFall, Chair of the Liaison 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 first of a series of follow-up sessions focusing on 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 the House of Lords committees, which reported last year.

I draw attention to the declarations of interest of Members who are participating in the meeting, and I will hand over to Lord Clement-Jones to question the witnesses. This will be done in three evidence sessions, the first of which is on the present and predicted impact of AI. The witnesses are Professor Dame Wendy Hall, Professor Michael Wooldridge and Dr Daniel Susskind.

Q2 **Lord Clement-Jones:** Lord Chair, I thank you and the Liaison Committee for making this the first of these innovative follow-ups. Before I welcome our witnesses, I declare my interests as a consultant to DLA Piper, a chair of the council of Queen Mary University, chairman of the board of the Ombudsman Services, and independent consultant to the Council of Europe Ad hoc Committee on AI.

I warmly welcome Dame Wendy, Professor Wooldridge and Dr Susskind. Two of you have appeared before us in the past, and I reprise a question that we asked three years ago which you will be familiar with. What are the biggest opportunities and risks associated with artificial intelligence over the coming five years? We will be able to compare and contrast at least two of your previous answers.

Professor Dame Wendy Hall: Thank you very much for inviting us back. It is a pleasure to see you again. Thank you also for the report, which had a huge impact when it came out. We have done many shows together since then on the roadshows.

Keeping it short for now, we see the opportunities even more clearly, particularly during the last six months. The impact of AI on health and well-being is well documented now. We have used that successfully and not so successfully during the crisis. Many companies are using AI, big and small; we can talk about that later. We have seen growth in the sector.

The risks are that we do not keep up the impetus as a country. We are third, after the US and China, in the global index of AI. Our ambition is not just to remain third, but to be a real player in the world in this. There are many things that we can lead on in the UK; we can pick that up later. The biggest risk to us is that we do not keep up the impetus. The Government funded a number of skills programmes. They have been successfully launched, but we need to keep that going. There is not just a money

impetus in this; we will have to keep that steady and increase it, if anything.

Lord Clement-Jones: Professor Wooldridge, you have written two books in the meantime. What is your perspective on this?

Professor Wooldridge: My perspective is much the same as Wendy's and I endorse everything Wendy has said. Today is an interesting opportunity to revisit what we thought was going to happen and what has actually happened. I would underline the predictions we made in 2017 as still being my predictions.

I still believe that healthcare is an immediate opportunity. We have seen a lot of advances in healthcare, as Wendy mentioned, since the initial report. Interestingly, to transfer AI techniques from the laboratory to GPs or practitioners using them in hospitals is a long and slow process if you do it right. You should not expect to see a sudden explosion in this. It is great to see a lot of systems coming through the process of getting regulatory approval. For me, healthcare remains the most exciting opportunity in AI. I believe that the UK is at the forefront of that.

I would identify two risks. This probably sounds like a record stuck on repeat, but one is data and privacy. Since we originally met, we have seen endless examples, every week, of data abuse. Here is the thing: for current AI techniques to work with you, they need data about you. There is an incentive not just for these companies to take data and sell it to somebody else. They need your data for the AI to work. That remains a huge challenge. Society has not yet found its equilibrium in this new world of big data and ubiquitous computing. We are still finding our way with that. That is the No. 1 risk.

The No. 2 risk I would highlight is complacency. You build an AI system to diagnose cancers or something like that. It looks like it is doing the right thing, so you immediately conclude that it must be a great program to do that, but this technology is really brittle. If we naively rely on it instead of relying on human judgment, there are huge risks. AI can be another voice in the room, not something to displace human judgment. It is very important not to become complacent about the technology.

Lord Clement-Jones: Dr Susskind, you have also written a book, but one that is rather more focused on the impact on employment and the jobs market. What is your take on the future?

