HOUSE OF LORDS

Communications and Digital Committee

Corrected oral evidence: Large language models

Tuesday 17 October 2023

2.30 pm

Watch the meeting

Members present: Baroness Stowell of Beeston (The Chair); Baroness Featherstone; Lord Foster of Bath; Baroness Fraser of Craigmaddie; Lord Griffiths of Burry Port; Lord Hall of Birkenhead; Baroness Harding of Winscombe; Baroness Healy of Primrose Hill; Lord Kamall; Lord Lipsey; Lord Young of Norwood Green.

Evidence Session No. 4 Heard in Public Questions 29 - 36

Witnesses

I: Professor Dame Wendy Hall, Regius Professor of Computer Science, University of Southampton; Professor Dame Muffy Calder, Vice-Principal and Head of College, University of Glasgow; Dr Jeremy Silver, CEO, Digital Catapult.

Examination of witnesses

Professor Dame Wendy Hall, Professor Dame Muffy Calder and Dr Jeremy Silver.

Q29 **The Chair:** This is the Communications and Digital Committee's inquiry into large language models. We have two witness panels this afternoon, and I shall start by asking the witnesses before us now to introduce themselves and the organisation they are speaking for—if you are representing an organisation, that is.

Professor Dame Muffy Calder: Hello. I am a computer scientist, and vice-principal and head of the College of Science and Engineering at the University of Glasgow, so my responsibility is for all the scientific and engineering disciplines.

Professor Dame Wendy Hall: Hello. I am regius professor of computer science at the University of Southampton. I am well known for the work I have done on the internet and the web, but in 2017 I co-chaired and wrote the report/review of UK AI strategy and was a member of the AI Council from that.

Dr Jeremy Silver: Good afternoon. I am the CEO of Digital Catapult, which is one of the network of Catapults, funded by InnovateUK and designed to act as an acceleration layer between academia and commercialisation. We have a long-standing interest in AI, and we have conducted a number of programmes in that area over the last five or six years.

Professor Dame Wendy Hall: And I was on his board for a long time.

Dr Jeremy Silver: And she was on my board, so I have to mind my Ps and Qs even now.

Professor Dame Wendy Hall: You are not on it anymore, so you are all right.

Q30 **The Chair:** It is obviously a small world in the AI/tech arena. For this first panel we will concentrate on policy and government strategy—whether it is good enough, whether it is still relevant, whether the policies proposed balance opportunity and risk sufficiently, and what the UK needs to do well if we are to be successful and meet the ambitions that the Prime Minister has set out for us.

I shall come first to Professor Hall. In your view, is the national AI strategy still relevant, and what progress has been made? I ask this in particular, because we are very conscious of an apparent shift by the Government, since the creation of the task force, to something that is perhaps more focused on safety than I think was originally conceived as a strategy.

Professor Dame Wendy Hall: You are absolutely right. To be honest, I could not tell you whether there is a functioning AI strategy. There is a strategy on the Web. I do not know, because the AI council has been disbanded, so I do not have an official role and do not get updated on

anything that the Government are doing other than through the press and the various other roles I have advising various departments.

The AI Council was disbanded in May or June, and the Centre for Data Ethics and Innovation board was also quietly let go in August. The task force was set up, and I think the Government's strategy seems to be focused on what the task force is doing. There is still an Office for AI; as I say, I do not get regular updates. I still have the title of national AI skills champion, I think, but no one has told me whether I am still it or whether there will be another one. I do not get automatic, regular updates on what is happening with skills. The Office for AI, I am told, is a lot bigger than it was, but it is focused largely on organising the summit at the moment. That is what I hear, but I do not have an official update on it now to give you.

The Chair: Right, okay.

Professor Dame Wendy Hall: Because there is no council anymore, so there are no meetings and we do not get updates.

The Chair: And is it your view that the original core objectives of the strategy that was published in 2021 are still relevant for the ambition that the Government set out?

Professor Dame Wendy Hall: Yes, definitely. It was a good strategy. After the review I did in 2017, we had a list of recommendations, the Office for AI was set up, the AI Council was set up, various other things were set up, and we implemented that. As part of that work, with the Office for AI, we developed the 2021 strategy. The council was advising on that and talking to other people, and that is still very relevant. My worry is that a lot of the things that were in that strategy are not the main focus of what is happening today. The focus is on AI safety around what they are now calling frontier models.

