



Communications and Digital Committee

Corrected oral evidence: Large language models

Tuesday 21 November 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; The Lord Bishop of Leeds; Lord Lipsey.

Evidence Session No. 10

Heard in Public

Questions 83 - 94

Witnesses

[I](#): Hayley Fletcher, Director, Competition and Markets Authority; Dr Yih-Choung Teh, Group Director of Strategy and Research, Ofcom; Stephen Almond, Executive Director, Regulatory Risk, Information Commissioner's Office; Anna Boaden Director of Policy and Human Rights, Equality and Human Rights Commission.

USE OF THE TRANSCRIPT

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

Examination of witness

Hayley Fletcher, Dr Yih-Choung Teh, Stephen Almond and Anna Boaden.

Q83 The Chair: This is the Communications and Digital Select Committee, and we are continuing our inquiry into large language models. We have two panels of witnesses today. First, we are going to hear from some regulators and then, later on, we have some representatives from the large frontier research lab businesses.

Just before we get going, I should say that the committee wrote to 10 regulators seeking written responses to a range of questions about how regulators are preparing to meet the expectations set out by the Government in their AI White Paper. If anybody wants to see those responses, they are available on our website.

Before us today we have three regulators which are members of the Digital Regulation Cooperation Forum, the DRCF, but also the Equality and Human Rights Commission, which is not a member, in order to get the perspective of a regulator as well that is touched by this technology but is not necessarily part of the normal digital-facing regulators.

I will ask each of you to introduce yourself and the organisation you are here to represent. Then we will move into questions.

Dr Yih-Choung Teh: Good afternoon, I am group director for strategy and research at Ofcom.

Stephen Almond: I am the executive director for regulatory risk at the Information Commissioner's Office, the UK's data protection regulator.

Hayley Fletcher: I am a director at the Competition and Markets Authority, and I led our recent work into AI foundation models.

Anna Boaden: I am director of policy and human rights at the Equality and Human Rights Commission.

The Chair: Thank you again, all four of you, for being here. We have quite a bit of ground to cover. I would ask each of you to try to be not just informative in your responses but crisp as well, so we make sure that we cover the ground that we want to cover. I am going to go straight to the Lord Bishop to get us going.

Q84 The Lord Bishop of Leeds: I am glad that the introduction referred to the AI White Paper in March 2023, where it was affirmed that the five principles would be implemented by existing regulators. Could each of you in turn explain to us what problems and challenges you expect to see and adjudicate on within your respective sectors over the next three years, and is this doable in the existing regulatory framework? Can we start with Anna Boaden?

Anna Boaden: It is fair to say that this is a new area of work for the EHRC. Some of the regulatory risks around equalities are probably more self-evident around discrimination, bias and the need for transparency than perhaps the more nuanced human rights risks. For us, it is critical

that we are able to work across our regulatory family to consider some of the issues but, also, that we are effectively resourced with the right technical capabilities to be able to respond in an agile way. We have a robust regulatory model; it has been around for a while, but it was very much built to help people. We are dealing with rapidly evolving technology, so it is important for us to consider what that means for a small regulator with an important remit around equalities in particular but also human rights.

The Lord Bishop of Leeds: You said a couple of times there that it is important; is it achievable?

Anna Boaden: Yes. If we are structured and resourced in the right way. There is something for us about making sure that, when the technology is developed, equalities and human rights principles are embedded so that we can then focus on outcomes. As you will be aware through the inquiry, this is a massive area, with huge human rights and equalities risks particularly, and we are small in size. There is a lot of will but not a huge amount of resource, so I think there is a challenge for us there.

Hayley Fletcher: It is incredibly difficult to reliably predict what is going to happen in the next three years, as you ask. From our perspective, we think about a spectrum of possible outcomes. In some spheres, we think that strong competition might be the outcome for AI, particularly in foundation models and large language models, that spurs innovation and growth. On the other hand, we could find ourselves in a situation where there is ineffective competition and poorer outcomes for economic growth and consumers.

The way in which we try to answer that question is to think about what the key drivers to those positive or negative outcomes might be. We have put forward some proposed principles precisely with that view. It is uncertain as to how the market will develop but we want it to be more in that positive end of the spectrum. Our proposed principles set that out with the hope to guide the market towards those outcomes.

Stephen Almond: I would start by saying that AI is not new. While generative AI, indeed large language models, have had their moment this year, considerations around how to regulate them have been around in the context of data protection law for some time. What is perhaps different in this context is more questions of the scale, the complexity of the technology and the rate of adoption rather than there being a core challenge around how it plays in with data protection law.

