



HOUSES OF PARLIAMENT

Joint Committee on Human Rights

Uncorrected oral evidence: Human rights and the regulation of AI (HC 1262)

Wednesday 17 December 2025

2.30 pm

Watch the meeting

Members present: Lord Alton of Liverpool (The Chair); Juliet Campbell; Lord Dholakia; Baroness Kennedy of The Shaws; Baroness Lawrence of Clarendon; Lord Murray of Blidworth; Alex Sobel; Sir Desmond Swayne.

Heard in Public

Questions 34 - 44

Witnesses

I: Ellen Lefley, Senior Lawyer, JUSTICE; Louise Hooper, Barrister, Garden Court Chambers; Dr Janis Wong, Policy Adviser, Data and Technology Law, Law Society.

USE OF THE TRANSCRIPT

1. This is an uncorrected transcript of evidence taken in public and webcast on www.parliamentlive.tv.
2. Any public use of, or reference to, the contents should make clear that neither Members nor witnesses have had the opportunity to correct the record. If in doubt as to the propriety of using the transcript, please contact the Clerk of the Committee.
3. Members and witnesses are asked to send corrections to the Clerk of the Committee within 14 days of receipt.

Examination of witnesses

Ellen Lefley, Louise Hooper and Dr Janis Wong.

Q34 **Chair:** Good afternoon and welcome to the 39th meeting in this Session of the Joint Committee on Human Rights. Today we will be continuing our interrogation of the human rights implications in artificial intelligence—AI. We will be joined in two sessions this afternoon by expert witnesses. In this first session will be Ms Louise Hooper, who is a barrister at Garden Court Chambers, Ms Ellen Lefley, who is a senior lawyer with JUSTICE, and Dr Janis Wong, who is from the Law Society.

In the second session—I will introduce them properly later—we will have Professor Roman Yampolskiy, who is the author of *Artificial Superintelligence: A Futuristic Approach* and director of the cybersecurity lab at the University of Louisville. Joining us online will be Professor Ethan Mollick, who is the author of *Co-Intelligence* and co-director of the generative AI labs at the Wharton School. He is a Rowan fellow.

Our committee's mandate is to examine the implications of public policy for British citizens. We are comprised of 12 members, six from either House, and we come from different political traditions. We scrutinise legislation and we conduct thematic inquiries such as this one. Our reports can be read on our committee's website. This is the second public session of the AI inquiry since we closed our call for written evidence back in September. Much of that evidence, including new evidence that we agreed earlier today should be published, is already on the website.

Today, members will examine the issue of the liability and accountability of AI systems. We will look ahead to what challenges we can anticipate in the near future, with a view to examining how to design regulation that effectively protects human rights.

In this first panel, we will hear from experts in the legal and justice sector about the impact that AI is already having and how the principles that the UK has signed up to in the Council of Europe's framework documents might be leveraged to address AI-specific challenges. In the second panel, our expert witnesses will explain their experience with and understanding of AI, and share their predictions for the future impact that the development of AI may have on society.

It is worth pointing out that our inquiry is not taking place in a vacuum, as a number of other parliamentary committees are assessing other aspects of AI, everything from copyright laws to the impact on employment.

Turning to our first panel, we have Louise Hooper, who is a barrister at Garden Court Chambers with a distinguished practice in public law, human rights and migration, and a strong international profile advising Governments, NGOs and intergovernmental bodies, such as the Council of Europe. She is widely recognised for her expertise in artificial intelligence governance, addressing ethical, legal and human rights

challenges posed by emerging technologies. Her work includes shaping global policy, drafting guidance for compliance with international conventions, and delivering training to legal professionals and policymakers worldwide.

Ellen Lefley is senior lawyer at JUSTICE, before which she practised as a common law barrister in civil, family and criminal law. She leads JUSTICE's current workstream on artificial intelligence, human rights and the law, and in 2025 published a rights and rules of law-based framework for AI in the justice system, which was co-authored with Sophia Adams Bhatti. Ms Lefley has experience on several advisory boards for external projects on behalf of JUSTICE, including the independent oversight group for the Home Office's 2025 public attitudes research into the use of AI in policing, and she contributed to the Cambridge University Press *Handbook of Generative AI and the Law*.

Finally, Dr Janis Wong is the data and technology law policy adviser at the Law Society of England and Wales. She leads on developments across artificial intelligence, data protection, cybersecurity and technology to support the legal profession and justice. Dr Wong holds an interdisciplinary PhD in computer science and was awarded best data protection thesis by the Council of Europe. Prior to the Law Society, she was a post-doctoral research associate at the Alan Turing Institute and research fellow at the Open Data Institute, focusing there on children's rights, data governance in education, AI ethics training for civil servants and addressing biases in healthcare.