Dr Susskind: That is right. I want to emphasise the data privacy and security risk that Michael mentioned, particularly in light of the current pandemic. Before the pandemic, there was a very lively public debate about issues of data privacy and security. That was reflected in your report. At the start of the pandemic, a "do what it takes" mentality took hold with respect to developing technologies to help us to track and trace the virus. Technology companies around the world were given huge discretion to collect smartphone data, bank account statements, CCTV footage and so on in a way that would have been unimaginable eight or nine months ago.

That was not necessarily a bad thing. The moment required it to help us to control the spread of the virus. There is an important task in the months to come, once the pandemic starts to come to an end, in reining back the discretion and power we have granted to technology companies and, indeed, to states around the world. That debate over the last few months has, understandably but slightly troublingly, fallen away from the intensity with which it was taking place before the pandemic began.

Lord Clement-Jones: We will certainly be picking that up later.

Dame Wendy, you talked about impetus, but you did not unpack it hugely. Do you agree that the data issue is something yet to be resolved? I know you are a big fan of data trusts, for instance.

Professor Dame Wendy Hall: When I talked about impetus, I meant for funding more skills. We are seeing a rollout, and we welcomed the Prime Minister's speech two weeks ago about lifelong learning and adult learning in this digital and AI space. The impetus I was talking about was keeping a rolling programme going for master's, PhDs, fellowships, so we keep our workforce and start upskilling everybody.

I totally agree with Michael and Daniel about data. [Dr] Jérôme [Pesenti] and I made data trusts the first recommendation in our review. As has become clear with Covid, there are issues in how we use personal data to do what companies and government need to do to analyse situations and to develop AI. How companies share data, where the asset is recorded and the legal and ethical framework in which we share data in all circumstances are major societal issues. The UK is in the lead in this space. We need to keep the impetus up there as well.

Q3 **The Lord Bishop of Oxford:** Good morning, Wendy, Michael and Daniel. It is really good to see you again in this context. I declare two interests: I am on the board of the Centre for Data Ethics and Innovation and I am doing some work with the Ada Lovelace Institute in its rethinking data project.

As you are aware, we stressed ethics in our report. It seems to me that society has been on a learning curve about how to develop ethical AI. The complexity of that grows year by year. What are the greatest barriers to the ethical development and deployment of artificial intelligence now? How can any such barriers be overcome?

Professor Dame Wendy Hall: It is good to see you again. I have shared your 10 commandments on this issue with a lot of people over the last few years. It was a very good list. In the UK, we have a lot of people studying this issue. Our big universities, Oxford included, have big programmes in this area, some of which sit with the philosophers, as ethics should.

We also need to develop practical guidelines. These are major issues for society that take a long time to think through. I like to make the point, which will probably resonate very well with you, that morals and ethics are not the same thing. We cannot expect government to do everything here.

We have to self-regulate. We have codes of conduct and moral codes and we have to get people to understand where their responsibility lies.

On regulation, I know you are hearing from Roger Taylor from the CDEI and Carly Kind from the Ada Lovelace Institute, which I now chair. That is a fantastic job to have. They are looking at how to advise government on this. We have to think very practically about how we get companies involved in this. "AI will be used for good" is a fabulous statement to make, but it is quite meaningless in terms of what will happen. We need to develop quite simple frameworks and audit arrangements for companies using AI that can be very simply applied.

There are whole new career prospects in this for all sorts of professionals—the actuaries, accountants, lawyers. Education is important, but we have to focus on how we practically help companies, large and small, to audit their algorithms and the data for bias. I talk about FATE AI these days. I got it from someone in Canada: fairness, accountability, transparency and ethics. It is just FATE. That is my latest mantra. We have to work on FATE AI in the UK.

The Lord Bishop of Oxford: That is really helpful. Thank you.

Professor Wooldridge: This is another interesting opportunity to reflect on where we were in 2017 and where we are now. Again, I agree with Wendy. Assume that I agree with everything Wendy says; that will make life easier.

