

John Lehane Memorial Lecture

Supreme Court of New South Wales

**The implications of the development of artificial intelligence on
the judiciary and the legal profession**

18 August 2026

Sir Geoffrey Vos, Master of the Rolls

Introduction

1. I pay my respects to elders, past and present, and to all indigenous people present this evening.
2. The Preface to the first edition of Meagher Gummow and Lehane's Equity Doctrines and Remedies, published in March 1992, is characteristically modest. It ends by saying that the authors hope that the book “whilst not a *κτηῖμα ἐς αἰεί* [*ktēma es aiei* which means ‘a possession for all time’] will **not** be considered in the phrase attributed to Lord Westbury, “difficult to read, disgusting to touch, and impossible to understand””. Lord Westbury was describing old conveyancing documents, but, parenthetically, I have the same aspirations for this lecture.
3. The phrase *κτηῖμα ἐς αἰεί* comes, of course, from Thucydides on the History of the Peloponnesian Wars, who said he had designed his account of the war as a possession for all time rather than as a temporary prize.
4. Sir Frank Kitto helped colourfully to explain in his Foreword to that 1st edition of Meagher Gummow and Lehane why the lecturers in this series have been invited from England:

Equity remains ... the saving supplement and complement of the Common Law at the ends of the earth as in England, prevailing over the Common Law in cases of conflict but ensuring by its persistence and by the very fact of its prevailing, the survival of the

Common law and the enduring influence of English jurisprudence as a whole in the history of civilisation.

5. It might be hard to do justice to John Leane's equitable legacy when addressing *[t]he implications of the development of artificial intelligence on the judiciary and the legal profession*. But I shall try.

The background

6. Your Chief Justice, Andrew Bell, and I have been exchanging views on the implications for our legal systems of rapidly advancing AI ever since we first met here in Sydney at a SIFoCC¹ full meeting on 20 and 21 October 2022. As it turns out, that was portentous timing, because ChatGPT was first launched as a public research preview just one month later on 30 November 2022.
7. My contribution to the SIFoCC meeting in Sydney that year was also somewhat portentous, because I said three things about how commercial litigation would, in future, be undertaken.
8. First, I said that commercial litigation would: “use artificial intelligence and algorithmic technologies to distil relevant legal precedents, to make sense of increasingly large amounts of case data, and to bring about mediated solutions”.
9. Secondly, on the subject of mediation, I said that: “[a]rtificially intelligent applications can suggest resolution of interim issues and issues arising along the litigation path as they arise”.
10. Thirdly, I said back in 2022 that work should be “done now to work out how AI can be used to handle the ever-increasing volumes of data produced in commercial

¹ The Standing International Forum of Commercial Courts.

cases, and how the process can be streamlined by integrating every dispute resolution method within a single online process”.

11. I am telling you all this just in case you might be thinking that I had changed my tune since the advent of generative AI.
12. Conversely, Chief Justice Bell has described his approach to AI as “relatively conservative”.² He points to: (i) the quality, integrity and currency of the underlying “training” data leading to “rubbish in, rubbish out”, (ii) embedded biases in algorithms, (iii) algorithmic failure, (iv) issues of confidentiality, privacy and data protection, (v) cyber-security risks, (vi) the still unsolved issue of hallucinations, (vii) the apparent confidence of Gen AI responses and the equanimity with which such confident responses may be changed, withdrawn or qualified when challenged, (viii) AI sycophancy and seduction, and (ix) users’ ignorance of the algorithmic underpinnings of the platform being used and the related problem of the black-box.
13. Whilst Chief Justice Bell and I may differ on some of these points, I suspect that, ultimately, we have more in common as to the future adoption of AI within legal systems than divides us. At the end of this evening, you may be able to judge whether that is true or not.

Three essential realities

14. I want to explain in this lecture three essential realities.
15. First, I shall explain why I believe that lawyers and judges have no choice but wholeheartedly to embrace and adopt the use of artificial intelligence so as to make their work outputs more efficient, cost-effective and more quickly deliverable.