Where large language models are trained on, fine-tuned with, or continue to process the deployment stage of personal data, the data protection law applies. The data protection law is designed to be technology neutral, and principles based. We apply those principles to whatever new and emerging context we find ourselves in. There are a number of areas where I would say there are basic expectations that people have around how large language models process their personal data, whether that is about being transparent about the personal data that is being processed or considerations around how people's information rights are respected, which I think are still important points

for industry to get right in this space. There are also slightly more nuanced points—and we might come on to some of those—around areas where there is still some thinking to be done about how best to bring data protection into the design of these models.

Dr Yih-Choung Teh: I would observe that our focus, of course, is on the services used by people in the comms sector rather than specific technologies itself. I agree with Stephen that the technology and consumer behaviours that we see are developing at an increasing pace. Of course, that brings huge benefits but with it what we see is an acceleration and exacerbation of some of the harms or risks associated with it.

AI is used across our sectors but to pick one example, in the next three years we will be implementing the Online Safety Act. That legislation has had the benefit of Parliament considering the implications of AI and generative AI, and those services use AI throughout.

To pick a particular example, things like content recommendation engines are algorithmically driven. AI is an important part of the answer to some of that in reducing the prevalence of illegal content. However, equally, of course, generative AI has the potential to create large amounts of harmful content as well. It is very much at the heart of how we consider the regime.

The Lord Bishop of Leeds: On one hand, you are saying that the technology is developing so rapidly that it is very hard to be clear about what might happen in the future—we are not clear about the risks that may emerge in the future—and yet, on the other, we are saying that the regulatory regime that exists at the moment is adequate to absorb it or cope with it. Is there not an inherent contradiction there?

Dr Yih-Choung Teh: No, I do not think so in the sense that it will always be a challenge because of the pace of change. We invest significantly in horizon scanning to try to understand what the implications of these technology and consumer behaviour developments are around the outcomes. With some of our regimes—not just online safety but telecom security—Parliament has given us regimes that put the responsibility on the platforms, on the networks themselves, to undertake risk assessments and then put in place systems and processes that mitigate those risks. Therefore, you get to a place where you want to see, by design, a more safe or secure environment and not just when the regulator sees something unattractive after the event. That gives me some optimism that these are the right kind of regulatory mechanisms.

Stephen Almond: Far be it from me to step in and talk about the Government's White Paper, but I would say that the approach that they have set out of relying on existing regulators is grounded in a lot of common sense. The approach that we need to be looking at in respect of AI, which is a general purpose technology where the risks that emerge from it will be quite specific to the context in which it is deployed rather than just deriving from it as a whole, means that we need to think less about how to regulate it in the round and slightly more about where the emerging gaps are in regulation that need to be closed off. For example,

we see some areas of regulation where it is exposing some challenges. I know there has been a lot of debate on this committee around copyright, for example. Closing off copyright and some of the challenges there and finding the right balance between different interests does not necessarily mean that we need to reinvent the model for how to regulate AI in the round. What is important is having—as my colleague said—the right processes for identifying where there are risks that are not mitigated.

To that end, I am very interested in seeing progress from the Government in relation to the central risk function, which is supposed to identify gaps that are not yet mitigated and work out what the solutions should be. Within our own regulatory parishes, we can talk about the extent to which we are able to rise to those challenges but there will be gaps that sit outside that framework.

The Lord Bishop of Leeds: We may come back to them a bit later, but I have one final question to Ms Fletcher. In your submission from CMA, you talked about an increasing concentration of power among big firms and market exits from challenger firms. Witnessing what we have just seen in America over the last 48 hours with OpenAI, how do you regulate against this? How, beyond wishful thinking, do we ensure that the big tech firms do not concentrate power and, therefore, enforce the exit of some of the challenger businesses?

Hayley Fletcher: For reasons that I think you will appreciate, I will not be able to comment on the specifics of the case that you mentioned. The facts of that are not yet clear. The way that we do that concretely is in the principles that we have set out, explaining to both businesses and consumers in the market what we expect and what the right conditions can be for competition to flourish.

For example, we have been absolutely clear that no firm should be subject to any anticompetitive conduct. We have to have markets where there is fair dealing. That is crucial for businesses to have the confidence to invest. We have also set out that it is important that the widest possible range of firms have access to critical inputs, such as data or compute or the relevant expertise.

That mechanism is going to ensure that a range of firms can bring the best models to market and guard against the risks in that version—as you rightly point out—of the world that I set out earlier on, so that we are not in a version where there is ineffective competition and this technology is in the hands of the few. We think we can do that concretely through these proposed principles.

The Chair: Thank you. Lord Hall has a supplementary before we move on.

Q85 **Lord Hall of Birkenhead:** A quick one to Dr Teh. You mentioned the rapidly changing consumer behaviours and demands caused by what we are discussing this afternoon. Could you amplify where Ofcom sees some of the tension points, some of the behaviours that it thinks are leading to new harms that it is having to deal with, and what do you think those are?

