AI Literacy in the Legal Context

Thank you, and thank you for that very generous introduction. And I also want to thank both you and Christoph as well for organizing my visit to Hong Kong. It's lovely to be here. So, hello everybody. I'm looking forward to the rest of the panels as well, and just finding out a little bit more about what's actually happening in Hong Kong. I did speak at Peking University in September, and again, I've freshened up a little bit in terms of what's going on in the Mainland. The developments, though, in Hong Kong are really and have always been quite idiosyncratic. At the moment, I'm an expert advisor to the Victorian Law Reform Commission, and in one of those roles, I'm actually looking at governance, having just written my university's AI policy, which includes setting up a governance framework. I also echo the comments, but I won't spend a lot of time talking about governance here, except to note that it's no longer sufficient to simply assume that an IT manager or a chief information officer can manage the ethical and other issues that emerge with AI. You actually need a much more sophisticated framework. And I congratulate some of the big law firms who've actually done that and have brought in a multidisciplinary team to consider ethical issues. I won't mention, or perhaps I will. There is one large corporation that actually has an AI advisor, ethics advisor, that is, in fact, a form of AI, but that's another story. And I have my doubts about a fox looking after the chickens.

So, we've come a long way in the context of courts and their use of technology. And it's always good to reflect on a few things before I get into the body of the talk. We've only had smartphones for 17 years, which really isn't very long. And for those of us who are older, of course, we remember having maps and having a whole lot of things that we no longer can even imagine. Let alone the birth of smartphones. It's been about 29–30 years that we've had technology in courts and, again, across the world it varies significantly. Some places really aren't digitized. Some places have been digitized for many years. Some law firms are digitized and paperless, but other law firms aren't. So, there is a huge difference in terms of how people have adopted technology, and I think that's important to remember in the context of my comments as well. It's only been 25 years since we didn't have to have a dial-up modem. If you actually click on that link, you'll hear that sound, for those of you who feel like they want to have a memory of what it was like to pull out the cord and stick it in and make that sound. We've had about 37 years of supportive technology that included, in the early days, actual hard copy videos that courts and others would produce, that people could put into a video machine, and they could watch what went on, so that they had a better understanding about how things could work in

courts. For 11 years, we've had fairly good textual analysis, and it gets better and better all the time. We've had 26 years of AI-using branching technology. Now, a lot of techno folks wouldn't necessarily regard that as being AI these days, but it's certainly AI when I was involved in two very large projects, one in 2003 and one in 2007, where we built massive decision trees that could give an advice on whether or not a person was likely to be found guilty in relation to a whole series of crimes in Victoria. That was pretty simple technology. It's basically “Yes”, “No”, “Maybe”. And then you branch out and you go through the rest of it. We've had 10 years of more sophisticated AI, about 9 years of wearable tech, and I'm going to comment more on that in a talk I'm giving tomorrow, but I will mention neurotech as well today. I sit as an advisor on the Australian Human Rights Commission with respect to neurotechnology at the moment.

And then I've got a question mark about what will happen over the next decade. It's always a bit dangerous to be a futurist. You know, every now and again, I say something, and I think I shouldn't have said that, because it might not happen. So, I'm going to touch briefly before I go into the tech, to talk about legal education, just to contextualize the two panels that are following this. In about 2000, in my own jurisdiction, and I suspect in Hong Kong as well, much of the legal education that was around focused on content. That is, you learnt legislation, you learnt cases, you had a problem-solving approach, and it was very much a lecture style of engagement with students. Earlier than the year in the US, they had already moved to some very different problem-based learning methods, and they had even begun to focus on and begun to look more at clinical legal education. My own Law School has been a leader in clinical education for more than 25 years, and that makes a big difference in terms of your understanding about technology, because you actually work on it in a clinical setting with real clients, or sometimes not with real clients, but building solutions for real clients. There's little focus on project planning. That's a much bigger focus in law school education now, and little focus on automation, although mostly there were reasonable precedent libraries that had been developed by most big law firms.