² See Chief Justice Bell’s 2026 Harold Ford Memorial lecture delivered at the University of Melbourne Law School: *Corporate responsibility and directors’ duties in the era of Artificial Intelligence* at [26].

This is all about access to justice and ensuring that the law remains relevant to the new generation of litigants and businesses brought up on social media and AI. (**Embracing AI**).

16. Secondly, I want to try to explain why it is an essential element of the rule of law to resist the introduction of autonomous judicial decision-making by machines without the fully informed consent of the parties. This is all about the preservation of people's right to an independent and impartial decision by a human judge. (**No machine-made judicial decisions without consent**).
17. Thirdly, I shall explain why I think that one of the most important things that judges and lawyers should now be doing in every self-respecting private law system is to be developing a clear strategy to explain to individuals, SMEs and the financial and the industrial markets how their legal system will deal with the critical questions of legal liability for harms caused by the use of AI. (**Legal liability for harms caused by AI and agentic AI**).
18. These three essential realities may seem haphazard, even unrelated, but they are not. The advent of generative AI has changed, I think forever, the parameters of modern life and of the legal environment. The law is going to be, already is, critical to the confidence that other industries have in adopting new technologies and to the way in which the utilisation of AI develops across consumer and business services. Understanding where the liability lies if AI causes loss and damage will play a significant part in developing that confidence.
19. Closer to home, the legal community needs also to lead the way in demonstrating how the adoption of AI can bring improvements in legal outcomes for those that lawyers and judges serve. But as this adoption progresses, it will bring with it challenges. The legal community needs to show that it has answers to the rule of law issues, as well as to the liability issues, that are raised when it becomes possible for machines to decide cases far more quickly and economically than can

be done by any human judge. We all need to consider how to avert any risk that our societies will drift into a position that might put the fundamental rights of individuals and businesses at risk.

A little more background

20. Before I come to the three essential realities in a little more detail, let me just mention five background propositions.
21. First, the usage of AI, for good or ill, has massive take-up across the globe. It is used in every professional environment. It is true that there are very distinct sovereignty issues caused by the fact that the most ubiquitous LLMs are US-based. It is true also that there are further major issues caused by the fact that these US-based LLMs are English-language centric, creating problems for users in Africa, India, and South America to name but a few. And it is true that further potential future issues are caused by the US or New York law foundation to the usage of these LLMs. But whatever may be the repercussions of these issues as they play out, I remain convinced that there is no turning back, so far as the massive take-up I have mentioned is concerned.
22. Secondly, it is the case, whether we like it or not, that every litigant, whether a consumer or a business, in fact feeds their legal problem into Claude, ChatGPT, Google or CoPilot before they ever contemplate consulting a lawyer. And many, if not most, lawyers do the same, using mainstream LLMs or Harvey or Legora or other specialised LLMs according to their means and their need to maintain their clients' data confidential.
23. Thirdly, LLMs in particular, and AI in general, are at an extraordinarily early stage in their development. We are not yet 4 years from the launch of ChatGPT in November 2022. The LLMs will inevitably get better and more powerful. Quantum computing will transform them. Limiting my observations to the law, we need, to say the least, to grasp the nettle and quickly to establish red lines for the adoption of AI

tools before the inevitable economic advantages of its ubiquitous usage answers every question we may – perhaps too late – be asking.