I do not know what the skills people are doing. Jeremy will be able to talk more about adoption by industry. I cannot tell you what is happening with these sorts of strategic elements. They are still incredibly relevant, and I worry that, as a nation, we are slipping in not being open about these very important aspects of AI: how we train people and give people awareness so that they can move into that sort of world; how they get training and find out about possible new career routes if their job is being threatened; what we are doing with education. Dame Muffy and I are both on a board with the DfE, but that is not directly connected to the Office for AI. In all those things, we are now slipping.

I do not know whether you are aware of the Tortoise index—I hate league tables, but you have to be aware of them. It is a general index of how AI-ready a country is. We were third to the US and China, and there is a long gap between the US, China and us. We are now fourth to Singapore, which has taken over from us in being third, and there are others snapping at our heels. My view is that by focusing on one aspect of AI, which is what the summit is doing, we are slipping on the other things.

The Chair: Okay. As we proceed through this session, we will obviously

come on to more specific questions about risks and opportunities for commercial investments and so on, but, before we do, do Dr Silver or Professor Calder want to say anything in a big-picture sense as to where the Government are at the moment in their overall approach to AI?

Dr Jeremy Silver: I am not sure that I have a view on where the Government are. If I do, it is probably a somewhat partial one. The sense is of a somewhat fragmented approach, I think it is fair to say.

The Chair: Is that how you feel within the industry?

Dr Jeremy Silver: Yes. The origins of the strategy that Dame Wendy co-authored still stand. We still want to generate a vibrant ecosystem. We still need to encourage more skills coming out of universities and more PhDs in this area. Those kinds of objectives are being supported, and there is a degree to which that is ongoing, so all is not lost on that front.

The surprise that everybody had last November when ChatGPT was launched and OpenAI suddenly became the thing that every conversation I had was dominated by has shifted the agenda and changed the nature of the focus, probably not necessarily in our favour at this point. Given the fact that there are many other players globally who were very active in this space, the opportunity for the UK is not necessarily to try to become some sort of arbiter of the direction of travel.

At the moment, as Dame Wendy has said, that is very much the preoccupation of many of the officials. Indeed, officials who were focused on other areas that my organisation has been working have diverted from their work to look at how to make a success of the summit. One can hope that the summit has a long-term agenda and that it is not as short-termist as just one event, but I think that is the perception at the moment.

Professor Dame Muffy Calder: I have a role in a particular area in national security, particularly within the IPA Act. The approach to AI there is quite advanced, very sound and quite different from what is happening in the rest of government, so I have no concerns there, and I have some responsibilities for leading work in that area.

The Chair: Okay. There are some who think that the focus on national security is at the expense of other more routine-type opportunities.

Professor Dame Muffy Calder: I am talking about how AI is being approached in some of those organisations. It is impressive.

Professor Dame Wendy Hall: I would agree with that. I just want to add one thing that I missed. UKRI, which is the UK research funding agency, is funding a lot of very good research and Centres for Doctoral Training in AI, which will keep us ahead in research and training PhD students. That has been really good, and that is a continuation of the strategy.

Dr Jeremy Silver: I suppose I should add that there is a significant programme called BridgeAI, which is being led by UKRI and Innovate UK. We, as the Digital Catapult, are maturing a number of others engaged

with leading that. That is very much about attempting to take AI capability into industry, and that is ongoing.

Professor Dame Wendy Hall: It is not like everything has stopped.

Dr Jeremy Silver: It is important not to create the wrong impression. There is activity. It is a question of degree.

Professor Dame Wendy Hall: Does that make sense?

The Chair: Yes. We will move on and explore some of this in more detail.

Q31 **Lord Lipsey:** Perhaps we could start with you on this question, Dame Wendy, because you are less involved in some bits of government than you were. We do not even know quite what you are involved it. It may be easier for you to take an objective view on my question. There is obviously a balance in this area between encouraging innovation and taking risks. Do you think the White Paper sets that balance correctly?

Professor Dame Wendy Hall: Ah. Now you are talking about the White Paper from the Office for AI*,* which is a good point to raise, because this happened on the cusp of the flip. The White Paper on regulation came out in March this year, 2023, just as ChatGPT was causing such a stir and everybody, certainly the Prime Minister and people at that level, were looking to OpenAI and the other companies on how we deal with the potential threat to security and existential threats as well as seize the opportunities.

You have to balance our White Paper against the EU's AI Act, because that, if you like, is the benchmark that people are measuring against. Nobody thinks that it is perfect, but it is out there ahead of everybody else in trying to find a way to regulate AI, given that in Europe, whether you are in the EU or the UK, we have none of these companies. All the AI companies that we are talking to are based on the west coast of America. They will look to Washington for regulation, and we know that Biden and Harris are talking to them at the moment about that regulation.