In about 2005, you've got internationalization, and there's this expectation that students will be doing very different sorts of research, and then they'll be doing much more comparative work. And there was still an expectation that most law students would learn skills on the job, and so there was a focus

on content. Then you've got dedicated tech law classes. They began to emerge in the sort of early 2000s. You actually saw internet law. You saw other courses which focused more on the legal implications, rather than the “How” of it, or the “What could students do?”. And the conversation about “Should we teach students how to code?”, began to develop, particularly in Australia. Thankfully, we didn't teach them how to code, because it's all a bit redundant now. But still, doing data analytical work is actually quite important. As there are lots of free courses, there's not necessarily a need for law schools to offer those courses. We did see this big integrated-skills-based learning focus, which had already, as I said, been a big focus in the US. And you've got different courses popping up in US universities. User-centered design, for example, at Stanford. Major projects at Stanford which really were about how we could interrogate the data about what's going on in courts and in the justice system, to improve and make better processes. And I just want to say a word about that, because often what I'm seeing in the AI and tech space, and courts at the moment, is that you simply replicate processes that don't actually work properly in the first place, before actually asking, “Well, is there a better way that we could do this, and what would users say? And how could we design this so it actually meets the objectives it's supposed to meet?”

So, I tend to break up AI in 3 ways, and when I do that, what falls into each category varies a little bit every year as things move on. But I will talk here about supportive AI, which is about, you know, how you enhance human capabilities and capacities. You don't need very complex governance structures when you've just got supportive AI. On the whole, that can be done within an existing organizational framework. When you begin to use replacement AI, which takes over specific tasks, that actually means that you've got to have much more accountability, much more oversight. And then you've got disruptive AI, which I'm going to touch on as well, which is much more complex in terms of challenges, and you need a comprehensive governance framework. I would actually suggest that all courts at the moment should be building a comprehensive governance framework, otherwise they will be left behind by the AI revolution. So, what am I talking about here? This is where you begin to get law courses actually focusing on supportive AI, that first category, which I'm going to go in. There's an expectation by around 2010 with the explosion of clinical programs that students will actually begin to understand how to use supportive AI. Now, a lot of the supportive AI was pretty ordinary. It's AI that you use to support your grammar, to produce documents, to create chronologies, to cut and

paste what was in there to begin to summarize. But by 2010, again, very early days in terms of tech and law students' understandings, and we did see, of course, by 2015, a huge shift towards online learning that kind of dropped away. It was really thought that everything would be online. Then we had COVID, where everything had to be online in terms of video conferencing. But the focus on online education has really waned over time. So, some students love online learning, and some universities are very good at it, but others really prefer to have that social interaction and the development of ideas that comes through having that social interaction.

And then 2020, you've got COVID. That's where we get our online classes and we begin to see Gen AI. I was already using Gen AI in 2019, before ChatGPT, and in fact, wrote part of the preface of my book on AI and judges. You know, of course, I disclosed it, but it was so badly written, really, but it was a reflection of the time in that the sources that it looked at were clearly the sources that tech programmers thought people should be looking at. So, you can just imagine it was not very thoughtful. At the same time, we've got driverless cars and this kind of promise. And I think, when we talk about AI and the kind of opportunity and dystopian sort of divide, the black and white divide, part of it is there because of the tech hype as well. People get disappointed. You know, we're not really seeing that many driverless cars out there on the roads, and I think it's a bit unrealistic that we'll actually see them for a while. I did write an article on this and tort law reform a few years ago, and it was abundantly clear then already that there was a lot of promise. Yes, and I know that there are places in China where we've got good driverless cars, but it's going to be a while when we're actually looking at aging infrastructure and the reality of what happens in a lot of cities.

We began to see drones emerging. And I think you can call this the innovation stage from 2020 to 2025. This is where we've had a lot, in the last five years, of really beginning to go up a curve. Now, I have to say, because, you know, I hang out with techies. My son's one. My daughter's one. You know, it's a family piece. Pretty much everybody who works in the tech field considers that we are only at the beginning of the upward steep curve. We are babies in terms of our tech development. Things will really change over the next years, and they will change significantly. So, I think this is interesting to think about what direction that will go in, and whether it's something that we should be scared about or whether it's something we should be excited about or happy about, and maybe it's a mixture

of all three. And this is why I think, increasingly, people are talking about guardrails and governance and how we can manage some of the challenges that we have going forward.