24. Fourthly, there is one crucial feature of AI that is often forgotten. AI can process data and reach conclusions in seconds. But humans are much slower than machines. Humans take days or months to process, understand and to accept conclusions or outcomes produced by machines. If AI reaches decisions affecting humans and their lives and businesses very quickly, we should not expect humans to come to terms with those decisions in the same time-frame.
25. This fourth point is the reason why we should not believe those who say that AI will spell the end of lawyers. It will not. Humans will need other humans, who are trained in legal reasoning and can explain legal advice and legal decisions to them in slower time, even if that advice and decision-making is given or assisted by machines.
26. My fifth and final introductory point concerns the importance of the justice system as a whole. The advent of new technologies raises questions about the **relevance** of justice and the legal system itself to our societies. Justice systems are provided by the state **for** the people. Justice systems cannot ignore the way people live their lives. Generative AI and social media have become essential elements in the lives of large swathes of our populations. Lawyers cannot just dismiss those facts and wish it were otherwise. It is a real imperative for all our legal communities to find ways to deliver justice within our new environment that are as valuable and relevant to the people of today as our old ways were to the people of yesterday.

Lawyers and judges must embrace AI

27. My first essential reality is that lawyers and judges must embrace AI. I doubt that that is really a controversial proposition. The questions raised relate really to the degree of reliance that can safely be placed on the products of LLMs.

28. In my view, many of the arguments that seek to find fault with current LLMs ignore a number of things, some of which I have already mentioned.
29. First, current LLMs are at an early stage and are improving very rapidly.
30. Secondly, there is a tendency to gloss over what LLMs do well: for example, when they save hours of drudgery by creating case summaries, translations, transcripts, case chronologies, shortening and improving prolix text, and undertaking basic legal research that gives lawyers a significant head start.
31. Thirdly, the classic “hallucination” argument is, I believe (if you will forgive the bluntness), entirely flawed. The proponents point to those lawyers who have included non-existent cases and citations derived from an LLM in their court documents or judgments. Any lawyer who allows false or inaccurate material to be used under their name deserves all the opprobrium and professional sanctions that are applied to them. That is so whether the false or inaccurate material emerges from the lawyers themselves, the lawyer’s pupil, a judge’s JA or from an LLM.
32. The first rule of the use of AI within the legal process is that none of its work product must be used without human checking. In the words of the Guidance that we judges have issued to the Judiciary of England and Wales: “The accuracy of any information you have been provided by an AI tool must be checked before it is used or relied upon”.³
33. Fourthly, the risks that LLMs pose to confidentiality are, at least in major part, a human problem. As the Guidance given to judicial office holders in England and Wales explains:

³ Artificial Intelligence Guidance for Judicial Office Holders (in England and Wales) of 31 October 2025: <https://www.judiciary.uk/wp-content/uploads/2025/10/Artificial-Intelligence-AI-Guidance-for-Judicial-Office-Holders-2.pdf>

Do not enter any information into a public AI chatbot that is not already in the public domain. Do not enter information which is private or confidential. Any information that you input into a public AI chatbot should be seen as being published to all the world.

34. Cyber-security is a universal problem caused by the fact that so many of our citizens live their lives online. It will not be ameliorated by setting our face against the utilisation of AI, even when it is economically advantageous. It is axiomatic that, in using AI, all lawyers and judges need to understand how the LLMs operate and the risks that they pose in terms of both confidentiality and cyber-security.
35. Fifthly, I think that there is much exaggeration in the complaints about the quality of the work product of LLMs. I use LLMs daily to find things out – always checking the outcome against documented material before using it. But I find the quality generally remarkable. Our experience in the UK and elsewhere is, as I have said, that almost every litigant puts their legal problem into one or more LLMs before they ever go near a lawyer. My colleagues have found that many of the AI-generated pleadings and skeleton arguments promulgated by Litigants in Person are of generally higher quality than was ever the case before. Of course, these AI-generated arguments need careful scrutiny, but so did the lengthy material generated without AI by the Litigants in Person of a previous generation. The Times newspaper reported yesterday that AI is creating a “kitchen sink” approach to employment claims in the UK. I recognise that approach from my legal practice that terminated now 17 years ago.⁴
36. Whilst it is also clear that more cases are being brought than was the case before the ubiquitous availability of LLMs, this change actually produces improved access to justice – even if

⁴ <https://www.thetimes.com/uk/law/article/employment-tribunal-claims-ai-hjjk86dvt>

we need to develop mechanisms to use AI to process the additional cases.