In Europe, we can think about how we can lead on the soft power of governance and regulation. Our White Paper was very much about innovation and enabling people to innovate in this area without going straight to mandatory law, which is what the EU is doing. The EU is taking a risk-based approach, which is a sensible place to start, but there is a feeling, as with GDPR, that it is too restrictive on innovation.

It is about getting the balance right. Our White Paper was very much focused on trying to redress that balance. When it was released, a lot of people said very good things about it. The AI Council, which was still there then, talked a lot to the Office for AI and helped to develop it., but it was released at exactly the time when Geoffrey Hinton was leaving Google, and Elon Musk published his letter to say, "Let's stop for six months. This is all far too dangerous". That, of course, was because he was starting his own company. There was a lot of spin, a lot of marketing, a lot of hype, a lot of scaremongering.

Since that March pivot, as I call it, we have pivoted away to looking down the tunnel of safety and security risks. That is a good thing to do, but it is a tunnel that does not look at the broader picture. The White Paper said, "How are we going to develop that?" The White Paper and the EU AI Act were both written before OpenAI released ChatGPT. It has happened so quickly. It feels like the last few months have gone in a flash and everything has changed.

Lord Lipsey: Perhaps I could just raise my particular concern about the White Paper. It says that we should focus on high-risk concerns rather than hypothetical or low risks. It seems to me that the risks have two dimensions to them. One is the possibility of them happening. The other is the seriousness if they do happen. For example, the probability of a huge meteor wiping out the earth is very low, but the effect if it did would be very high. That is why many people adhere to the precautionary principle. I wonder whether the Government have not ducked that crucial question about what we are going to do about the perhaps remote possibility that AI will massacre the whole human race once it is let loose.

Professor Dame Wendy Hall: It is not going to do that, for starters.

Lord Lipsey: That is reassuring. Thank you.

Professor Dame Wendy Hall: I think you can rest assured that it will not do that in our lifetimes., but there are hypothetical risks. The key thing, I believe, is that the Government have gone down that route. They have gone, "This stuff is dangerous. We need to manage the risks of generative AI". I am firmly of the belief that we need to manage it, but it is not as important to me as, say, regulating face recognition, which is a very mature technology that we have still not regulated, or managing the threats to democracy of the deepfakes of this world, which is nothing to do with generative AI. Well, it can be, but deepfakes have been around for a long time.

These are the sorts of risks that are in the here and now, so we have to get the balance right in looking at hypothetical existential risks, possible though they are. I firmly believe that we could be leading in the debate about global governance in this space. We as the UK cannot do that, but we could be leading on it. I firmly believe that we need those sorts of debate at a global level. The UN is pulling its advisory board together to start doing that. We as a country also need to focus on the things that could do us more immediate harm.

Professor Dame Muffy Calder: I think it is helpful to be clear about whether we are talking about new risks or existing risks that might be amplified, or indeed reduced. It is crucial to separate the different classes and to ask whether we already have a risk in that area and whether it has been amplified.

The Chair: Can you give us an example?

Professor Dame Muffy Calder: Yes, facial recognition. Here is another one: ChatGPT and students plagiarising essays. We always had that risk, but students had to pay for the service, so not a lot of them did, and we

had ways to find out. Now, it is free and it is rampant, so that risk has been hugely amplified. Then we have a new risk: the students who are not cheating are now worried that they will be accused of cheating because we are having difficulty in distinguishing between the two. That is not a risk that I had thought of before.

Dr Jeremy Silver: Perhaps the most important distinction here, since the focus of this conversation is on LLMs, is the fact that AI is a much broader field, and LLMs are of one of the most current manifestations of the preoccupation that we have with it. What Dame Calder and Professor Hall are saying—I would agree—is that there was a whole set of attendant risks associated with AI before the advent of LLMs. The fact that we are now in a position where we have real-time capability that is far more accessible to far more people changes the nature of the risk.

That is the thing that has triggered the level of concern that we have seen and has escalated the nature of the conversation. It is not necessarily that even the nature of the risks themselves is that much greater than before, except for the fact that it is now speaking in natural language and is capable of producing code. The level of accessibility and the speed with which it can deliver all that to a far larger number of people also creates a greater risk.

The Chair: If I understand the three of you correctly from your contributions so far, the White Paper was developed before the arrival of large language models, so its proposals for regulation were fit for AI prior to the more popular awareness of LLMs, and now that we have these new LLMs, the regulatory model proposed is perhaps not sufficient for the risks.