The good news is that the debate has moved on a long way since 2016. There has been a huge amount of water under the bridge and a lot of discussion. Many codes of conduct are now released by companies and organisations such as the IEEE—the Institute of Electrical and Electronics Engineers—that have hugely informed the debate.

There is a lot of awareness of these issues, and I sense that from my students. When I talk to my undergraduates now, they are really engaged with the issues of ethics in a way that was unthinkable 20 years ago. You could not have dragged a student into a lecture on ethics 20 years ago. There is a complete change in attitude, so they are demanding it. That is very welcome.

There are still some barriers. The two headline barriers for me are people and naivety. This echoes what I said earlier. We can build an AI system, to use a UK example in Durham, for making judgments about whether people should be held in police cell custody overnight. You build that system and it looks like it makes the right decisions, so you start to rely on it. That complacency, the assumption that the technology must be doing something better than a human being, is very dangerous. It could be doing some things better than a human being, but we need that human in the loop. Complacency and naivety are really important issues.

We need to be vigilant. The debate has not ended. We have many codes of conduct now. Many people are actively researching these. This is an

ongoing debate and we need to stay on top of it. We need to take to task the people who abuse these things and show that we really mean it, with respect to ethical developments of AI.

Dr Susskind: I have two observations, which echo some of the comments so far. One is the importance of the shift towards providing practical advice and guidance to the companies and engineers developing these systems and technologies. If we are honest, the finest computer scientists, although they might be technically incredibly able, are not necessarily hired for the sensitivity of their moral reasoning. Therefore, often through no malign intention by the designer, these systems have ethical and moral consequences.

Taking the growing body of thinking in academia about the ethical development and deployment of AI, and making it practical for companies and employers to use, is very important. That is exactly what has started to happen over the last few years, and we need to accelerate it. There is another part of the story—the users themselves. There is also a burden on companies and engineers to make these technologies as transparent as possible, so that users can scrutinise them, understand how they work and the consequences they are having.

One development in the world of artificial intelligence is that, 30 or 40 years ago, these systems very closely followed human reasoning. They were often based upon the explicit rules that human beings follow when they solve a particular problem. For that reason, it was very easy to understand why a system gave you a particular medical diagnosis or parole decision. In the way the systems work today, they are far more opaque; they are far less transparent. This point has been made many times, but another part of this is allowing and empowering consumers, users and recipients of these technologies to scrutinise them effectively.

Q4 **Lord Hollick:** I would like to declare the following interests. I am a director of Honeywell International, I chair the independent advisory board on a UKRI project into AI at the University of Bath, and I sit on the Royal Society advisory board. It is very nice to see the three of you again.

I want to move on to jobs. We took a great deal of evidence during our inquiry and there was a general acceptance that the advent and implementation of AI would cause widespread disruption and changes within the labour force. Estimates ranged from 10% to 50% of jobs being at risk to AI. That was three years ago. This year, we have seen a rapid move from the analogue world to the digital world as a result of the pandemic, home-working and all that goes with that. Dr Susskind, could you give us an update on your assessment of the impact on the labour market?

Dr Susskind: The short answer is that, for various reasons, the pandemic has increased or is likely to increase the threat of automation. There are various reasons for this. One is less to do with the fact that we find ourselves in a pandemic and more that we find ourselves in a recession.

Evidence suggests, particularly in the US, that recessions are the moment when automation often picks up.

The distinctive feature of labour markets around the world is that they have been hollowed out over the last few decades. Low-skilled and high-skilled jobs have grown as a share of employment, but middling-skilled jobs have declined. Why? One popular thesis is that technology has eaten away at the middle of labour markets around the world. If you look at when those jobs were lost, though, in the US about 88% of those middle-skilled job losses took place within one year of a recession. The evidence does not apply equally well around the world, but it is something to bear in mind. That is just one observation: that we find ourselves in a recession.