Okay, so that was a funny slide. AI, that's already here. So, our supportive AI, that's what we've already got. We've got remote engagement, we've got Gen AI, and the Gen AI is improving all the time. We've got basic apps, and our students are building apps. You know, that's part of what they do in our law clinics. They'll actually build an app. So, we had students build, very successfully, a domestic violence app at my own Law School, in consultation with 2 wonderful staff, and it was directed at the offender. So, it was directed at behavioral modification as a supportive tool. So, these are the sort of things that students are doing. And, I think even our "Know the Law" app was really exciting. You know, you could easily translate it into another language so that international students could actually access basic legal information. You've got digitization, which means that you can have massive tech reforms, because without digitization, you don't have them. I spoke recently in a conference in Japan, and I was really surprised to hear that so much there isn't digitized in their court system. They're still handwriting some things, which just amazed me. In a way, it's not a surprise, because the handwriting is so beautiful that you would want to, you know, retain a sort of culture of handwriting. But it's really hard to look at digitization and really hard to look at tech reforms if you don't have tech. Of course, this will change as there will be more focus on what you create from a person saying or a bit of video. So, you know, query what will happen to text in the next five years.

Then you've got supportive AI, being sort of devising, summary, collation, doing research, doing presentations, and you've got the, what I would refer to as chatbots. But agentic AI is a little more sophisticated, so I tend to refer to chatbots as being quite simplistic, able to answer questions, but not able to solve problems, and I'll give you an example soon. But there is this growing expectation that you will get 24/7 service with chatbots, and that's happening in law as well. So, these changes have implications for law firms and significant implications. And this idea of having a responsive client interface as well, so that a client can go in, can be updated on what's going on with their matter, can find out if anything's happened, can also have a discussion potentially with an AI assistant, although that's a bit more advanced again. And then you've got all the variation of generative AI. And there's just some of them, but there are literally dozens now. And of course, the big publishers, legal

publishers have been very much engaged. So, Thomson Reuters, Co-Counsel, and LexisNexis have a model. The extent to which they've been rolled out fully in Hong Kong, I'm not sure about. But in some jurisdictions, they're very much rolled out and very much part of what's been going on. All of these have different uses, different error rates, and different levels of hallucination. Some will reference. So, Perplexity references. What else do I use? The Stanford one references. So, some of them are actually quite good, even in my own judicial work, I'll occasionally ask a question of, say, Perplexity, if I'm looking at country information, and I want to find out, for example; you know, what happens to an illegitimate child in Malaysia or a particular state in Malaysia, I can look it up on Perplexity, and I'll get references, and the references will take me to articles so I can check those. So much easier to do that, to do that research, compared to even a few years ago. Again, you know, the reliability issues are there, and there are other Gen AIs. I've just put a section up there of the ones that I'll tend to use more than others.

And then you've got, which is really quite interesting, big shifts in China with the micro court. And I did have a colleague pull up some more recent stats for me on this, but they're not great. But again, you've got about 72% of cases completed in less than 15 minutes. It's done entirely on your mobile phone. I think there are some issues with the micro court, but the micro court itself has actually serviced, and it's very hard to see, by March 2020, 437,000 cases. So, the use of AI and technology in the court system can really vary, and sometimes for simple matters, supportive technology can be sufficient to actually deal with a large range of more simplistic matters.

Today, this morning, I was drafting a continuum on judicial use of AI. And what might be in a kind of gray zone, what is the no zone, and what's in the okay zone? It's actually quite a challenge, but most jurisdictions around the world have not looked at replacing judges with a form of AI. Some jurisdictions have. I'm not suggesting that the micro court necessarily does that. And there are questions around how you triage and how you don't. And there are, of course, democracy and other political issues with replacement of judges. And then there's some supportive technology that is banned. So, this is myself and Professor Laurence Boulle. Does anybody know him? They probably would know Prof. Laurence Boulle from when Nadia lived in Hong Kong. And this is us at our campus at Newcastle, just mucking around with some facial recognition software. This sort of software is

more or less banned, really, in Europe now. Facial recognition software; it's a real no-no, but I actually think it's really interesting software, and it's what I would call supportive software. And when I've used this with judges, so I rolled it out with some Singaporean judges a year or two ago, and I gave them a whole lot of visual cues, and then I tracked what their faces might tell somebody if you were using visual cues, and some of them looked, you know, when you're sitting on the bench, and you have to be a bit aware of it at the moment, you tend to not want to show much emotion. You know, you're trying to really be supportive and not show much emotion, because you don't want to give sort of an indicator of where you might be heading, other than to encourage people to talk to you. However, the judges in Singapore, when we use the facial tracking systems, were actually giving expressions which showed that they were disgusted, appalled, and it was completely unconscious. Now, these are reasons, I suppose, and I'm going off on a bit of a direction, but these are reasons why perhaps we should not actually be looking at banning everything in the way that perhaps the EU has suggested in the context of high risk. Because I actually think it's very useful, from a training perspective, for people to have this kind of input. And you might rule it out and say, “Look, it's not that accurate, and it's not that great”. Even so, if it makes you think about how things work in a human-to-human sort of context, it can be useful. Now, you can see that my face up there was analyzed with the facial recognition tracking as being a cross between happily surprised and disgustedly surprised, which is probably close to actually how I feel about this software. I have major concerns about, of course, the police using it, and about it being used in other ways. However, just because you might rule it out of some applications, that's actually a governance question about what goes in and what goes out. It's also a question about risk, rather than having, you know, one uniform category.