Professor Dame Wendy Hall: I would not agree.

The Chair: Tell me what you think.

Professor Dame Wendy Hall: I do not think it took account of what we have had since the release of ChatGPT and, as Jeremy said, the availability and accessibility of that. I do not think it makes the principles behind the White Paper incorrect.

The Chair: Okay, good. Do you agree with that as well, Professor Calder?

Professor Dame Muffy Calder: Again, I think it is all about sector and purpose. It is not just about the technique; it is to what end you are using the technique. Even before LLMs, we had the big problem of authentication of information that is available on the web, from video to text.

The Chair: So the White Paper is still relevant to large language models.

Professor Dame Wendy Hall: The principles are.

Q32 **Baroness Harding of Winscombe:** I apologise for being late. You said specifically that the principles in the White Paper still stand. Do you think that the light-touch regulatory approach of the White Paper still stands?

Professor Dame Wendy Hall: That is hard to answer. I think that it has sharpened the need for more regulation rather than less. I still like the approach of the White Paper, which enables innovation, but—the others might disagree with me—it was basically saying, "We'll use existing agencies to regulate". We need to think about whether we need a specialist agency to pull it all together. Who is co-ordinating? Who is conducting that orchestra? I am quite open on that, though.

Dr Jeremy Silver: I do think that there are two levels to this. In a sense, we are talking about regulating something where our ability to regulate it is very limited for all the reasons that have been set out, which are to do with the location of the businesses and the nature of the technology. We ought to be thinking about our contribution to some kind of global regulation. The lack of any global institution being capable of carrying out that task at the moment is notable, so there is perhaps a role for us in trying to get to that point.

That is one strand. The other strand is: what does this mean for the UK economy and the UK population? The challenge that I feel the White Paper does not play into enough concerns the question of how we will make sure that our industry, including our domestic industry, adopts these technologies effectively and uses them to become more competitive.

To that extent, the involvement of regulators is helpful inasmuch as they have more domain knowledge and are more sector-specific. However, even there, there is a trade-off, because the idea that each of those regulators would act individually and not work together when there are so many cross-cutting themes is something that needs to be addressed.

There are those two levels. The balance that we need to strike here is one where the UK's capability in innovation in this space is preserved—in fact, supported and encouraged—so that we do not end up smothering our own potential for leadership here, because there are more and more generations of this all to come; the job is not done.

Professor Dame Wendy Hall: LLMs are not the end of AI. There is a lot more coming down the pipeline.

Dr Jeremy Silver: Right. So it is about balancing the understandable concerns that we have about protecting consumers and protecting workers' rights in factories in the face not just of automation, but of much closer scrutiny. All those agendas are really important in all this, but the White Paper probably does not quite get the balance as right as I would prefer to see it in thinking about how we can encourage more businesses and encourage industrial adoption.

The Chair: We will come on to that. Before we do, I am sorry to press this point again, but I want to bear in mind the task force, which is very much about generative AI, and its focus on safety. Do you think it needs to be broader in its approach or agenda?

Professor Dame Wendy Hall: That would be the wrong thing to do. The task force was set up with a particular aim. It was about whether we should develop our own sovereign capability in LLMs. We still need to

revisit that, but the task force has very much gone off on safety, risks to national security and so on. That is good, but it is very narrow.

I do not think that the task force should have this broad aim. I would say that, of course, but I would bring back an AI council with a much broader remit generally to hold the Government's feet to the fire on our national strategy across the board. We need something like that. It does not have to be exactly the same, of course; it could have different people. There is no independent advisory board taking this issue to government at the moment.

The Chair: That is really helpful and really clear.

Q33 **Lord Hall of Birkenhead:** Can you help us on the question of where research and innovation in generative AI is best placed and best funded? We have a lot of data on what the Government have provided in funding for research and commercialisation. A lot was announced in March; I am not too sure what was announced fresh and what was just reconfirming, but it looks like a lot of money is being put into research, with UKRI also funding. Is what the Government are putting into research and innovation going into the right places? Is it the right quantum, or should more be done commercially to boost research and innovation? If so, what?

Professor Dame Muffy Calder: To reinforce the point, there has been a lot of well-placed funding by UKRI. A lot of it is through collaboration between universities and industry, and rightly so. That is positive.

There is one unaddressed question at the moment. For a lot of AI research, universities are good places to do it and have suitable resource. However, when you are talking about certain types of LLM, you need massive compute power. An unaddressed question at the moment is how we will make that available to academic researchers.