Before turning to my colleagues, let me set the ball rolling by asking the panel a first question on the regulatory challenges. Perhaps Dr Wong would like to go first on this. Could you give us an overview of the biggest challenges that Governments are likely to face when seeking to regulate AI? Tell us whether the existing law is adequate, without going into granular detail because we will have a chance to do that later. Can it all be done by regulation rather than, for instance, having new legislation at all?

Dr Janis Wong: Thank you, Chair and honourable Members, for the invitation to give evidence today. Regulating AI is no easy task; otherwise, we would not be in this room. Through our work at the Law Society, our research and engagement with solicitors and other stakeholders, we see three key main challenges when it comes to the regulation of AI.

The first is balancing how much we wish to regulate. Under-regulation risks not having the appropriate safeguards to protect key values, going to the heart of the question, including those around justice and human rights. Over-regulation also perhaps inhibits innovation, making the UK, in this context, a taker as opposed to a maker in terms of technology and AI.

Secondly, when it comes to thinking about the challenges, a key question is whether AI regulation should be sector-specific or sector-agnostic. For

the former, regulating in a sector-specific way, in the context of the legal profession, for example, we are able to rely on and use the expertise of the sector to make sure that AI adoption and development is as useful and effective as possible. At the same time, this might result in fragmentation across different sectors, which may harm inadvertently when it comes to justice and human rights consideration.

Thirdly and finally, one of the questions is how AI should be defined. Using overly technical definitions may lead to regulations not being as effective or perhaps as useful in the short term. If AI is defined too expansively, that might involve regulating more traditional machine learning and low-risk applications of AI technologies and tools, which we also do not necessarily want to see.

At the Law Society, in terms of whether AI regulation will need new legislation or how we should approach AI regulation, we believe that any regulation must take a principled approach that promotes ethical consideration of the values that we believe to be true, including fairness, accountability and explainability, to ensure that we are bringing in the unique expertise of solicitors and the legal profession, leading by example. As a result, the Government need to consider underpinning regulation with these core principles, supplemented by assessments that ensure fairness and accountability in a measurable way.

Chair: Thank you very much. That was a very helpful introduction to this afternoon.

Ellen Lefley: In addition to what Dr Wong has highlighted, from the perspective of JUSTICE and what a rights-based approach to regulating AI means, we think there are some idiosyncrasies of artificial intelligence that are really important in understanding the challenge that it poses in terms of adequate regulation and protection of human rights.

First, by its nature, it is human intelligence mimicry, which can be impressive, even uncanny, with an unhelpful veneer of neutrality, and can be very tempting to resource-pressed public services that make important decisions.

Secondly, it is fuelled by big data, the processing of which tends to flatten human experience and make generalisations about groups of people rather than individuals. After all, they are statistical and probabilistic tools. This can embed bias, but, critically, it can also eliminate human dignity and autonomy, which are really key principles within the concept of human rights and what we are concerned with today.

Finally, the vast majority of AI is made by the private sector. The relationship between the private sector and the public sector is unavoidable in this area when we are considering its regulation. The private sector has commercial interests that will not always align with the protection of human rights, democracy and the rule of law. When AI is used by those in power to impact those who have none, these characteristics can cause harm in some areas more than others. Later, I

am sure we will get into what the high-risk areas are and where we consider the risks to lie in the justice system. That is an addition that I would make at this stage in terms of challenges.

Chair: Thank you. Again, that is a very helpful introduction. We are going to be having a session in the new year with the private sector to ask it some of those questions. This is helpful to us as a committee.

I will turn to you, Ms Hooper. You might also address the question of whether we have enough regulation. Can we do it all by regulation now or will it require a new Act of Parliament?

Louise Hooper: If I can start with the biggest challenges, they are who you are regulating, what you are regulating and how you manage enforcement. Those three questions are fairly fundamental.

If you look at what we have already and break it down into different groups, you have to look at business-to-business regulation, business-to-consumer regulation and public sector regulation. There are gaps. It is not my area of expertise, but I did have a look at consumer protection regulations. Arguably we could govern some artificial intelligence systems through those, but there is a problem with what we call a product and what is not a product.

We have some way in through procurement regulations. We have technical standards that we could use to govern those. The current Private Members' Bills of both Lord Holmes and Lord Clement-Jones are a really good starting place for looking at overarching new legislation that would cover both the private and public sectors and bring in those elements that my two colleagues are talking about in terms of transparency, accountability, explainability and fundamental rights-assessment processes. They would fill some of those gaps.

Chair: Some colleagues may want, later in the proceedings, to question you further about whether Private Members' Bills are the correct route to do this or whether the Government should be doing it.

Q35 **Sir Desmond Swayne:** Ellen, I was very interested in your second point, the flattening and embedding that you mentioned. Is that any different in principle from what we have always done with social sciences, particularly economics, where you rely on statistics and large numbers to predict behaviour in order to eliminate the caprice of individual human behaviour? Is there anything really new there other than the fact that we have machines to do it very quickly?