37. Sixthly, let me mention bias. Of course, LLMs are prone to bias. But then so are humans – both consciously and unconsciously. The important thing is for humans to be aware that they may be biased and for humans also to realise that machines may be biased.
38. We cannot turn the tide, and we should not try to do so. Lawyers and judges should be in the vanguard. They should showcase how AI can reduce costs and delays, without reducing standards, and how it can safely and advantageously create increased access to justice. If I had time, I would explain the ways in which AI can improve pre-action dispute resolution within what we call in England and Wales, the Digital Justice System – but that will need to wait for another lecture.
39. So, to conclude on the first essential reality, we already accept deficiencies and risks from human actors. We need to be wary of illogically holding AI to a different standard. Standing out against proper careful usage of AI within the legal process risks damaging the trust and confidence that citizens and businesses have in the legal community and in justice itself. Of course, lawyers and judges must be careful to ensure that AI is used responsibly and appropriately.

No machine-made judicial decisions without consent

40. If you accept my first premise, it is inevitable that AI will be used extensively in the legal system by lawyers and judges alike. I also believe that it is inevitable that AI will be used, not merely to facilitate and inform judicial decision-making, but also as the judicial decision-maker in some fields in some places. But this inevitability does not come without consequences.
41. The first thing to understand is that humans already happily accept **some** machine-made decisions. For example, eBay

decides 60 million disputes on its platform every year within 24 hours. Of these disputes, 90% are resolved, we are told, without any human intervention.

42. It is equally likely, I think, that litigants will come to have confidence in machine-made decisions once those decisions are shown, as I believe in time they will be, to be **as** or **more** reliable than human decisions. Machine-made decisions will, of course, be cheaper and quicker to obtain. In such circumstances, litigants may well **choose** to have their cases determined in minutes by a machine at little or no cost rather than waiting months or years to receive the decision of a human judge. This may start with decision-making models capable of largely algorithmic determinations such as the assessment of personal injury or medical negligence damages or with share valuations in minority shareholder disputes.
43. The AAA-ICDR (the American Arbitration Association – International Centre for Dispute Resolution) (AAA-ICDR) announced in December 2025 the availability of an AI arbitrator for two party, document-only construction disputes up to a certain value.⁵ It is not quite clear whether that ceiling is US\$25,000 or US\$50,000. But that may not matter. The model includes a “human in the loop”, whereby the AI arbitrator will evaluate the merits of claims and prepare draft awards, but a human arbitrator will review it to test the reasoning and evaluate the decision, in order to safeguard trust, transparency and due process.
44. There is seemingly no reason of principle to prevent competent commercial parties from agreeing to such a dispute resolution method. Before too long, we will see how popular such mechanisms prove to be in practice.
45. I can imagine that the attractions of machine-made arbitration decisions may be reduced in intensity by the possibility of challenges to the enforcement of the award under the public policy provisions in Article V(2)(b) of the New York

⁵ <https://www.adr.org/ai-arbitrator/>

Convention.⁶ In New South Wales, there could be a challenge to the arbitral award under 34(2)(b)(ii) of the Commercial Arbitration Act 2010 (NSW), on the grounds that the award was in conflict with the public policy of the State.⁷

46. But either way, consensual machine-made decisions are one thing, and non-consensual ones are quite another. There are, I think, important reasons of legal principle which might prevent the courts from imposing machine-made judicial decisions on non-consenting parties.
47. In my part of the world, article 6 of the European Convention on Human Rights, applied domestically by the Human Rights Act 1998, guarantees to everyone “a fair and public hearing within a reasonable time by an independent and impartial tribunal established by law”. I have expressed the view, with which some disagree, that an AI can probably never constitute an “independent and impartial tribunal” within article 6.
48. In Australia, there is no equivalent directly applicable legislation, but the common law has a similar effect. So far as statute is concerned, section 80 of the Commonwealth of Australia Constitution Act provides for jury trial in Federal matters as follows: “The trial on indictment of any offence against any law of the Commonwealth shall be by jury”. It is hard to imagine that a machine could ever properly be regarded as a “jury”.
49. And article 14 of the International Covenant on Civil and Political Rights, which is applicable but is not directly applicable in national law, here in Australia, provides that: “[i]n the determination of any criminal charge against him, or of his rights and obligations in a suit at law” that “everyone shall be entitled to a fair and public hearing by a