We need some kind of national resource. There are lots of ways in which it can go technically. Some of the reasons why we need it include the fact that, when you are drafting a research proposal, it is difficult at the moment to anticipate what your needs will be. By the time you get into it, things have changed. We need to consider what kind of national resource we have—one that is accessible to both university researchers and industry.

There are always difficult questions about things technically, whether it is single-science or multiple-science, but the real question here concerns access. How do you define fair access? Everything about LLMs concerns access to data and access to compute power. In both of those, you have to ask: where is the concentration of power, and where is the notion of fairness? We need to have a conversation. I would go to the learned societies and UKRI to start to answer those questions.

Lord Hall of Birkenhead: You mentioned compute power and access to it. Should that be publicly funded? The debate is out there. The E*conomist* argued recently that the price tag is far too high for taxpayers and that it should be provided by other commercial investors—ones with

deep pockets, I imagine.

Dr Jeremy Silver: First, quite a lot of money has recently gone in and been announced. Bristol and Edinburgh have both received around £900 million for additional compute power in this space, so the Government and UKRI are not unresponsive to that requirement. We have been here before, though, in the world of what was previously called high-powered compute. It was always expected that that might help SMEs outside of academia in some way or other, but it was never possible for SMEs to access that kind of capability. So there is a question here: where is this kind of capacity placed, and how is it made accessible? That is one thing.

Lord Hall of Birkenhead: What is the answer to that, in your view?

Dr Jeremy Silver: There are two aspects to this. One is that the level of investment that we have seen Microsoft and Google making is extremely difficult to see the public sector ever engaging in, because it is just too significant. There is one view that says, "Work with those proprietary organisations".

Lord Hall of Birkenhead: And give up doing something else to work with them.

Dr Jeremy Silver: And do not even try. Those are proprietary solutions. The other view is that there are open solutions, that the means of training those could be done for less and that, once they are trained, those models could be made available to people in different sectors and different companies in order for them to enhance themselves and customise, if you like.

The awareness of those two different routes has probably not been explored that much, because it is only just emerging now. However, it does offer us an opportunity to think about how we can move UK entrepreneurship as well as academic research forward. The danger we have is that it is easier for us to solve the academic problem, in a way. The commercial problem is more challenging, so we tend to put the money into the research side. The danger we have is that we end up producing really interesting spin-out businesses or businesses that come out of graduates from universities but do not last long in the UK because they end up being acquired and going to the US. That does not help us.

Lord Hall of Birkenhead: Those routes are interesting. If you are the Government and public investment in what is necessary is just too big, would our alternative be to say, "We have a chunk of money here, but we're also going to pull in money from other commercial sources, because we want to have something that is equivalent in scale to what is on the west coast"? Would that be a route?

Dr Jeremy Silver: We are in an unfortunate economic moment to have that kind of conversation with industry. That is a real problem right now. Most of the programmes that come out of Innovate UK, for example, require matched funding from industry. Our experience at the moment is that, understandably, industry is not thinking about innovation as its priority when it is budgeting and making cuts. That is one of the places

that suffers. So the timing is not great, but the direction of travel is something that we should actively explore.

Professor Dame Wendy Hall: Picking up on both of those points, as well as the compute power, you mentioned data. That is hugely important here. We have some of the best universities in the world here in the UK. We have a fabulous legacy in AI research at a number of universities across the UK. However, apart from compute power, you need access to vast amounts of data to train the models.

At the moment, it is all in the big companies. They take all our data and use it to train the models that we will use. We have to redress that balance somehow; the Government have a role to play here. I will not go into detail—Muffy, please pick this up—but we are a world leader on open data and the accessibility of data. Some of this data is clearly deeply secret, such as stuff for defence, security and health, but we need to learn how we can use that data. Take the National Health Service: why are we going to give that to an American company to play with? We need to learn to manage that data securely so that privacy is protected and we can use it to generate good stuff.

Professor Dame Muffy Calder: Do you mind if I pick up on this subject? I am on the Prime Minister's Council for Science and Technology. Earlier this year, we looked at the issue of data valuation and public sector data in the context of LLMs.

We have fabulous resources in health data, ONS data, geospatial data, environmental data through museums, et cetera. We have a really precious resource. Our point to the Government was that we need to be able to value this. At the moment, it is not clear how to value it and how to make decisions across government so that we do not have fragmented views about what relationships are built and what are not. If there is one message that I want to get across today, it is exactly that: this is about dealing with our public sector data.

Two questions come from that. First, who has access to that data? Secondly, who has access to the models that were built on that data? It is not just about a sovereign model; it is about sovereign data. I am not talking from a security point of view. I am talking about this being a huge national resource.