Ellen Lefley: We have machines to do it very quickly with, as I said, this veneer of neutrality when in fact all these processes make generalisations about populations, types of people and groups of people. If any of those conclusions, predictions or generalisations are going to impact how you treat a specific person, they need to be combined with individual circumstances. It is a core public law principle that a decision about an individual needs to be based on their individual circumstances.

I am focusing mainly on decision-making by people in power about those who are not. The state-individual decision-making process is the prime example of that. That relationship is the primary consideration of all human rights.

We must not forget equality law. We have laws that prohibit discrimination. We have laws that prohibit decision-making based on the type of person you are or the group that you belong to rather than who you are. Specifically, when we consider decision-making in justice, criminal liability and sentencing are things that cannot be done solely with statistics. There always needs to be that additional individual analysis and assessment.

Sir Desmond Swayne: The Government have set out their action plan for regulation, but we have also signed up to the Council of Europe's Framework Convention on Artificial Intelligence, although we have yet to ratify it. Are those two things consistent with one another?

Ellen Lefley: The AI opportunities action plan from January and the convention are very different pieces of work. The AI opportunities action plan is essentially an industrial strategy and its primary focus is growth. They are not particularly compatible. Human rights are not namechecked within that industrial strategy document. The phrase "human rights" does not appear in that document.

The need to protect vulnerable people does get one sentence, but it is caveated by, "That said, that mustn't get in the way of growth". We would say that is the wrong hierarchy. The only growth worth having is one that protects human rights. They are the basic standard.

Sir Desmond Swayne: At the very least, there is a significant agenda that we need to address.

Ellen Lefley: Yes, definitely. The other issue with that document is that, because the plan for growth is linked to AI and the embedding of AI within the public sector, it treats AI as the solution and then it looks for a problem rather than saying, "There are public sector delivery problems. We are not meeting people's rights as best we can; we are not delivering for the public good. How can we do this better?" That is the problem, and AI is one of the solutions. We think being problem-led, rather than solution-led with AI as the solution, is the right approach.

Chair: It might be helpful, in developing the question that Sir Desmond has asked you, if in writing, subsequent to the meeting today, you could give us some more examples of the incompatibility between the framework document and the Council of Europe proposals.

Ellen Lefley: Yes, I can do. I have mostly set out the thematic issues, but I can certainly follow up with more detail.

Chair: More detail on that would be helpful to us.

Q36 **Alex Sobel:** Louise, under the framework convention, states would have

to ensure that private actors respect human rights in the use of AI. How could this be achieved in a UK context? My personal view is that our legal framework is wholly insufficient, but would it be enough to rely on existing laws or would we need to introduce specific legislation?

Louise Hooper: I am inclined to agree with you. The starting point is that the state is the duty bearer when it comes to enforcing human rights and would be responsible if private companies breached individuals' rights via the Human Rights Act and other mechanisms.

I touched on this earlier. There are ways of directly incorporating a requirement for human rights adherence into procurement contracts, for example. That could be done through technical specifications or through explicit provisions within the regulations or the contracts that the Government use to contract with private companies.

In respect of consumers, as I say, there is a gap in product safety legislation. Human rights are not necessarily embedded into that. In business-to-business contracts, it is just left to the parties to work out what is reasonable and what is not. At the end of the spectrum, you end up with AI almost being built by default into products that were not AI products. That is not regulated at all. You get AI by default or human rights breach by default. That area ought to be looked at by government.

The private sector's view on trade secrets and confidentiality is often pushed as a barrier to regulation. I would ask you not to be hoodwinked by this. We do not need to know your commercial secrets, but we do need to know that you are using AI systems. We need to know what you are putting into them, what you are doing with them and that you are testing the outputs properly for human rights breaches. We need that to be transparent. I am getting really fed up of listening to people say, "We need transparency" and then not saying what it is. At a very basic level, that is what we need from business and government.

Many of those decisions are political. For example, you can code for individual fairness or group fairness but very rarely for both. Why are you making that decision? How are you mitigating the results of that in any AI system that you are building? In terms of confidentiality, we have adequate mechanisms that could deal with it, such as the ICO. The business and commercial courts deal with these issues in side hearings all the time.

The overarching point is that, yes, we probably do need some amendments to legislation to cover those risks.

Q37 **Baroness Lawrence of Clarendon:** Good afternoon, all. My question is for Dr Wong. What can be learned from the EU AI Act? Is it the right way to regulate this area? In practice, will UK companies have to comply with the EU AI Act if they want to do business in Europe?