Dr Yih-Choung Teh: Yes, of course. My comment on consumer behaviours is that it is not just the rapid development in technologies. We invest quite a lot in consumer research to understand how people are responding to that. A year ago, we saw ChatGPT get to 100 million users within two and a half months, rather quicker than for TikTok or Instagram before it.

When we think about how effective our existing regulatory regimes are, my mind turns to public service broadcasting and the importance of availability and prominence of public service media. In our last public service broadcasting review, we made recommendations to the Government on how we can see consumers—and we all have our stories from our own experience—increasingly watching content on demand. If we want availability and prominence of duly impartial or duly accurate news, the regime has to recognise that people are watching online. We have the Media Bill introduced into Parliament to address some of those.

Lord Hall of Birkenhead: Dr Teh, I get all that. I was rather driving in a slightly different direction, which is: what are the new harms or behaviours that you are worried about as a regulator, where you think, “That is behaviour that we need to either catch up with or think about whether we have the powers to deal with it”?

Dr Yih-Choung Teh: Our concerns typically are those that are set out in our statutory duties and, therefore, the outcomes look similar. The thing that strikes me is that the application of those principles or the outcomes we are trying to achieve can look quite different as the technology and the behaviours change in the way that I described.

To pick another example, when it comes to news, increasingly we are seeing that younger generations typically are looking to social media for their news: the likes of Facebook and TikTok. Curiously, the trust levels there are of the order of 30% and not above 80% as we would see with the PSBs or Sky News.

Of course, those online intermediaries are recommending content through AI in an algorithmic way that raises some interesting questions about: do we think that there is the same diversity and media plurality that are the things in our legislation that we care about? I am not sure that the outcomes we are seeking to uphold are necessarily different but the behaviours that we see are moving sufficiently that that calls into question whether our frameworks are effective.

Q86 **Baroness Harding of Winscombe:** One thing which makes generative AI large language models particularly tricky to regulate, and which we have been hearing about in this inquiry, is their complex end-to-end supply chain. Where in that supply or value chain is the best place to achieve impact from regulatory action? Is it around the upstream model developers or the downstream users or application developers? How does that value chain map to your existing powers and remit? Let us start with Ms Boaden.

Anna Boaden: I would say both. There is obviously a greater challenge for us from a jurisdictional perspective if we are looking at the upstream

development. Part of the challenge for us, but also part of the opportunity, is working with other regulators to address some of that. There is also a challenge for us in terms of people's interactions with AI from a public sector service delivery perspective. There is a big information angle to this in ensuring that when using AI technology, particularly in the delivery of public services, people are aware—the citizen receiving the service and the person delivering it—of some of the inherent risks around bias and discrimination that exist within it.

Baroness Harding of Winscombe: Just to be clear, do you have the remit and powers to go end to end through the value chain?

Anna Boaden: We do. Depending on where the value chain starts, we would have to explore the sort of the jurisdictional risks around that. There are challenges for us in our strategic litigation approach around our ability to act in an agile way and take targeted litigation as and when we see fit. It is challenging from a cost perspective as well as just having the technical legal expertise within our organisation. That is probably where we are on the start of the regulatory journey in relation to AI.

Hayley Fletcher: We certainly think about the value chain as a whole, precisely because it is so interconnected. The kinds of models or the way that models develop directly impacts on the choices that businesses face and the models that they can choose and deploy, as well as the impact on consumers. We are certainly thinking about the entire value chain for the reasons that you suggest. We believe that our competition and consumer protection powers allow us to do that. Indeed, that is what we have done in the CMA's recent report on foundation models.

Stephen Almond: As I have alluded to before, the data protection law applies to both model developers and deployers inasmuch as they are processing personal data. We are, indeed, able to act across the entirety of the value chain. That creates choices for us about where to have the greatest impact. By paying attention to model developers, for example, we may be able to tackle risks upstream around bias that would flow downstream into the deployment of those models into various contexts. Equally, lots of the challenges that surround how large language models are used come from the context in which they are deployed.

As a regulator, we have to make careful prioritisation choices about where we can have the greatest impact and where we will be tackling the greatest harm, based on our intelligence and our powers, but we do have the ability to work across that value chain.

Dr Yih-Choung Teh: As a sectoral regulator, I would observe that our focus is on those services that people engage with and use cases. There will be something about that that looks more like downstream. Having said that, I think that the focus on outcomes means that we can take into account various inputs that are considered.

To pick on an example such as online safety, the risk assessments leading to changes in systems and processes will then, we hope, drive a consideration of how algorithms need to be amended to reduce the likelihood of them servicing harmful content. They already use AI, of

course. Increasingly, as they use frontier models, my expectation would be that, by design, that will start to have an influence on those models as they are developed in terms of what their outputs look like.