Professor Dame Wendy Hall: It is a huge opportunity for the UK. I cannot say enough how important it is for us to get this right and lead on it.

Professor Dame Muffy Calder: Let me just say that the letter is not public at the moment, I believe, but it will be published by the end of the year and has gone into government.

Lord Hall of Birkenhead: I have a final point to make, which Dr Silver mentioned. There is this data resource. It is very precious and something that the UK could really build on if it did all the things that you have just talked about. How do we ensure that, as Dr Silver suggested, we do not have a number of start-up companies that are very successful and then, hey presto, they are bought up and disappear off to the States? Does the

quote, "We are open to business", matter?

Professor Dame Wendy Hall: Yes, it does.

Dr Jeremy Silver: Hugely so. This is a problem not just in this area but broadly across the UK economy, particularly in the tech sector. It is a complicated problem. It happens because of a combination of the greater acquisitive culture coming out of the US than exists here and, frankly, a desire from UK investment to make a quick return.

There are two things that we have to challenge in those areas. One is whether there is a way in which we can encourage investors to stay longer and not sell at the first opportunity. The notion of patient capital is designed to address that, but it needs more than that.

We also have to try to create a greater culture in UK industry of the kind of investment that we see the big tech companies making. However, that requires a real enhancement of its own expertise. UK companies—UK industry—do not really address this and use it as a market opportunity for themselves, because they are either a multinational or a subsidiary of a multinational, and the decision as to where to make those investments is made not here but globally. So we are in a global competition, from that point of view. On the other hand, there is a kind of cultural reticence, frankly, and a sense that innovation through acquisition is not part of UK industry's strategy.

Professor Dame Wendy Hall: It is my understanding also, although I am not an expert in this, that all roads lead to the Treasury on this one, because we have to change the tax regime. That is my understanding of how the venture capitalists are taxed on investments. I am not an expert, so you will have to ask someone from the Treasury about it, but there are ways that mean that our investors go to the States, because they get a quicker return.

Lord Hall of Birkenhead: If I understood what Professor Calder was saying earlier about us valuing the data that we have, could that be helpful in the sense that it is UK data, or is that of no consequence when it comes to stopping it being used?

Dr Jeremy Silver: Companies go to the US because it is the best country in the world to scale a business; there is no question about that. But that does not necessarily mean that the ownership has to go as well. That is the distinction that we should try to make.

Professor Dame Muffy Calder: There can also be a role for government procurement processes and looking in the round when awarding contracts, et cetera.

Lord Hall of Birkenhead: That was really interesting. Thank you.

Q34 **Lord Kamall:** I just want to pick up on the issue of data and access to data, not only public-sector data but as much data as possible if you are really going to train your systems. One concern that we have heard previously has been the whole issue of barriers to access to data, particularly things like copyright. We understand that copyright is very

important, for various reasons, but is there a way to access that data while respecting copyright, or is this a very difficult barrier to overcome?

Dr Jeremy Silver: One of the challenges of LLMs, as we know, is that they are all about language. They have also trained on music, as we have seen. The approach has been to act first and ask for forgiveness after; I think that is the phrase. It is interesting that the European draft legislation encourages those who are training LLMs to identify the sources upon which they are training. It does not encourage them necessarily to clear copyright. It points rights owners to where their fight is, by having that information highlighted, but it does not actually come to their aid.

It is interesting that the very nature of the way that LLMs are trained means that, technically at least—this has not been argued in court yet, although there are a number of cases going through at the moment—it is not making a copy. It is not even making what you would define as a derivative work. Therefore, the current form of the legislation does not really provide the protection.

Most of the rights-holding companies are trying to argue, understandably, that the current copyright law does provide that protection, because certainly the spirit of the law would seem to. We have seen huge debates about this, particularly in the music industry, as you know. US copyright law is slightly different from that in the UK. In the US, there is a right to protect one's own likeness, which does not exist under UK copyright law. My belief is that we will need to change that legislation, that it is not adequate, and that is one of the few areas that we can see really clearly needs to be acted on.

The point about this is that there is also a huge amount of other data that is not copyrighted; it is not owned by rights holders. It has the challenge that it is very often personal identifier data and therefore has a whole different set of concerns around it.

Lord Kamall: Do the other witnesses want to add any more on copyright or IPR? No? That is fine.