Dr Janis Wong: If I may, I will incorporate a brief response to Sir Desmond Swayne's initial question around statistics and problem finding. That is particularly important in the context of justice because there are

considerations around what is fair for individuals as well as what is fair for the collective. It may be possible that in certain circumstances the use of automated decision-making and algorithms on an individual level is broadly free from bias or discrimination, very broadly construed, but there may be patterns on a collective level that we are not able to spot.

Furthermore, in thinking about fairness in the context of justice and the legal system, both procedural and outcome fairness are key issues that are being considered, which brings us to considering what we can learn from the European Union's Artificial Intelligence Act. The Act is certainly one of the first attempts, if not the first, to recognise the potential harms that artificial intelligence can bring and provide a regulatory solution to address that, which should be lauded. At the Law Society, we also believe that a risk and principles-based approach should be taken, should further regulation take place domestically as well.

To provide a very brief summary of the EU AI Act, the Act recognises that certain types of AI systems should be banned, such as certain uses of biometric ID, emotion recognition, social scoring as well as facial recognition databases. There are also specific transparency obligations for general-purpose AI systems and some limited AI systems, such as those that interact with individuals and people. More broadly, there are obligations around AI and digital literacy that apply more generally to different stakeholders. The Act came into force in August of 2024, but it is taking a staggered implementation approach through to August 2026.

More recently, taking more contemporary learnings, with the developments around the EU's digital omnibus and broader stakeholder engagement, concerns have been expressed around the challenges to compliance, touching on points that panellists have already mentioned, regarding uncertainty around the Act, the availability or lack of standards, other guidance and national governance processes, and uncertainty around the breadth of the horizontal AI literacy requirements put on organisations with regard to training staff members who use and engage with AI.

Specifically on the question around compliance for UK and domestic companies and organisations in practice, the EU AI Act will apply to UK companies and organisations if they are providers or developers of AI technologies and tools, and possibly to users or deployers of AI technologies and tools. The EU AI Act is extraterritorial. Any organisations that license, deploy or have outputs in the EU will fall under the EU AI Act. The extent to which the provisions and articles apply will depend on the organisation and the extent to which it engages with those means.

For the legal profession, for example, we would expect that large multinational law firms will be not only complying with the EU AI Act but certainly advising their clients, organisations and individuals on the legal ways to be compliant with the Act and the ethical ways to do so in order to meet their client obligations and professional obligations.

Put another way, when we are considering whether government should be regulating in this space, if, for example, the UK Government were suggesting that certain purposes of AI that are being banned in the EU will be permitted in the UK, that would not only be a deviation on the AI regulation position but also a deviation of values related to justice and human rights considerations.

Q38 **Alex Sobel:** Ms Hooper, what specific concrete mechanisms are there for ensuring that transparency is in place? How would they work in practice?

Louise Hooper: For government in particular, the recommendation that I would make is to put the algorithmic transparency recording standard on a statutory footing. It has been mandated, but we know it is not being used properly. Not all algorithmic use is on there.

The Online Safety Act is good in terms of the labelling of artificial intelligence-generated content. We need better election protection, so clear labelling of AI content for elections. The use of algorithmic systems by private companies should be made clear. There are some suggestions on how to do that on the website of the Financial Conduct Authority, but I am happy to come back to you on that later.

Chair: If you could write to us on that, that would be really appreciated.

Alex Sobel: We will have an elections Bill first thing in the new year, so maybe we could amend that.

Chair: There we are. You are already prompting some of our colleagues to think about amendments.

Q39 **Baroness Kennedy of The Shaws:** I am going to come to you first, Louise, if you do not mind. I want to talk to you about data protection. We have updated the rules in the Data (Use and Access) Act in relation to automated decision-making. Has the right balance been struck in that legislation? What do you think?

Louise Hooper: If you talk to anybody from the NGO-type sector, they would say that the Data (Use and Access) Act has watered down the existing protections that we had in a way that is damaging and unsuitable for automated decision-making and AI systems for the people who are impacted by them.

Baroness Kennedy of The Shaws: They have been watered down. That is what I have heard as well. There is a sense that the protections that were there before for data protection would have worked better than what we have now. Was that a deliberate decision or is this the law of unintended consequences? How did it come about that we have ended up with less rather than more?

Louise Hooper: That is a political question, which it might be better for politicians to answer rather than me.

Baroness Kennedy of The Shaws: I am concerned about it. I am interested in the extent to which we are bowing the knee to the United

States of America and its new absolute determination to be unregulated, and whether we are, because of our need for trading relations and not wanting to fall out with our great ally, allowing the United States to dictate to us. The complaint before was that the European Union was dictating to us. We are now bowing the knee to an even greater and more belligerent or determined partner.