Baroness Harding of Winscombe: Just to be clear, is it your expectation that a developer of a large language model that was ultimately used in an application that was accessible to children would have an obligation under the Online Safety Act to conduct a risk assessment upstream?

Dr Yih-Choung Teh: Potentially. It depends on how that model is used. For example, if a large language model is used to generate content in a stand-alone way but that content is then taken and uploaded to a user service, that content will be caught under the regime. In a different example, if a frontier model is embedded or integrated—and we see some signs of that within the services as a whole—yes, there is a potential for that to be subject to the same requirements for risk assessments and then mitigations.

Baroness Harding of Winscombe: Are you confident that there is sufficient clarity about liability? Just take that specific case—I think it helps to make it a bit more concrete rather than theoretical across all regulation—is it clear where liability might lie down that value chain?

Dr Yih-Choung Teh: My expectation is that the primary focus is more on the downstream in the services that are offered, whether it be user-to-user services, search services or pornographic services, and then indirectly that has implications for what they are drawing on. I am sure there will be boundary cases and, of course, as we work through live cases around supervision then we will learn and understand that space better.

Baroness Harding of Winscombe: Just to make sure I understand: I think what I am hearing you say is that if I am an app developer embedding a large language model in my app that is used by children in the UK, as the app developer, I will be liable if the large language model is causing harm to children.

Dr Yih-Choung Teh: My understanding is that within the Act, whether it is Part 3, on user-to-user services, or Part 5, on pornographic services, that is the locus for the regulatory focus.

Baroness Harding of Winscombe: Does the developer of the large language model have any liability?

Dr Yih-Choung Teh: It will depend on the service that is engaging with the consumer and how that relationship works.

Baroness Harding of Winscombe: It looks to me like a complicated world if you are a developer of a large language model, but maybe I am wrong.

Dr Yih-Choung Teh: I appreciate that. You will be aware that we set out our first set of consultation documents just last week. It is still early days in getting feedback on how the regime will work, but it is my hope

that through putting out that guidance we will give increasing clarity to the industry on how this will work.

Baroness Harding of Winscombe: I will spread that out to the other three regulators. How are you thinking about liability in that whole value chain?

Stephen Almond: I will start with the ICO's position and reference how that interplays with some of the considerations around Ofcom. In data protection law—as other attendees at this inquiry have said—the allocation of liability comes down to the organisation that is determining the means and the purposes of the processing. Obviously, different types of processing activities might go on as part of the AI value chain, so the considerations of who is determined the data controller—in data protection speak—would be based on that. It would be done on a processing operation-by-processing operation basis.

For example, you may have a model developer who is collecting data that is used to train a model and they are the controller for the processing that is undertaken to train that model, but then they sell on that model to a third party. Then the third party has some responsibilities for the data that is inputted to that model, which at that point may just be a series of weights and numbers, and then they have responsibilities themselves for how the data is processed.

I am not going to claim that this is straightforward. It is a process that people can go through to work that out. It is also an area where we are looking to provide greater regulatory clarity. We have said in response to Sir Patrick Vallance's review of innovation in AI, and the regulation of it, that we will try to clarify our rules on this and our interpretation of the law to ensure that it is crystal clear for organisations.

That is how it translates in practice in terms of how we would allocate responsibility and where we would focus our attention. That does mean, as I was saying earlier, that we can focus some of our attention on developers but also on organisations deploying AI, generative AI or, more specifically, large language models.

To the points that were under discussion just prior to this, I would also highlight that some of the allocation of responsibility comes down to who is doing the different activities. While it is quite easy to separate out the process of a model developer versus somebody who is using it, there is a much longer series of activities here that may involve, for example, an app developer working with a model developer and asking for their model to be fine-tuned and made into a bespoke version for their particular context. These questions of how liability is allocated really come down to some of the finer points of who is taking on what responsibility and what role in that process.

Baroness Harding of Winscombe: That is very helpful. Thank you. Ms Fletcher, is there anything to add?

Hayley Fletcher: It is an incredibly complex question. In the way that Mr Almond said, it depends on the situation, precisely because there are there are circumstances where a firm might use a large language model

or foundation model off the shelf or it can use its own data, its own expertise, personalise it for its own customers and do more with the model in a fine-tuning process. It depends on the individual circumstances, but it is incredibly important that people are clear that everybody has a role to play in making sure it is clear who is accountable when something goes wrong. That is an issue that we are seeking views. It is a complex and open question.