The Chair: Coming back to what you were saying about US acquisitions of start-ups and that sort of thing, I was on a panel recently, and a fellow panellist argued that we fixate too much on exits and that a US firm buying up a UK start-up does not necessarily mean that we lose the value of that start-up here in the UK. You have touched on this already, but I would be really interested to know your view. It was one of those arguments where I thought, "I'm not sure I've heard that before".

Dr Jeremy Silver: It is a great question.

Professor Dame Wendy Hall: DeepMind is the obvious example.

The Chair: You would argue that it is still here in the UK.

Dr Jeremy Silver: If we were in some way able to control that, so all those talented people who were in an organisation or a company that was acquired remained in the UK every time and continued to be employed here, to pay tax here and to spread their expertise around the

UK ecosystem, you might say, "Okay, that's less of a concern". If you look at what actually happens, that is unusual.

DeepMind is still here and did not move largely because of the numbers of people in the company when it was acquired. It simply becomes impractical to move 300 or 400 people to the US. However, if the company, as in most cases of early-stage companies when they exit, has turnover of between £10 million and £20 million—that is the average sweet spot for an exit—at that point, the company probably has maybe 100 staff. In tech companies, it can be much fewer than that. It might be as low as 20 or 30 staff, and that makes them much more mobile and much more transferable. So that is part of it.

The other piece of this is about what scale they reach before they are acquired, which is something we have seen repeatedly here. Just as an example, we ran a programme very successfully for three or four years called the Machine Intelligence Garage, which was helping early-stage AI companies to access some compute, mentorship, strategic guidance, and ethics and responsibility work. One company went through that. It was a group of 12 graduates from UCL who had no product and no revenues and went through our programme. Within a month of finishing it they were completely acquired by one of the large tech companies and were all shipped off to California, where they now reside.

Q35 **Baroness Fraser of Craigmaddie:** I need to declare an interest, because Dr Silver and I sit on the board of the British Library together.

I want to come back to league tables. You mentioned that the UK has slipped and Singapore is now number three. What, in your opinion, does the UK need to do to capitalise on its strengths? What is the right ecosystem for the UK?

Professor Dame Wendy Hall: I think we were on the right track and we have slipped off that track. I hope this inquiry will help us to get back on track.

Baroness Fraser of Craigmaddie: Do you mean by that the balance between safety and development?

Professor Dame Wendy Hall: And all the other things, like skills, adoption, awareness, entrepreneurship. There is no focus on that at the moment. That is what I am saying.

Professor Dame Muffy Calder: I discussed access earlier. Many companies do not know where to turn on AI. Do they have a question that AI could help them with? It is not clear where they should go at the moment. They do not have a research project in mind. They need some simple advice and guidance. It is very hard to reach those companies, but it could be done better, I propose.

Professor Dame Wendy Hall: Before DSIT was created, the Office for AI was across BEIS and DCMS. I am not saying that that was perfect, because it is quite difficult to be an office for AI when you are trying to advise everybody in government, industry and funding agencies and you are straddling two departments. But we are now in a position where the

Office for AI has gone into DSIT, so it is very much part of that innovation, science and technology ecosystem.

Every department in government should have an AI strategy for its sector, just as every department has a data strategy, remembering that data and AI are inextricably linked, so they should be joined at the hip. Every department in government should have an AI strategy, someone responsible for AI in that department, and there should be an ecosystem that brings all that together in some way and links it to the Office for AI and to whoever in government is leading on AI, whether that is No. 10 or however it is managed. That is an ecosystem that we should move towards.

Baroness Fraser of Craigmaddie: Can you say a bit more about what you think about developing a sovereign capability.

Professor Dame Wendy Hall: That is a whole different question.

Baroness Fraser of Craigmaddie: To me, it also follows on from the conversation we have just had about data.

Professor Dame Wendy Hall: Yes, it does, definitely.

Baroness Fraser of Craigmaddie: You have spoken about the incredible value of the data sources we have, but there is a danger: the stable door opens, the horse bolts, and we have not done anything about it.

Professor Dame Wendy Hall: Your special adviser designed a plan—so ask him—for how we develop our sovereign capability in LLMs. I am definitely in the camp that says that we should really consider this. That is why we all supported the setting up of the task force and the money to do that. I hope it will go back to thinking about that as well, because, as I said, we missed cloud; we have no sovereign cloud capability. So if any of our security agencies or the health service want to use cloud, they have to buy American—they are not going to buy Chinese—cloud capability.

We will be in exactly that position with generative AI if we do not do something about that now. Is that a problem? That is what I was hoping the task force was going to address.