Another reason, more directly related to the pandemic, is that it has created a very strong incentive to automate the work people do. A machine, after all, does not fall ill. It does not have to self-isolate to protect customers or co-workers. It will not have to take time off work. For now, that incentive has been kept at bay, particularly in the UK, by the remarkable government interventions we have seen to encourage employers to keep their employees on the books. At one point, up to 9.4 million workers, about a third of the total number of employees in the UK economy, had up to 80% of their wages paid by the state.

As these measures are inevitably wound down, I worry about the consequence of this incentive. I worry about lower-paid workers in particular, and this theme is important to emphasise. One of the tragic ironies of the pandemic is that these jobs have been hard to automate in the past because they involve interpersonal interaction or manual dexterity that needs to take place indoors. Those are the very reasons why these jobs have been so hard hit during the pandemic. It is through face-to-face interaction, gathering and carrying out manual activities together that the virus spreads.

I worry about low-skilled workers, because the incentive to automate the tasks that people do applies particularly to them now, in a way that it did not in the past since their work was so hard to automate. That is not the case for many white-collar workers, who, as we all know, have been able to work remotely, retreat to home offices and use technology to do their work differently instead. On balance, my view is that the threat of automation has increased because of the pandemic.

Lord Hollick: It is a threat and a reality, by the sound of it. Professor Hall, could you give us your take on this? Perhaps tell us how well prepared you think people in the UK are in terms of level of understanding, public trust and reskilling to take advantage of the benefit of AI.

Professor Dame Wendy Hall: I am not as expert as Daniel in this. For blue-collar workers in the factories and manufacturing, I understand what Daniel is saying. It is a big risk. Companies have to make that investment and it is a really difficult time now. In some ways, Covid has slowed that down, because they will not have the money to do that. Generally, the rate at which the forecast saw the jobs going was too great. Over the last few

years, we have seen that it takes a long time. Some companies are not even digital, let alone using AI. It is really hard sometimes for management to introduce that type of technology and it does not happen that quickly. The jury is out on how that will play out.

On skills, we are nowhere near ready. We talk about new jobs that might come in for white-collar office workers, such as audits, checking data for bias and the huge issues with making sure that datasets are complete, a lot of which cannot be done by AI at the moment and needs to be checked with the human in the loop, as Mike said. Those jobs are not there yet. We are still preparing people for jobs that do not actually exist yet.

We are in a very strange world of wanting to prepare people for something, although there are not that many AI jobs to apply for unless you are a machine learning programmer. We need to prepare a future workforce, but the jobs they might be doing are not there. I am very pleased that we introduced something that was in the Hall-Pesenti review—MSc courses on AI for non-technical people. The Government funded that last year to get more diversity into the workforce and bring in different disciplines, so it is not all about coding; it is about all the things we have been talking about. I am not as pessimistic as Daniel. Covid has both accelerated it and slowed it down.

Lord Hollick: Professor Wooldridge, are we in the slow lane or the fast lane in terms of disruption?

Professor Wooldridge: I am probably closer to Wendy's position than Daniel's. I would emphasise not so much the way people will lose jobs or jobs will be created as a consequence of AI. Both those things are true, but I would emphasise that AI will change the nature of work.

Let me give you an example. Geoff Hinton is one of the researchers behind the current boom in AI. He is often described as the father of deep learning. In 2016—I do not have the details—he said approximately the following: it is obvious that deep learning is better than radiologists analysing X-rays, so we might as well stop training radiologists now. Naturally, radiologists were rather upset about that quote. If you google "Geoff Hinton radiology", you can find some quite outspoken views on the subject. How many radiologists have lost their job to AI since 2016? Geoff, who is a very smart individual, clearly thought this was imminent. I predict none. I do not think anybody has lost their job.

AI in radiology will make better radiologists. It will be another tool that radiologists can use. It will change the nature of their role, but it will make them smarter. It will allow them to focus on the stuff that is important and requires human consideration, judgment and skills. Dealing with patients is a big part of being a radiologist. It is not just a matter of looking at an X-ray and analysing what is on that.