⁶ Article V(2)(b) of the Convention on the Recognition and Enforcement of Foreign Arbitral Awards (New York 1958).

⁷ In the UK, the challenge might be on the basis that there was a serious irregularity under section 68(2) of the Arbitration Act 1996 in that: “the way in which [the award] was procured [was] contrary to public policy”.

competent, independent and impartial tribunal established by law”.

50. Whether or not it is correct to say that a machine could never be regarded as “a competent, independent and impartial tribunal established by law”, the right to a human judicial decision in crime, civil, family and administrative cases is deeply entrenched in our constitutional settlement. It would incontrovertibly require legislation to change it. So, however, attractive it might be for the parties to a personal injury claim to have their damages assessed in minutes by a capable AI, human judges cannot just be abandoned in the courts without legislative change. And such legislative change will take time. And there will be significant disagreement as to what should be changed and how. Should, for example, a final appeal to a human judge be retained in every case?
51. The further impediment to machine-made judicial decisions is that the necessary dedicated LLMs have not, in most countries, been either created or tested for reliability and absence of bias.
52. The problem that I find most difficult in this area is the line that is to be drawn between machine-made judicial decisions and machine-assisted judicial decisions. One gets a variety of answers from Claude, ChatGPT, CoPilot and Google to a question about whether machine-made judicial decisions infringe the rule of law, but all of them accept that they might. But then they assert that **machine-assisted** judicial decision-making does not have the same issues of legal certainty, equality before the law, judicial independence, transparency and due process.
53. As I have said in several early lectures on this subject, the line is hard to draw because AI is rapidly improving and will be able to give “assistance” to judges in a much shorter time than human judges will need to take to check what they are being told.
54. There is probably no ‘rule of law’ problem with judges using Harvey or Legora to assist with legal research, but these

programmes are also capable of answering the ultimate question raised by a legal dispute as well as just providing background legal information. Like the AAA machine-driven arbitrations, the Chinese model for criminal sentencing retains the human in the loop. But as confidence in the programmes grows, one wonders how much judicial time and attention will be paid to checking the machine's advice. Ultimately if a judge is presented with 1,000 decisions, economic and time pressures will inevitably lead to judicial approval of the machine's recommendations.

55. It is, therefore, quite predictable that there will be pressures on all our justice systems to move towards allowing machines, at least, to **advise** judges on the decisions they have to make. Even article 14 of the European Union's AI Act⁸ does not ban machines advising judges. It only makes justice systems into "High Risk AI Systems" that must be: "effectively overseen by natural persons". Whilst responsible guidance issued by judiciaries in many countries cautions that the judge is always personally responsible for their decision-making, whether or not assisted by AI, it is possible, even perhaps likely, that some judges will be tempted to use their LLMs outside the guidance.
56. I am fairly sure that the use of AI to **advise** judges as to case outcomes, in common law countries at least, is likely to be some way ahead. Human caution and circumspection is likely to mean that changes will be slow at first, but when they come, they may accelerate. And it is quite likely, I think, that humans will come to accept, on economic grounds, machine-enabled resolution of small disputes at first instance.
57. SIFoCC is doing some ground-breaking work in this area, seeking to draw up some principles that can achieve broad consensus as to what is acceptable, and what is not acceptable, as to the use of AI in relation to machine-made or machine-assisted decisions in commercial and arbitration cases.