Dr Jeremy Silver: It is interesting. The Norwegians are in the process of developing their own LLM, because so much of their documentation is in Norwegian, which, funnily enough, is not captured by the American organisations. So they can see a real value in doing that. The extent of the investment is immense. It took the entire academic compute capability of the nation for three months to make the first version of their algorithms for this. Norway is a relatively small country compared with the UK, so the challenge for us in doing that is immense.

Baroness Fraser of Craigmaddie: And there is the question of whether we think it is worth it.

Dr Jeremy Silver: Because of the English language, of course. In some sense, that is our advantage.

Professor Dame Wendy Hall: There are other reasons that I can think of why we should.

Dr Jeremy Silver: Absolutely, yes.

Professor Dame Muffy Calder: It is also not one LLM. It is LLMs, which align with the different datasets. I keep separating the models from the datasets. There are sovereign aspects to both of them, but they are separate issues.

Q36 **The Chair:** I want to come back to when the AI Council still existed. Did you as a body at that point raise with the Government the arrival of LLMs and what the implications of them were going to be?

Professor Dame Wendy Hall: Yes, we did. As I know, the Turing institute did as well. The AI Council wrote to the Minister at the time. It was not a public letter, so I cannot tell you what we said.

The Chair: Can you tell us which Minister it was?

Professor Dame Wendy Hall: I think it was George Freeman at the time. So we had started that discussion. This was all, in a way, swept away because of ChatGPT. As for the criticism that nobody in the academic community in the UK was looking at LLMs, we have a number of experts in the UK researching this area, and people, like the gentleman to your right, were raising these issues with government and the funding agencies. So it is not the case that we did not see it coming. Google did not see OpenAI coming with GPT. If you want to talk about existential threats, OpenAI was an existential threat to Google. They are coming back, and there will be interesting territorial wars between these companies.

We have, as I say, a fabulous research legacy, and people around the table of the council knew that this stuff was coming and were starting to talk to government about it. We had plans about sovereign capability before ChatGPT emerged.

I want to say one thing just to make sure that the record is correct. The council has gone, but, as I understand—we got letters to say this—DSIT was creating a pool of AI experts to consult on AI, in the same way that SAGE worked on Covid. But I have not heard anything since the summer. We were all going to be in the pool of experts, but we have not had any follow-through on who is there and how we will be consulted.

So I cannot say that there are no plans. They may be able to make that work—it is a reasonable model—but consulting a pool of experts is very different from having a set of people you call on and get to know over a period of three years, or however long the term is, and then refreshing that model. That is a much healthier model. A pool of experts is supposed to be being created.

Lord Kamall: I just want to be absolutely clear on this. When the AI Council was disbanded, were there individuals from it who are now on the Frontier AI Taskforce? Has it taken that expertise from those individuals?

Professor Dame Wendy Hall: I am not 100% sure who is on the task force.

Lord Kamall: Right. Maybe we should look at that.

Professor Dame Wendy Hall: As to whether there was any overlap, I would have to check who was on the task force. I do not know of anybody. There might be one or two.

The Chair: Okay. And do you have any view on the role of the Turing institute?

Professor Dame Wendy Hall: We led the world in establishing a national institute for AI, and I think we should have a national institute for AI. Turing is going through a review at the moment, and I hope what comes out of it is that we have a national institute for AI that is fit for purpose, which is a national institute for AI co-ordinating some of these things and being the place where people can go to say, "Where do I find this out?" I firmly believe that research happens in universities.

Professor Dame Muffy Calder: I always had the vision of the Turing institute as a resource for government, and it is not clear that it has been able to fulfil that role.

Professor Dame Wendy Hall: I never thought of it like that. I thought of the Turing institute as being a resource for the UK generally.

Professor Dame Muffy Calder: Which includes the Government.

Professor Dame Wendy Hall: Well yes, of course.

Dr Jeremy Silver: I have one further thought on that. I very much hope that the new version of Turing succeeds and flourishes, but I do think that individual industrial sectors need specific support in this area, and I am not convinced that that will come from a largely academic institution, which I think should be doing the leading-edge work that it applies across the field. We need to then be able to apply that into industrial sectors, and we need another group of bodies, some of which may already exists. The catapults potentially perform some of that function.

Professor Dame Wendy Hall: TechUK is also very strong.

Dr Jeremy Silver: Potentially. The idea, when it was founded, was that it would have industrial application, and it has become pretty obvious that that has been very challenging to achieve. I would like to see a more integrated approach between Turing and agencies that can then lead industrial adoption.

The Chair: This has been a very rich session, so I am grateful to all three of you for your testimony today and for giving up your time to come and talk to us.