Okay, AI that is emerging and developing, and this is replacement AI. This is really where everything has been for the last year and going forward for the next five years. And when I'm talking about replacement AI, I'm talking about agentic AI, paralegal AI, accounting AI and AI workers. And this is where everybody is investing and spending their money. And it's already developed in some fields, and it will grow enormously. This is probably an area where there will be the most significant impact on courts, tribunals and on legal practice, and I don't think, as a profession, we've necessarily engaged sufficiently with how to use agentic AI and how to actually build agentic AI in a way that makes sense

for us. It's now so easy to build your own specialized large language model. It doesn't take much. It's all available. My son built one with International Property Law and Trade Practices Law. It took him two days to specify the libraries that he was going to use and draw upon, and it was actually pretty coherent, the responses that you got from Gen AI. So, that capacity to build closed large language models has just gone astronomical in terms of the use. We've also got the creation of scenarios, creation of clones. We've got the development of a lot more in robotics, a lot of investment going on in relation to robotics, and much more algorithmic decision-making, which is more complex than it was in the past, and, of course, more problematic. A lot of what courts are dealing with now, the decision before a matter came to court may have actually been made, or at least developed, through algorithmic decision-making tools. So, a lot of insurance companies use algorithmic decision-making, and have for a number of years now.

Now, this is a photo of a hotel that you check into in Japan and it's a bit of a joke, really, because when I show it to people in courts and I say, “Is this how your front counter should look? Should you have these two robots at your front counter?” It's the robot hotel. And I don't know if anybody's ever stayed there, but you check in and you never see a human. And for some reason, there is a dinosaur on one side. And this robot girl who can, you know, speak in different languages to you. The robotics craze is a very interesting one, because these are some of the companies that are leading the robotic surge. And, you know, you've got Spot, everything from Spot the robot dog, and then you've got the humanoid robotics, you know, from Ameca and Sophia, and then you've got the whole sort of robotic systems. Probably, most of the tech funding is really focused on agentic AI and robotics, the new funding that’s gone into the system. And this is really robotics that's developed as part of AI. You couldn't have these forms of robotics without having fairly sophisticated systems of AI. Now, what difference will they make to lawyers? Probably, not much. However, they will make a difference to your clients, and they'll make a difference to your clients in terms of how they operate, even in the service sector. And they'll make a difference to you, probably personally, in terms of how you operate and what your expectations are. And, because they'll make a difference to the human experience more broadly, that actually means the clients might expect something a little bit different. Now, some of the robotics that have been around for a little while now are just ridiculous. They're toys. They're nothing serious. However, they will become much more able and adaptable in the near future.

What will really have an impact on lawyers, though, is the development of sophisticated AI workers. So, this is an example of an AI worker, because it's not a real person, it's a made-up person. They will often have an image and be called Bob or June or Jan, or whatever it might be, and they'll actually communicate with you. This is one in a medical capacity, you know, appears up on you. “Hi. My name is Jan, and I've just completed the lab tests.” You know, it's agentic AI. “I'm just letting you know that I've sent the results through to the referring doctor, and also sent an email to the patient, letting them know”, so they've already done all of that work that you might otherwise have to do. Just think about this in the legal context. “There's only one matter that's unusual. The patient has very low Vitamin D levels. I can send an email to the patient with information about this after I talk with the referral document.” None of this is done by a human. This is all just happening. “I'm gathering the other results, and we'll have them with you in the next 10 minutes.” Essentially, the agentic AI is checking in with the medical practitioner just to make sure that they're doing what the expected work would be. And, the medical practitioner doesn't have to consult with the referring doctor, doesn't have to email the patient, doesn't have to check on what might be unusual in the results. Now, apply this in a legal setting and you can see how important agentic AI will be as a support for lawyers. So many lawyers will expect to have somebody who works with them, Jan, Bob, that's actually agentic AI. That is an AI assistant, and we've already got it built in some areas. We already have paralegal AI. We already have accounting AI. Over the next few years, this will become more and more important. And then we're going to be challenged, of course, by the development of a whole range of these crazy systems. And, in fact, we were just talking about a minute ago, the case in Arizona, the manslaughter case, where the daughter of the deceased put together an avatar to appear in court. And, of course, there's been another one where the US courts just laughed at the fellow who had an avatar appear on his behalf. I've actually been making an avatar because I'm giving a lecture in a couple of weeks, and I thought it's so easy now. I'll just like, insert a bit of me as an avatar, and I can have a seat while they present up the front. But this, of course, has implications in the context of evidence.