⁸

<https://artificialintelligenceact.eu/>

58. I believe the legal community needs to prepare itself urgently. We need to have an informed debate about where the lines are to be drawn so as to preserve the humans' rights to an independent and impartial human judicial decision, save where there has been freely and fully informed consent given to machine-made decision-making.

Legal liability for harms caused by AI and Agentic AI

59. My third essential reality is to deal with the question of legal liability caused by LLMs and agentic AIs. These questions are likely to occupy much of the time of the lawyers and judges of the future. It is surprising that, in the nearly four years since the launch of ChatGPT, so little thought was given to this question, until quite recently. Reasonable certainty about liability for AI harms will, as I say, give users confidence to adopt AI solutions across commercial life.
60. The European Union was quick to regulate AI with its AI Act.⁹ Article 14 of the AI Act requires, as I have said, “high risk AI systems”, including justice systems, to “be designed and developed in such a way, including with appropriate human-machine interface tools, that they can be effectively overseen by natural persons during the period in which they are in use”. On 11 February 2025, the European Commission announced the withdrawal of its proposed Artificial Intelligence Liability Directive. Instead, the European Union is relying on the formal adoption of its new updated Product Liability Directive (EU) 2024/2853 on 23 October 2024.¹⁰
61. As the European Law Institute's (ELI) first supplement to its Draft of a Revised Product Liability Directive published in June 2026 explains, there are three potential routes to liability

⁹ <https://artificialintelligenceact.eu/>. The EU's AI Act came into force on 1 August 2024.

¹⁰ The Product Liability Directive was brought into force on 8/9 December 2024.

for defective software and AI under the New EU Product Liability Regime.¹¹

62. Providers of social media platforms and AI chatbots may be liable under the new Product Liability Directive, depending on the true meaning of the word “product” that has been newly expanded to include software and AI systems. ELI expresses the view that strict liability may exist where: (i) the digital offering itself qualifies as software that has been placed on the market or put into service, (ii) the digital offering is characterised as a service but is supplied through defective software operating in the background, and (iii) where a digital offering that primarily constitutes a service may nevertheless qualify as a component of software that is clearly a product, such as an app used to access the service. For this evening’s purposes, it is enough to know that these complexities are likely to lead to legal difficulty and even confusion – the very antithesis of legal certainty.¹²
63. What then of the position in the USA? I will leave to one side the gigantic lawsuit starting in California today with the attorneys general for four US states as plaintiffs and Meta as defendant: there are claims there for the violation of local child and consumer protection legislation.¹³ In the US more generally, the position is complicated by the First Amendment to the US Constitution protecting free speech. It seems that AI algorithms of all kinds, including LLMs, are regarded as embedding human judgment and will, therefore, *prima facie* at least, fall within the scope of the First Amendment. The question of whether the First Amendment will protect AI

¹¹ https://www.europeanlawinstitute.eu/news-events/news-contd/news/eli-publishes-supplement-on-liability-for-defective-software-under-the-new-eu-pld/?no_cache=1&cHash=9337faff2f8233da2951de3b732ece3a

¹² Article 22 of the General Data Protection Regulation provides that: “[t]he data subject shall have the right not to be subject to a decision based solely on automated processing, including profiling, which produces legal effects concerning him or her or similarly significantly affects him or her”. This adds a further potential layer of uncertainty to AI decision-making within the EU (but also in the UK).

¹³ <https://www.techpolicy.press/four-states-take-meta-to-trial-in-california-the-latest-front-in-sprawling-multidistrict-fight/>

producers from claims is much contested in cases across the USA with decisions going both ways. The US Supreme Court has yet to reach a definitive ruling.¹⁴ Here too, therefore, it can be said that there is legal uncertainty and a lack of clarity as to liability for harms caused by AI and agentic AI.