Anna Boaden: I would echo my colleagues. From the Equality and Human Rights Commission's perspective, our focus in regulatory terms is on the sector in which the technology is being used. We would look at that, be that the NHS, police, local authorities, if that is a service provider who has responsibilities under the Act. There is a challenge in that the burden is on the end user to prove discrimination, which is particularly challenging given all the links in the value chain and the complexities you are setting out. That is a concern for us because it puts a high onus on an individual to understand complex algorithms and information that are not immediately apparent, even to the people developing the technology. There are risks in being able to exercise rights or challenge discrimination. This is concerning for us from a regulatory perspective.

Q87 Baroness Harding of Winscombe: It seems that bias and hallucination is a fact of current large language models at their training level, right at the beginning of the value chain. How concerned are you about your ability to reach all the way down to that level and hold the developers of large language models accountable for that bias?

Anna Boaden: It is untested for us at the moment, so it is difficult for us to answer. We can see some risks and challenges, but we do not think our regulatory powers are insufficient to be able to address those. However, there are challenges. Because we are new to the work, that is an untested area for us.

The Chair: Thank you. We will come on in a moment to the way in which you each as regulators will relate to this central function and some of the functions that it may provide that you might need to work with or rely on. Before we move on, can I come back to Ms Fletcher on a couple of things you said earlier on in response to questions on the CMA's current abilities in regulating this space? It is to clarify something that seems, from what you said, inconsistent with what was in the letter that we received from the CMA—it may be that you want to look at this when you get back to the office and write to follow up. In the letter that we received, the CMA suggested that you would have quite limited ability to interrogate AI when investigating risks, particularly in tests around algorithms. Is that true? When you initially responded, it sounded like the CMA was on top of it all and you thought it was all going to be okay.

Hayley Fletcher: I was trying to convey that competition law applies across the value chain and, similarly, all firms have to make sure that they comply with consumer protection law. As we set out in the letter, when you are considering technologies such as large language models, they are quite complex to understand, and it is important that we also

have the right information. There are some limitations to that, but to date, we think that we have the tools that we need to understand them. I would be happy to clarify that, but that is what I was trying to convey. I did not want to create the impression that competition or consumer law did not apply across the value chain; they do.

The Chair: Okay. That is great. Thank you. If there is anything later on that you want to follow up with in writing, please feel free. I know that you were dropped in to come and give evidence to us today at rather the last minute, and I am very grateful that you are here. We will move on to Baroness Healy.

Q88 **Baroness Healy of Primrose Hill:** You all agree that the pace and challenge of AI will be difficult, even though you all strike me as being quite confident as well that you can cope with it. How do you respond to concerns that regulators do not have sufficient professional and technical capacity to address the potential volume of issues that will arise from generative AI? The Government are putting so much pressure on you to perform. You seem to be the main people that will have to deal with all this. There is no other body that I can see that will have an effect on it. How would you address that, Dr Teh?

Dr Yih-Choung Teh: I can speak only for Ofcom, of course. Notwithstanding that the magnitude of these challenges is significant, we feel that we are in a pretty good place. I say that partly because, as we have had time to prepare for online safety, we have benefited from being able to expand our capability when it comes to data science and engineering, and I genuinely think we have been able to put together a world class team, pulling in talent from tech platforms as well as academia and other places.

Incidentally, I would observe on that point that being an independent regulator has been an important facet of that, where our ability to set salaries—albeit not at a level that competes with the private sector—has enabled us to recruit scarce talent and when coupled with a purpose-led mission, it has been quite an attractive proposition. I have been encouraged by the talent that we have brought in. Even on an international stage, we have assembled a very significant set of individuals to take on this challenge.

Right now, I am optimistic about our capabilities and our strength of expertise in areas such as natural language processing or computer vision when it comes to facial recognition for age estimation or a number of specialist topics. However, I do not want to be complacent. There is, of course, a lot of work to do.

Stephen Almond: As a digital regulator that constantly has to look at different areas of emerging technologies that are processing personal data, it is a challenge for us. It is not a new challenge. We always find ourselves in the position where we are competing in the market for scarce talent that is valued extremely highly by commercial entities.

It will always be part of our business model to try to attract that talent into the organisation, recognising that the public sector will never be able

to compete with some commercial entities on salary, but it can compete on mission and impact. We find that we are able to attract strong talent into the organisation on that basis. Of course, like any public authority, if we have more resources we can do more, but that does not mean that we are not able to achieve good impacts with what we have right now.

Hayley Fletcher: We have a specialist data unit within the Competition and Markets Authority that comprises data scientists, those who are experienced in machine learning. We have significant internal capability, but I agree with the comments of the other panellists. It remains a challenging environment to ensure that we retain and attract staff with specialist skills.