Louise Hooper: I am going to answer that in two parts. The first part is a general comment on companies' views of data protection, the need for data and the fact that they do not know what they have, whether it is covered by Data Protection Acts or how they are using it. I cannot talk about the actual action, but in the US there is currently some issue around that. The EU AI Act has started to enable separate processing of special category data for bias testing by producers but not by equality bodies, the very people who are charged with working out whether we are being discriminatory.

In relation to your second point, which is a significantly more political point, I have extreme concerns. From June of this year, we have had one executive order that requires any debiasing measures to be removed from tech and we have a national AI strategy, which came out about a week ago. In my view, the very basis of what tech companies in Silicon Valley are now required to build is arguably directly contrary to British democratic liberal values and human rights, in particular in their attack on gender, race and LGBT issues, and diversity more generally, which has to be factored into any regulation that we are looking at, as well as data protection issues.

Baroness Kennedy of The Shaws: I feel quite despondent about this because we are going to have to stand firm and be bold, but I am not sure that we are seeing signs of that. We are up against a very difficult and rather rigid attitude from those who are basically at the heart of the creation of this stuff. Would that be right?

Louise Hooper: Yes.

Baroness Kennedy of The Shaws: Does data protection law do enough to ensure that public authorities and private actors in our country inform the public that AI has been used? Is it currently adequate?

Louise Hooper: At the moment my view is no. The protection for knowing whether an algorithmic system is being used depends on what type of data is used in the system, not on the system itself being used. Proxy data, which can be as discriminatory as direct personal characteristic data, is not covered by our current data protection law. Dr Wong might be better than I am at answering this question.

Baroness Kennedy of The Shaws: Dr Wong, maybe you have more information for us.

Dr Janis Wong: Perhaps to broaden the question slightly, in the context of thinking about human rights and justice, just informing is insufficient.

Data protection law does that. It is one vehicle for doing so and it is a broader means of regulating AI more generally. Certainly, we have found this throughout the developments of previous iterations of the digital information Bill and the current Act. However, data protection alone certainly does not prevent unethical use, particularly where individuals may not even know that AI tools and systems are being deployed in the context of justice in particular.

I will spell out more clearly to the committee the ways in which the Data (Use and Access) Bill has changed our relationship with automated decision-making. With those changes, no individual consent is required for ADM. A legitimate basis, as a lawful basis for data processing, can be used to justify and explain how and to what extent personal data is being used by AI systems. Again, we would reiterate that simply informing is certainly not sufficient.

Q40 Baroness Lawrence of Clarendon: AI has emerged and we have been trying to catch up with it. While you were speaking I was trying to understand who polices AI, thinking about the questions that Baroness Kennedy was asking. Who polices that? All these things are out there. Who polices them? Who keeps an eye on this to make sure that AI is not running too far away without us catching up with it?

Ellen Lefley: This comes back to the differentiation that Dr Wong made right at the start about how our approach to AI can be sector-specific or sector-neutral. The sector-neutral approach would be to have a new authority that has that policing oversight of AI across different sectors. That is not the approach that has been taken here. That is the EU AI Act approach. The UK Government have decided on a sector-specific approach with the idea that the policing—we may replace it with regulation—is incorporated into existing regulatory remits. There are pros and cons to that.

The sector-specific approach means that existing regulators know their existing area. They can really home in on those context-specific risks. The CMA is the body that understands the AI risks within that consumer protection remit. It also poses some really important questions about where gaps emerge, where we do not have existing law or regulators to oversee the uses of AI. One of those gaps has emerged in our recent work in policing. I have the report here; I can share it with the committee. We found a lot of regulation and governance structures around policing, almost too many. It is a disaggregated landscape.

In terms of who is policing the policing use of AI, we in fact found there to be a complete wild west. The 43 different territorial police forces are able to deploy AI in 43 different ways, with differing basic standards of accuracy and no central testing capacity on police data. They are often reliant on marketing claims of reliability. In the context of facial recognition, they are writing their own policies on what a human-rights-compliant use of live facial recognition technology looks like in terms of where and how it can be used, potentially creating 43 different approaches to human rights protection in policing use.

There are areas that are constantly cropping up across different sectors where there may be ostensibly some regulatory oversight, but there are still gaps. In our sector-specific approach, that needs to be a matter of constant vigilance. In Lord Holmes's Private Member's Bill, there is this idea of an AI authority that has an ongoing duty of vigilance over the sufficiency of existing regulatory schemes and existing law. Do we need additional consumer protection laws? Do we need additional product safety laws? We think that is a good middle ground where you embrace the benefits of a sector-specific regulatory approach, but you also have someone charged with the job of spotting those gaps and filling them.

Chair: That is quite a lot now on the regulatory approach and oversight. People who have been in correspondence with the committee have raised concerns about what happens to them when they are adversely affected. Lord Murray is going to ask you a little bit about that.