There are many different law reform approaches going on around the world to deal with this. It also has implications in terms of human-to-human contact and how we deal and work with one another into the future. And of course, most of you will remember the deepfake scam with Arup, given that it

took place in Hong Kong. You're all familiar with it, yeah? So, $25 million; a lot of money, a lot of money, and we'll see more of that with lawyers as well. So, one thing that tech will actually do is generate quite a lot of work for lawyers, because there will be questions. For example, in agentic AI, do you have capacity? You know, have you actually discharged your obligations in a contractual sense? If your agentic AI has dealt with something, if you've delegated the making of an agreement to a form of agentic AI, you know, is that agency or have they acted in an improper way? I mean, there's already been that one case of Moffatt and Air Canada where the chatbot incorrectly advised a person about their entitlements in relation to a reduced airfare. And in Moffatt and Air Canada, Air Canada lost. It was like Air Canada saying, “Oh, it was the chatbot. You know, it's not us.” The court, of course, didn't accept that.

So, then you've got the third level that I talk about, which is on the horizon, if not already here, and that's where you've got neurotech, what's called “the living intelligence”. That's the convergence between human and technology. It will happen when we have more quantum computing, which is just in its infancy at the moment, but again, hold onto your hats for the next five years. That's when we get Web3. That's when we get the growing power of the web giants, if you like, the tech giants. And, they are seriously powerful in the context of the amount of money that they make, the workers that they have, and the influence and power that they exert across the world. It's also where you get much more complex, independent, agentic AI, which responds straight away. So, where does it start? You know, you've got Neuralink. But there's actually better, I think, better forms, much less invasive forms of neurotech. And, of course, there's all the wearable neurotech now, so that you can wear a hat and not be depressed. Or you can have your stress dealt with in a particular way. Or, there was the; did anybody see the headbands that the children in China were wearing for a short time, until parents and teachers were so appalled that they required them to be removed? They showed a green light if a child was focusing, or a red light if they weren't. And, I've often laughed and thought, “Well, how would that be for lawyers?” You know, you're in a courtroom, and the judge says, “Here you go. Can you just put this headband on?” And, you know, how would it go if the judges were actually required to wear the headbands as well? So, there's all sorts of very interesting neurotech that is developing that's quite simplistic, but at the same time, the capacity for the brain-computer interface is fascinating. I actually think Synchron in Australia, which has been around for seven years now, is

slightly more interesting. Because there it's like, you know, when people have a stent put in their heart, what happens is it passes up, instead through a vein, into your brain and actually can read your thoughts. Now that is much less invasive technology than what you're talking about with Neuralink, but of course, there's a few more variations of Neuralink.

This is big, and this will change how we operate. So wearable tech will have an enormous impact, but some of this neurotech will also have a big impact, and Elon Musk predicts brain chips will be in hundreds of millions of people within the next years. So, for those of you who are getting a bit older, like me, there's hope for us. Let alone the notion that by 2050, which was a prediction made a few years back, we're supposed to be able to download our brain and feed it into a collective consciousness or something else. I don't think anybody will want my brain by that age.

Okay, so what does all of this mean for lawyers, these tech developments? One of them is that you've got this court-lawyer-client disconnect. You've got legacy court systems that are limping on, and are often not appropriate for the real challenges of our modern world. And, then, that effectively means that there are threats to courts. And, when there are threats to courts, not that all lawyers litigate, but that means that there are threats to lawyers as well, because it impacts on litigation, and it impacts on how and whether people trust and respect lawyers in the same way. You've got the increasing privatization of justice. And, you see this with some of the tech giants who've set up their own systems offshore, whether it's a form of arbitration or it's a form of something else, and you've got different literacy needs. So, you've got, what I call the digital natives and the old dogs, you know, you can't teach an old dog new tricks. I've actually got a cartoon somewhere about that, which shows that you can. So, it's really the point, though, is that when we're looking at AI literacy, we have to understand that people's requirements will really differ according to not just their age, but according to their interest, capacity, and experience.