It is not something that is entirely new to us. We also face this challenge when it comes to finding, say, economists and lawyers, because we are generally competing with the private sector. That requires us to ensure that we present the entire offer that we have. We have a strong mission and a purpose. We find that that attracts talent, and they can feel and know that the work that they are doing has a real-world impact. So far, we have found that to be a real draw for some of the most significant expertise out there.

Anna Boaden: This is probably one of the biggest challenges we face from a regulatory perspective. Our budget is small and has not been increased over the last 10 years, so it is a real-terms cut of about 30%. We absorb that across a wide remit, across equality and human rights, in the delivery of services and the public sector.

That makes it difficult for us to act in the way that we might like to and to bring in the technical expertise that would make us more effective. That is probably borne out mostly in our ability to take strategic litigation, because we do not have deep pockets by any means, and we expect the cost of litigation in this space to far exceed what our budget would allow us to afford.

Q89 Baroness Healy of Primrose Hill: As you are not part of the Digital Regulation Cooperation Forum, do you believe that it would be helpful to have an agreed standard for auditing LLMs and, if so, who should lead that work—regulators, government or third-party private sector? I am interested to know whether it would help you to access more expertise if there was a standard that everybody could relate to.

Anna Boaden: I will defer to my other colleagues on a standard. It is fair to say that we feel our relationship with the DRCF is good and we are working effectively with it, albeit a smaller player in our resource and scope in the regulatory landscape. We do not think that that is a specific challenge for us because we have good co-operation. However, as the work evolves, our thinking on that will as well.

Baroness Healy of Primrose Hill: Ms Fletcher, do you have a view on an agreed standard for auditing LLMs?

Hayley Fletcher: There is ongoing work with DRCF and partners examining the role that third parties could play to support exactly in the

way that you suggest. Certainly, I think that they play a role, but that must be considered in addition to other factors.

If we have a particular case or investigation, we have to consider the role that large language models have played in the round alongside the way that companies are behaving, their internal policies as well as the underpinning mechanisms of the large language models. Having some standards is an important part of that overall picture.

Q90 Baroness Healy of Primrose Hill: Mr Almond, I would like to ask you, because I know that you are an expert in sandboxes, I think. You mentioned the Vallance review. What progress has been made on developing this cross-regulator sandbox as recommended by Patrick Vallance? Do you have any updates for the committee?

Stephen Almond: Certainly. In September of this year, the Government announced that they would award the Digital Regulation Cooperation Forum £2 million to develop what is titled an AI and Digital Hub, a service to support innovators—whether AI innovators or digital innovators more broadly. Where they have questions that straddle one or more regulatory boundaries, they are able to ask them of us and get a joined-up regulatory response.

We are clear among ourselves that, if we think that the answer to some of the challenges of AI is to have strong, individual, context-specific regulators or sector-specific regulators, we need to work together seamlessly to ensure that our regulatory approaches are joined up. Having that direct joint interface with the public will be critical in making sure we get that right.

Baroness Healy of Primrose Hill: Dr Teh, do you have anything to add on that?

Dr Yih-Choung Teh: No, but I would agree with Stephen that it is an important endeavour to promote innovation in an area that is potentially hugely beneficial to us. Reflecting on Baroness Harding's question, I am not so concerned that large companies will be able to navigate the regulatory scheme because they have good resources to do so, but the smaller innovators, the start-ups, are a concern. Piloting to see whether we can make that easier is an important step.

Q91 The Chair: Before we move on, could I ask both the ICO and Ofcom for your views on an agreed auditing standard for the technology? It seems to me relevant to what you were saying in response to questions from Baroness Harding about the difference between the impact of this technology on a service down the line versus a product that might come to market, and which would sit within your sectoral responsibilities. Do you have views on whether there should be a standard? As Baroness Healy was asking your colleagues from the CMA and EHRC, if you do, who should be responsible for it?

Stephen Almond: Perhaps I should go first. As I think you would expect me to say, I welcome all forms of industry-led governance that help deliver our regulatory objectives, be they standards, certification

schemes, or codes of conduct that can be developed. We at the ICO try to promote the development of such schemes. They are complex. I do not think there will be one standard to fit all circumstances in this scheme, because there will be a difference between, for example, management standards and governance standards that may be needed in this area, along with technical standards.

The development of standards has to be led and driven by the market. It has to reflect the processes by which the market wants to get that governance right. We play an active role in trying to steer and engage with some of the standards bodies to see those develop. We are keen to see continued progress in that area.

Dr Yih-Choung Teh: I would agree. There are benefits of standards. I was thinking about some of our work in telecoms security where technical standards are important globally to try to advance that. That is attractive. The only thing I would add is, given that AI is a set of technologies being used across a whole range of sectors, there is a question about being clear on what sort of algorithmic assessment we are talking about and what it is doing. I want to hold on to the objective in that, coupled with then trying to drive for as much standardisation as possible within those objectives.