64. In the UK, the UK Jurisdiction Taskforce (UKJT) (which I chair) published in July 2026 its Legal Statement on Liability for Harms caused by Artificial Intelligence under English law.¹⁵
65. It might seem obvious to common lawyers that the normal principles of duty of care and the law of negligence would apply to harms caused by AI. Likewise, it might seem obvious that misrepresentations, breaches of contract and defamation in the context of AI would be treated in the same way as such claims made in other contexts. But, even though the lawyers and judges may think that that is clear, many others have expressed genuine uncertainty about how those well-established common law principles would be applied in the AI context.
66. The UKJT's legal statement answers three legal questions in the simplest possible terms: first, when will the developers of large language models and others involved in the AI supply chain be liable under English law for harms caused by AI? Secondly, in what circumstances will the professional, business and individual users of these models be liable for such harms? Thirdly, when will professionals be liable for **not** using AI to protect their clients' or their patients' interests?
67. The statement expresses the view that businesses are likely to be liable if they use an AI chatbot to provide false information to consumers, causing them harm. And app developers may well be liable if they market chatbot friends, for example, to

¹⁴ See, for example, Francesca Procaccini and Wendy K. Tam on *Algorithmic Speech* 115 Calif L Rev (forthcoming 2027).

¹⁵ <https://lawtechuk.io/ukjt/liability-for-ai-harms-under-the-private-law-of-england-and-wales/>

teenagers without taking sufficient care to build in controls that regulate how the bot responds to questions in sensitive areas such as self-harm.

68. I would not presume to suggest the position under Australian law to this audience. But ChatGPT tells me that the answer to the question of whether an LLM is a product under Australian Consumer Law is “probably, in many contexts—but it has not yet been decided by an Australian court”. I was interested to see a report last week in the Guardian about an AI agent improperly bumping its principal, Andrew, up the queue for booking gym classes here in Australia.¹⁶ The article reported that the director of the University of Melbourne’s Centre for AI and Digital Ethics had acknowledged that “there is ... a good deal of scope for legal and ethical “murkiness” around actions perpetrated by automated agents”.
69. Lawyers and judges will be instrumental in working out who, in our modern world, is to be responsible for damage caused by new technologies.
70. I cannot over-emphasise how important it is to have early clarity as to liability for harms caused by AI under widely chosen private law systems. Where there is clarity, private law systems will be chosen in order that commercial parties can have all-important legal certainty if things go wrong.

Conclusions

71. I will freely admit that I have gone on too long. But I wanted to explain, as best I could, the plusses and minuses of generative AI for our legal systems, our lawyers and, most of all, those they serve, those that use legal services and need to

¹⁶ https://www.theguardian.com/technology/2026/aug/13/ai-agents-arent-legally-responsible-for-any-harm-that-they-cause-experts-say-so-who-is?CMP=Share_iOSApp_Other

have their commercial and consumer disputes resolved by domestic courts.

72. I do not suggest that the story is one of undiluted positives. And I accept that there are big issues as to how we utilise AI within the legal system and as part of the process of reaching judicial decisions. But none of that should lead lawyers and judges to appear technophobic or, worse, Luddite.
73. Legal systems and legal professions depend on public trust and confidence. It would damage that public confidence if those providing legal services and dispute resolution services appeared protectionist or self-serving in the face of ever-more-capable AI systems that can provide reasonable, if less orthodox, assistance to lawyers and judges.
74. We need to work with the grain of technological advances. Lawyers and judges need, as I have said, to be in the vanguard. To ensure legal certainty in the face of rapid technological advance, they need to work out the circumstances in which there will be legal liability for harms caused by AI. There needs to be an open debate involving the entire international legal community as to what decisions can and what decisions should not properly be assisted or taken by machines rather than human judges. These questions need to be seen as connected, not disjointed.
75. The ways in which our justice systems resolve disputes and deliver just outcomes must move with the times. If they do not, they will risk becoming irrelevant. An irrelevant justice system would be the greatest of all possible threats to the rule of law.
76. I hope to have given you some food for thought, even if not a *κτῆμα ἐς αἰεί* (*ktēma es aiei*) or a possession for all time.
77. I look forward to your questions.