Q41 Lord Murray of Blidworth: Before I go to my question on redress, can I just pick up on that last question? On Monday, there was a parliamentary conference on AI held by the Parliamentary Assembly of the Council of Europe here in London. At that gathering there was a German academic who described something called the pacing gap, with which you no doubt are familiar. It seems to me that he was making quite a powerful point that the regulatory systems in AI always try to play catch up to the problems posed by AI, given the pace of development. I just wonder whether the panel have any particular thoughts on the best way of having a dynamic regulatory system that avoids the creation of a gap. Is that gap inevitable?

Ellen Lefley: The pace issue is a huge one. In our most recent report, we considered this in terms of ever-emerging technologies and their use and policing versus our ability to catch up and set standards. It was very clear within our work that saying that primary legislation is the only way is just simply not workable. That is a very long and involved process.

Rather, we can set principles and outcomes that regulation is trying to achieve—for example, putting on a statutory footing the current ethical framework for AI that is within government policy—and potentially some key ways in which that is assured, such as algorithmic impact assessments within the public sector. Those are the headlines and key linchpins of regulation.

Within policing, we identified the uses of codes of practice. There was precedent for that. Our Police and Criminal Evidence Act sets out the key principles in policing and the detail is filled in within codes of practice on the conduct of PACE interviews, for example. That kind of balance allows for democratic input into the key principles as well as a faster mechanism through which updated rules can be maintained. Technical standards can also be updated through codes of practice.

In addition, if we cannot do everything by primary legislation because it is just going to take too long, we need to think about whether that creates a democratic deficit and whether we need permanent public

participatory mechanisms that are not just our representative parliamentary system, our representative democracy. We proposed a permanent public participatory mechanism, such as a citizen's panel or a citizen's assembly, which could iteratively inform policing's adoption of AI. That was a key part of our recommendations as well.

Louise Hooper: We are probably about 10 years behind. We are only now talking about technology that was built years ago. That is really problematic.

By its nature, regulation takes time. One of the difficulties of adopting a risk-based approach to regulation is that you are regulating backwards. You are saying, "We know these things create risk, so we will regulate them". You then wait until the next lot of things come along that have already caused harm. Adopting a rights-based approach might be more effective.

In terms of what Lord Alton said earlier, I agree that a Bill from government is required rather than Private Members' Bills. My suggestion is that you could potentially combine the two, have those as an overarching framework, and then leave some of the detail to regulators, which can react more quickly and more appropriately as time goes on.

Dr Janis Wong: I will add two brief points to that, moving slightly beyond thinking about regulation and its pace. By design, regulation is slow. That is not a bad thing.

There are two things that we are working on at the Law Society. First, our response to the Government's plan for AI growth labs, which my colleague Mia Leslie is currently leading on, proposes a cross-economy AI sandbox. Our position there is focusing on how regulation is not necessarily a barrier for innovation; rather, it provides legal certainty and to some extent enables early adoption of technology.

Secondly, in the context of the cat and mouse game of regulation and development, in our responses to the Government's industrial strategy, one of the key points that we highlight is the importance of an open technology market, particularly when thinking about lawtech and legal technologies, to encourage domestic entrepreneurship and leverage technology as a tool for access to justice to ensure that we are a competitive country and a competitive market in the context of the digital economy. We have also been working closely with LawtechUK, which is supported and funded by the Ministry of Justice, and thinking about what a blueprint for developing AI and technology solutions for access to justice could look like.

Lord Murray of Blidworth: I am sorry to cut you off, but we are limited in time. This topic could take days. It is very interesting. If you have any further thoughts after the meeting, please do write to us on that question.

Can I ask you now about redress? If there is a violation in the existing

legal framework in the context of AI, is it useful? Is there enough in the current framework to establish where liability lies for harms caused by AI? Is there a need for specific legislation in relation to that?

Louise Hooper: This is one of the biggest difficulties. In terms of the public use of AI, it really needs to be thought about. I say that because at present insurance companies in the US are asking their regulators for permission to exclude widespread AI-related liabilities because they cannot quantify the risk. If that is the case, that would mean private companies do not have insurance for the risk that they are shifting on to the people they sell their products to, including government. My concern around that is that you are then pushing responsibility on to the taxpayer. That is not reasonable.

One thing the EU AI Act does reasonably well is look at what producers are required to do, what deployers are required to do and how they interact with each other. That is very much about how you apportion responsibility. Much of that will be in contractual terms. I will just do a plug for the Society for Computers and Law, which has done a lot of work on contractual terms in AI contracts precisely to look at this.

Overall, this is an issue that will require some kind of legislation in terms of who should be held responsible and where the burden of risk should lie. This is really about who has to insure against it.

Lord Murray of Blidworth: Before we go to your colleagues, can I just ask you one supplementary question to that? Should there be different liability considerations at different stages of the AI life cycle? I am conscious that the EU AI Act is a whole-life-cycle approach, whereas in the model envisaged by the Government it is all end-user-focused.