The Chair: Would you not see the need for auditing the standards for large language models?

Dr Yih-Choung Teh: Potentially. I am just making the point that I would want clarity on the objective of that assessment so that you make sure that the standards are working to achieve the outcome you want.

Q92 **Lord Kamall:** Very quickly, you touched on assessment algorithms there. One thing that we are hearing from this inquiry is that a lot of people are asking for transparency, not just of datasets but also of algorithms. Even if we have transparency of algorithms, who understands those algorithms and how do you make sure you transmit that understanding of the algorithms, so it is understandable to other people?

Stephen Almond: This is an area we will continue to explore. We have some good tools. For example, under the Online Safety Act, we have information notices to drive transparency over some of the governance processes going on around us as well as technical, or the ability to observe how the algorithms are dealing with certain datasets in real time to see what the outputs are when given certain problematic content. I very much take your point. Staring at lines of code is not necessarily the answer here. It is about real-world impacts.

The Chair: Thank you. We will move on. Lord Griffiths.

Q93 **Lord Griffiths of Burry Port:** This is constantly an education for me. I have heard and read a little about your respective sectors. Over the years, you have developed your methodologies, you have identified your problems, you have looked at how to deal with new stuff that comes your way, and all this is contained within your respective disciplines. Of

course, we had responses from 10 regulators. If they were sitting here, each of them would just enlarge the spectrum. I do not know if Meta is God to you lot, but Meta says that “the presence of multiple regulators in the UK could give rise to competing policy objectives, with significant challenges for companies and a possible detriment to development of the AI ecosystem”. That is definitely a danger, perhaps a risk or whatever you want to call it.

The other thing is this central risk function. I do not know how you set up a function, but it is clearly intended to pull together or to give a place where you can together identify gaps, for example, and address them. I imagine in your respective sectors, each of you has come up with gaps that you have had to deal with. How are the gaps that stretch across the spectrum, as it were, identified when you have 10 regulators? How would you take the necessary steps to address it? I would like to start at the human rights end.

Anna Boaden: At the moment, we have not identified gaps, but we see that collaboration and co-operation in this space as critical. We have clear principles around equalities and human rights, the public sector equality duty, and our legislation from 2010, which is very well established and tested. As I have said already, it is largely untested in this area because we are new to it, and it is evolving. This is a small part of what we do, but we see collaboration, co-ordination and being able to provide a consistent, reliable and clear message—not only to developers but to people interacting with AI technology—as an important way to mitigate some of that risk.

Lord Griffiths of Burry Port: I have no problem with that. When we were looking at the Online Safety Bill as it was and issues arising from it, we were astonished to see how many government departments and areas of work and so on are touched on by the emerging concerns and questions. It is about getting departments of government, that are used to working within their own parameters, to collaborate—collaboration does not just happen, does it? You have to identify the absence of collaboration as itself a gap. I am interested to know who says on the phone, “Hey, there’s a gap. Shall we talk about it?” Where do the gaps emerge from? I can imagine a gap in your field, your field, your field and your field, but those gaps that affect us all that then need to be addressed together collaboratively is the process that I am interested in.

Stephen Almond: Perhaps I could come in and illustrate with some real examples of when we are picking up the phone in this regard. For example, the ICO has known for some time that fairness considerations around data protection law are challenging for organisations. It is an area that we wanted to provide greater clarity to the market on. We did. Earlier this year, we provided new guidance on how to build fairness considerations into the AI design lifecycle, but we knew up front that any sort of regulatory position that we took in this area would have direct read across to considerations of discrimination in equalities law.

Immediately, we knew that right at the outset we needed to reach out and work with our colleagues at the EHRC. We needed to make sure that

any guidance that we produced had exactly the right signposts to relevant guidance from the EHRC and that any subsequent work that we have done—and we are now collaborating with the Centre for Data Ethics and Innovation on the fairness innovation challenge—would have to reflect our joint positions to the market.

Similarly, when the CMA put out its report on foundation models, we could see early on in the responses that people were debating the extent to which data protection law and competition law might be in tension in the debate around open source versus closed source models. We probably agree that the tensions are invariably overhyped in that space, but it is an area where we recognised that we would need to provide joined-up positions to the market. We very quickly said that this was an area where we would put out a joint position statement in the coming months to be able to say exactly how considerations around what you need to do under data protection law, and what you would need to do under competition law, would converge. I can assure you that we are very alive and responsive to these sorts of issues.

As I perceive it, the issue around the central risk function needs to be considered separately from the regulators. We need to play into and advise a central risk function that government sets up on what we see as potential gaps in our area, where we think it is not just a case of us joining up but there is probably a legislative gap in the space—we just cannot cover it with either of our powers.