We will see AI embedded throughout the workplace. I do not know if it is true in the House of Lords, but I am certain it is true at the University of Southampton and in Daniel's job. A huge part of my job involves endless

web forms that I need to press buttons on to complete, to send them through to the next person in some long chain of workflow and so on.

AI will improve all that. It should make it smoother, simpler and so on. AI will become embedded in everything we do. It will not necessarily make huge numbers of people redundant, but it will make people redundant. There is another quote from Andrew Ng, an eminent AI researcher. He said that any decision a human being can make in one second, AI can make.

That is a huge number of decisions. If your job involves just looking at a form with name, date of birth, salary, address—details of people's lives—and making a simple decision on that basis, that can be automated with AI. A lot of jobs involve that kind of thing. I am not suggesting they should be, but they can easily be automated with AI. Those kinds of jobs are at risk. Over the next 10 to 20 years, I am not at all optimistic about jobs in call centres where the only part of human intelligence is understanding what somebody is saying to you, but otherwise you are largely following a script.

The pandemic has demonstrated the importance of state-of-the-art computing technology. What we are doing now would be impossible without it. It may not be obvious to you, but in Zoom, which we are using now, there are an enormous number of AI techniques under the hood to do image processing so that we can all see each other, and so on. As that illustrates, AI can do an awful lot for us. Something we see every morning and will see another 10 times today is somebody who starts speaking but forgets to unmute. Why has AI not solved that problem? It should have done by now. If somebody is looking at the camera and talking while they are muted, it is not too smart a thing to realise that they probably want to be unmuted.

There is a lot of scope there, and the pandemic has brought focus to those issues. In a sense, it has spurred a wave of innovation, and I would expect to see that roll out over the next few years.

Lord Clement-Jones: We live in hope of that.

Q5 **Lord St John of Bletso:** Good morning. It is good to see you all. I do not have any direct interests to declare. I am a director of Empati, which promotes the efficient use of renewable energy, and we have been looking at how AI strategy can address climate change. We have seen encouraging developments in apprenticeships, large research programmes and PhD training programmes.

Is the United Kingdom an attractive place to learn about and work in AI? What mechanisms can be put in place to enable more collaboration between universities and industry? How can our AI research centres be more inclusive and co-ordinated?

Professor Dame Wendy Hall: There are three questions there. On industry and academia, yes, we are. Amazingly—I expect Mike will support me on this—we have seen international students coming into the UK, even during the pandemic, to study AI at the universities. That flow does not

seem to have stopped, even from China. Our biggest problem is that we do not have enough teachers to teach the courses. We are still at the top of the league as an attractive place to study AI in normal times.

As we go through Brexit, we have a huge opportunity to remodel the way we bring in international students. Until now, we have had a huge difference in fees between UK and EU PhD students and anybody from outside the EU. We have a chance now to attract international students from all over the world, including Europe, in ways that are affordable. The UK should seize that opportunity. It is possible in ways that were not possible when we were part of the European Union. We need to do that, because we are a sought-after place.

To your third point, I am a fan of the Alan Turing Institute. How it will be funded is currently being looked at. I will not go into the details, but the Alan Turing Institute can be a co-ordinator for all our universities doing AI. It is not a university, but it can support the cohort of universities involved in AI research and education. I would definitely support the development of the Alan Turing Institute.

On industry and academia, I am very keen to train people to do PhDs who are based in industry. We used to call them engineering doctorates. We need to make a move towards that, so that they do not have to leave their job for three or four years to come and work at a university. They can do the PhD while they are working and we supervise them jointly with industry. We should develop more schemes like that. I do not know if that helps, but that is one idea.

Lord St John of Bletso: That is hugely helpful. Thank you.