Louise Hooper: There absolutely should be liability apportioned throughout the life cycle and not just end of use. It is not appropriate for deployers of AI systems to bear the burden of risks created up stream, which they may not properly understand because they have not been explained properly, or for those very large, multinational and almost ungovernable entities to have the power to say, "You will pay out for what we get wrong". It is just utterly inappropriate.

Lord Murray of Blidworth: You could envisage very complex causation arguments.

Ellen Lefley: I would add two points, the first on liability and the second on access to redress.

First, to echo the point that Ms Hooper has just made around the spread across the supply chain, that absolutely needs to be in place. There are some examples that we can use where we have done that already. I mentioned knives in our written evidence. It is illegal to manufacture certain types of knives. It can also be illegal to sell certain types of knives. With respect to legal knives that are not banned, it is unlawful to sell them to certain groups of people. There are numerous laws that

restrict what people can do with knives. They cannot harm or threaten people. They cannot carry them without a good reason. There are existing examples where things pose risk and we can segment liability across the supply chain. That absolutely needs to be the approach with respect to AI and the risks that it poses.

On redress, the preliminary point is that a good offence is the best form of defence. Taking proactive steps to stop human rights harms accruing is much better than relying on redress and complaint mechanisms. Taking that into account, when we are talking about individuals seeking redress for harms that have been inflicted on them by AI, it is important to separate public harm and privately inflicted harm. With respect to public harm, it comes back to our data protection point. You have to know that AI has been used in the first place in order to make a complaint. That is such a simple point, but it is not reliable at the moment. You cannot mount redress to an AI-related harm if you do not know that AI has been used in the first place. You do not know what you are asking for. You do not know what you are looking for.

The changes that were made to the automated decision-making rules were accompanied by safeguards. If ADM has been used, someone has a right to information about it. If AI has been integrated into a decision-making process that has a human in it, there is no commensurate right to know. That transparency is a big block.

We also need to understand access to redress within the current context of barriers to access to justice in all other circumstances. No approach to human rights accountability for AI-related human rights harms can be thought to be adequate with simply theoretical access to judicial review alone.

Lord Murray of Blidworth: That is only in respect of public bodies anyway.

Ellen Lefley: Indeed, it is not a safeguard in itself. There has been a vast amount of automated decision-making and AI use by public authorities and police. We have had one case about the police use of facial recognition. That is not because it has not been used but because of the inaccessibility of judicial review, the amount of time and the cost burden on individuals.

We come back to transparency. In order to complain that use of facial recognition has been discriminatory, you would need to see an evaluation of what is being used. Two weeks ago, the Home Office released evidence that shows that retrospective facial recognition algorithms, which have been used for over a year by police, are 100 times more inaccurate for black women than for white women. They have been used. We do not know on whom. There is no obligation to disclose the use of that facial recognition tool in a particular person's case. There is no way for a particular person who has been arrested to know whether that has informed the intelligence case against them. I will stop talking.

Lord Murray of Blidworth: I am sorry to cut you off.

Chair: We do need to move on. These are really interesting points that the committee will want to explore further, but I am conscious that we are trespassing into the time of our second panel. I do not want to cut my colleagues short. My life would not be worth living if I did.

Q42 **Juliet Campbell:** A couple of you have touched on justice already. My question is to Louise Hooper. The Minister of Justice's AI action plan for justice sets out that AI should be adopted responsibly across the justice system. Does AI really improve access to justice?

Louise Hooper: It has potential beneficial uses, but it has to be very carefully thought through. My own experience with legal software has been uninspiring so far. I am aware of recent research that suggests that even the best legal research software has error rates of between 17% and 30%.

Baroness Kennedy of The Shaws: That is right.

Louise Hooper: My research for this hearing showed several very clear errors in my AI-enabled legal research. I am an experienced and now ageing lawyer. If you let litigants in person loose on what is basically a multi-guess system, you are going to end up with the problems that we are seeing in court. I am not sure that, in the long run, it yet saves us time or benefits access to justice.

What it possibly does is give us a sop to say that we do not need to invest in lawyers or courts. If you talk to judges, they want working sockets to plug their computers into. We are a really long way off.

Baroness Kennedy of The Shaws: There are some basics that we could do with getting sorted first.

Louise Hooper: There are some basics that need to be in place. Rules-based systems that answer specific legal questions for litigants in person can be helpful. In a commercial context where you might be more willing to accept outcomes because it saves you money, it could be helpful. I would very strongly suggest, in any circumstances, trying it in commercial courts first and on commercial litigants first. If they have problems with it, do not use it on people. That would be the worst.