So, the very dense slide, and I've got a couple of them, so I'll be quick. These are examples, because we're going into education and AI, not just a legal practice AI piece, about how content can get embedded. And these come from Australia, originally developed through Bond University. I took them on when I was a Dean of the Law School, and I built it a bit more, to say, “Okay, well, what, what sort

of content do you need to add?" Now, of course, you have to have skills assessment. That's going to be relevant. But you need to import some sort of, you know, connectivity, you know, you need to look at Gen AI. You need to look at 3D printers. You need to look at cybercrime. You need to look at, for example, deepfakes. You need to look at targeted surveillance. You need to look at all of the challenges that emerge in our modern technological era. So, we have foundational courses where that's seated, and I think this is a much better approach, my personal view, than having a tech law sort of subject. You actually need to make sure it's embedded in every course that you teach. And then in the intermediate range, you've got another sort of range of courses. They've got the foundational, then you've got the intermediate. You know, here you're sort of drawing on different case law, different questions, and then you've got the capstone courses as well. And they're the ones that really sit towards the end of a program, and ideally are integrated with the clinical program. So, in my own university, the last two years of the law school education involves clinical education as well, in our own pro bono clinic, as well as sending students out to do other work. If they're really going to engage with tech, we need to make sure that it's in our law schools. It not only needs to be in our law schools, but of course, we also need to consider the needs of lawyers who are out there and who are practicing. So, there's very different groups.

And, for lawyer training, and don't you love the way Copilot produces a picture? At least this one doesn't have a gavel in it. I keep saying to Copilot, when I get it to produce an image for me, we don't use gavels in Australia, and then it just hides them somewhere in the image. Oh, this is so frustrating. But lawyer training, you know, big law is already training. They've got their own courses. Sometimes, they've got their own systems. Sometimes, they're very well-developed systems. Sometimes, they have quite elaborate and useful governance arrangements that are in place. The legal publishers, of course, are training people all the time. And, it's no surprise, you know, of course, they want somebody to use Thomson Reuters or LexisNexis, or whatever else. And then you've got training in relation to what tools are actually out there. And that is often focused much more on the generative AI, not much focused on, you know, how easy it is to build, for example, a large language model yourself, because you can do it now. It's actually not that hard. But looking at what tools are out there for lawyers, what ethical principles should apply? When we get to ethical principles, there's pretty much a common set of about ethical principles that apply in relation to AI. Most governments around

the world have adopted something fairly similar. They look at fairness, they look at bias, they look at accuracy, they look at a range of issues with those principles. Then you've got legal use cases, and they're really important for lawyers to understand. You know, what happens when somebody's produced something from ChatGPT? What happens when an expert report has been reliant on ChatGPT? What's happened when, for example, somebody's attempt to use an avatar in court was accepted? In Arizona, the judge was happy to hear from the avatar, which was an avatar of a deceased person, about the victim impact component.

Then, you've got AI workers, and then you've got governance and regulation, and with AI workers, that's probably going to be a big area. Although, realistically, and you know, you could look at science, sci-fi movies, some of the AI workers will really be almost trained by the Gen AI that's already out there; so, that will be working with the AI workers. Again, they're not physical, the robotic AI workers; you know, are they really necessary in law? Probably not. But you know, if there is a need for them, then the agentic AI will be there as a platform to take it on in the same way that basic digitization has been critical for Gen AI development.

You're all looking at me a little bit blankly. Have I gone over time? No, I'm exactly on time. I think. So, I was going to ask, are there any questions? I tried to traverse both of the very interesting panels that are following, and I do agree with the approach, in relation to multidisciplinary governance, it's really critical, and it ought not to be just the judges. That's the hierarchical approach that's operated in the past. So, when I'm looking at courts and governance, I'm thinking, it shouldn't just be judges. You should have an AI expert or two on a committee. You also need to have an ethical expert, and you probably need to have a user on the panel and maybe a lawyer or two who's going to contribute to making decisions about how technology is used in courts in the future.