Professor Wooldridge: The UK has always been and remains an extremely attractive place to study AI. We have been there from the very beginning. The moves we have seen, the introduction of the centres for doctoral training in AI and the various other initiatives that Wendy has alluded to, have been a tremendous boost to this area. That was exactly the right thing to do and it is having an effect. Interestingly, when I went back and studied my evidence in 2017, both Wendy and I called for an increase in capacity. We were desperately short of capacity and we are addressing that issue. We still have a long way to go before we get there.

Let me put a little grit into the discussion. From the perspective of somebody who spends a lot of time trying to persuade top researchers to come and work at my university, the UK has always benefited, on the wider European scene, from being seen as being a stable, liberal, cosmopolitan, meritocratic society. All those things have made us an attractive place for researchers to come and work. If you are a top AI researcher, you can take your pick. It is a given that you will get an outstanding salary anywhere in the world.

Those attributes have always been an enormous bonus to us as a nation, in attracting talent, not just in AI but across the board. We still have all those attributes, but our reputation has taken a bit of a battering for

obvious reasons over the last few years. I recall Lord Ridley giving me a very hard stare when I made those comments last time, but I stand by them. It is important that we do everything we can as a nation after Brexit to make sure that we are not putting up barriers to talent and we are not sending out "foreigners are not welcome" type messages. It is extremely important, because top researchers will go elsewhere. They are the most mobile people in the world.

Lord St John of Bletso: I totally agree with you.

Dr Susskind: Higher education, while critical, is only part of the story. One of the challenges is making sure that we are getting the next generation of researchers excited about the field and the discipline. What are we doing in primary and secondary schools to capture the excitement there is at the frontier of the field in a classroom setting? My sense at the moment is that the curricula we have developed do not capture it. They do not reflect what is really happening in the world of artificial intelligence. That is a shame.

Q6 **Baroness Rock:** My interest to declare is as a board member of the Centre for Data Ethics and Innovation.

As you will recall, we asked our witnesses for one recommendation they would like the Committee to make. One thing we recommended was that the Government have a very clear road map for success. I would like to ask our witnesses for one thing, just one, that the Government could do to improve our approach to artificial intelligence in promoting opportunity and mitigating risk.

Professor Dame Wendy Hall: Is it just one?

Baroness Rock: It is hard, I know.

Professor Dame Wendy Hall: One thing that has happened since the reports came out is the development of the Office for Artificial Intelligence and the AI Council. It is really important that the Office for Artificial Intelligence is strengthened and works across government. It is hard for it to do that. Please, government, listen to your AI Council. That is what we are there for. We have developed a road map, which is not a public document yet. We are looking to launch that. It builds on everything that has happened over the last three years and makes recommendations going forward.

Professor Wooldridge: I give a big thumbs up for what has happened nationally in AI in the UK over the last few years. We have done the right things. The one thing I would ask is to understand that this is a long-term project. It is not a fleeting thing. We have to stick with this. We are trying to make a fundamental change to training in AI, our capacity in this area and, more broadly, the technologies around that. Let us hold this course and be aware that this is not a one-year or 18-month project. We are doing the right thing; we have the right course. Let us stay with it.

Dr Susskind: Finally, on the labour market, take seriously the threat of a world where there is not enough well-paid work for people to do because

of the remarkable technological changes taking place. That is not to say that there will be some big technological bang in the next few years, after which lots of people wake up and find themselves without work. Because of the changes taking place, I worry that, as we move through the 21st century, more and more people will find themselves unable to make the economic contributions to society that they might have hoped or expected in the 20th century to make, given their education and background. It is a less dramatic and perhaps less shocking threat, but none the less one that I would encourage people to take seriously.

Lord Clement-Jones: I will unfairly slip in at the last minute with a one-word answer question. Do you think we are as internationally competitive now as we were when we looked at this as a Select Committee?

Dr Susskind: No.

Professor Wooldridge: No.

Professor Dame Wendy Hall: Yes.

Lord Clement-Jones: We will unpack that at a future date. I thought it might elicit some interesting answers. Thank you very much indeed. That has been a brilliant reprise to give us the overall position.