That is what I think generally about AI and public service. We are starting on people with no power. We need to start with the people who have all the power: the people who are avoiding tax, the people who are evading tax, the people who are in the commercial courts. If they have problems with you imposing AI solutions on them, it is probably for a very good reason. If it works for them, there might be a reason to roll it out in welfare benefits, housing allocations or the criminal justice system. Start with the people who have the power to complain about it.

Baroness Kennedy of The Shaws: That was a great answer. I just want to comment that that was a very important answer.

Ellen Lefley: We think access to justice can be assisted through AI technologies, but we have also identified a real risk that the majority of the AI lawtech market at the moment is targeted at lawyers. There is a real risk that there is lots of AI for lawyers that makes their lives easier, but the people who really need quality additional assistance do not get any and the access to justice gap widens.

The other point that I wanted to mention with respect to the impact of AI on justice and the fairness of proceedings is the fact that artificial intelligence already—and this will continue to increase—makes up part of the evidence that courts are considering. For example, in criminal proceedings this may be AI forensic tools. It may be consumer AI that is introduced, such as evidence from a Fitbit that someone was awake at a particular time. It could be deepfake evidence because we all have deepfake machines in our pockets now.

AI is also being considered for integration into disclosure duties. After the Horizon/Post Office scandal, this is an area that we really need to be careful about. That scandal showed that we do not have adequate facilities currently within our criminal justice processes to make sure that the most disempowered can properly scrutinise the reliability of computer evidence that is coming into courts.

Chair: Thank you very much. You have certainly given us food for thought about the fairness within the system. We will need to reflect more on that.

Q43 **Alex Sobel:** Dr Wong, the Law Society and Bar Council have issued guidance to solicitors and barristers about their use of AI. How concerned are you about the effects on the justice system if AI is used?

Dr Janis Wong: Thank you for acknowledging that we have guidance on our website specifically on generative AI and how the profession should use it. Our most recent update in September acknowledges some of the cases that have gone through the courts and addressed some of the courts' concerns with regard to the use of AI by the legal profession.

We are aware of and hope to address the impact that the misuse of AI, knowingly or unknowingly, has on the justice system. For example, the Law Society responded to the House of Lords Public Services Committee inquiry on the use of AI in interpreting and translation services in the courts, which we are happy to send through and provide further information on. There are questions and challenges around the accuracy of these AI tools, whether AI can be an interpreter and how it interacts with other requirements, such as for the free assistance of these services and for human expertise and knowledge.

We are working towards and addressing these issues appropriately. For example, our flagship 21st century justice project, which also touches on the previous set of questions around what a justice system and human rights-based approach looks like in the modern day, recognises the consumer harms, while recognising that a lot of modern generative AI

and AI tools are particularly useful for improving efficiencies and specifically for looking at the unstructured data that perhaps may emerge from bundles or other legal documents. The use of those tools can increase the speed with which justice can be done for specific clients. There are concerns that need to be addressed from both a technical perspective and certainly a legal one.

Chair: Thank you very much indeed. Of course, you can continue to correspond with us afterwards on any of the questions that you think might need to be expanded on or issues that you want to bring to our attention. I am going to turn to Lord Dholakia for our last question.

Q44 **Lord Dholakia:** This is the last question. If you had to choose, which two recommendations would you like this committee to make to the Government? Can we have a brief answer from each one of you?

Ellen Lefley: On transparency, mandating the algorithmic transparency recording standard not just for central government but all public sector bodies is important. I would also include in that first one about transparency letting people know when AI has been used in a decision that is made about them.

As I said, offence is the best form of defence, so I would integrate proactive and concrete human rights assurance mechanisms. Let us start in the public sector. Lord Clement-Jones's Bill on that, particularly the obligation to conduct an algorithmic impact assessment, is a very good place to start. If I was to choose only two, they would be those.

Louise Hooper: The first would be the two Bills that exist as Private Members' Bills. I will incorporate those into one. The second one is to be brave.

Chair: We will try.

Dr Janis Wong: Mine are less pithy but I hope still brief. The Law Society has two recommendations. First, prioritise solicitors and the legal profession, not only us on the panel but others who have lived experienced, knowledge and expertise in advising on and leading by experience in terms of how AI can be used in a responsible and ethical way, particularly in the regulated professions.

Secondly, consider the first question around regulation at both a domestic and an international level. The UK Government have already developed momentum, been part of international conversations and led on the discussions about what it means to incorporate and consider justice and human rights through the lens of technology. The hope is that the Government should continue to foster those conversations and discussions, and lead on the international agenda to make sure that human rights considerations are considered in the context of technology and AI.

Chair: I speak for the whole committee when I say thank you to the three of you for addressing these issues this afternoon so thoroughly and

with such interest. You have given us a great deal to go away and think about. Dr Wong, Ms Hooper and Ms Lefley, thank you very much indeed for being with the Joint Committee on Human Rights today.