We need to flag those to government, but of course inasmuch as there may be legislative solutions, it is really for government to consider whether to bring those forward. I am keen to see any sort of central risk function that is established here either be close to or within government such that it has the ear and the influence to be able to say, “There is a gap over there in how these regimes intersect. This thing needs to be closed off”.

Lord Griffiths of Burry Port: I hope it happens. I have to say that we produce our reports regularly to Parliament and, in a sense, we are looking for gaps that we can identify for people to fill too. Sometimes our findings take several months to be heard on the Floor of Parliament at all. The dynamic, if this were an example, is not particularly encouraging.

However, I can see that in this business that there are potential areas. The mood and the tone of your reports—the abbreviated forms that we have been given to brief us—are so varied. Ofcom is pretty glum about potential harms and the rest of it. The markets want innovation and proper competition as well. I am delighted to see human rights here. I go to Strasbourg to the Council of Europe, and it is nice to have a proper seat at the high table. It is the practicability of it all. It is a fine ideal. In other areas of our work, such as the press, for example, the voluntary bodies that are supposed to regulate the press leave a little bit to be desired, let us put it that way. Who is the watchman or the watchwoman? Who is making sure that those things do not get dropped? Do just the two bigger ones of you converse about a problem that you have in common without this being looked at across the board?

The Chair: There is obviously the DRCF. Were you going to—

Dr Yih-Choung Teh: Indeed. I would make a couple of points. As Stephen was saying, when it comes to intersections, that is something that we can deal with reasonably well through things like the DRCF, but it is not just limited to that. Thinking about online safety, that Act gives us the duty to consult with the ICO because of the recognition of the interface between some of what we are doing and data privacy when it comes to children or other issues.

The challenge of genuine gaps is bigger, particularly when you want to move at pace. That is critical in the current environment. Like Stephen, I think that it is right that government leads on this central function to try to understand this, not least when the risks that you probably worry about most are things such as national security, which do sit with government, building on things like the national risk register. Exactly how that will be put into practice is an excellent question and I look forward to seeing the Government's response to the White Paper.

Lord Griffiths of Burry Port: We have had a nice seminar. Thank you very much.

The Chair: Do others on the panel agree with Mr Almond that the central risk function as proposed by the Government should sit within government? Does anybody have a different perspective on that? No. Clearly, regulators always fiercely defend their independence from government. It is particularly interesting that you, Mr Almond, said that you think that it should sit inside government. There is a sort of "who" question whether the central risk function will rely on tech companies to inform them of how things are developing. If they are the repository of information about developments and you rely on that function to a certain extent to help you in the way in which you regulate, do you feel that that you have that sufficient gap between yourself and the risk function? It is particularly interesting for the ICO because you actually regulate the Government in their information. I wonder whether this conundrum has been considered by any of you.

Stephen Almond: First, to clarify my original comments around this, I said, "close to or within government". What is critical is that it is able to influence government around where, for example, changes to policy or legislation are needed to close off gaps that as regulators we cannot mitigate. It has to have that influence with government. I fear that making it too distant from government such that it can provide recommendations that can be easily ignored creates its own set of challenges and risks.

In respect of the way in which that function needs to operate, I would be extremely worried if only technology firms had its ear. I would certainly hope that we as regulators had the ear of a function like that, with us flagging up candidly where we see gaps and risks. That goes for a much wider range of stakeholders too, including civil society, which is very happy to highlight risks to us that they think either we should mitigate or where they think there may be gaps in our powers. I would see it as being critical to the success of any function like that that it has the input

of a wide range of voices. I simply do not think you can assess risk without having that full range of perspectives.

Q94 **Baroness Harding of Winscombe:** We are talking about the central risk function in the future tense. I am intrigued. The White Paper was published in March. This world is moving very fast. You all sound as if you are working hard building the capability. Am I right in my assumption that this central risk function does not exist yet, to your knowledge?

Dr Yih-Choung Teh: My understanding is that individuals in DSIT are working on this and we are talking to them. Exactly what shape that will take and some of the practicalities, I am yet to understand that.

Baroness Harding of Winscombe: How worried should we be that this is not obviously in existence yet and all these different participants cannot yet engage with it publicly? Do you think that government is being too slow to set this up?

Stephen Almond: We are keen to see progress in this area.

Baroness Harding of Winscombe: Very diplomatically put, Mr Almond. Thank you.

Dr Yih-Choung Teh: I would say that we are not struggling to engage with DSIT as it progresses its thinking and works towards its response to the White Paper, but there are clearly some questions still to be answered.

The Chair: I am conscious of time. I will draw it to a close and say thank you very much indeed to all four of you, and especially to Ms Fletcher for standing in at short notice.