

REPORT ON PROCEEDINGS BEFORE

**PORTFOLIO COMMITTEE NO. 1 – PREMIER AND
FINANCE**

**INQUIRY INTO ARTIFICIAL INTELLIGENCE (AI) IN NEW SOUTH
WALES**

UNCORRECTED

At Macquarie Room, Parliament House, Sydney on Friday 8 March 2024

The Committee met at 9:15 am

PRESENT

Mr Jeremy Buckingham (Chair)

The Hon. Robert Borsak (Deputy Chair)

Ms Abigail Boyd

The Hon. Stephen Lawrence

The Hon. Jacqui Munro

The Hon. Cameron Murphy

The Hon. Chris Rath

PRESENT VIA VIDEOCONFERENCE

The Hon. Dr Sarah Kaine

The CHAIR: Welcome to the first hearing of the Committee's inquiry into artificial intelligence in New South Wales. Firstly, I would like to acknowledge the Gadigal people of the Eora nation, the traditional custodians of the lands on which we are meeting today. I pay my respects to Elders past and present, and celebrate the diversity of Aboriginal peoples and their ongoing cultures and connections to the lands and waters of New South Wales. I also acknowledge and pay my respects to any Aboriginal and Torres Strait Islander people joining us today.

My name is Jeremy Buckingham. I am the Chair of the Committee. I ask everyone in the room to please turn their mobile phones to silent. Parliamentary privilege applies to witnesses in relation to the evidence they give today. However, it does not apply to what witnesses say outside of their evidence at the hearing. I urge witnesses to be careful about making comments to the media or to others after completing their evidence. In addition, the Legislative Council has adopted rules to provide procedural fairness for inquiry participants. I encourage Committee members and witnesses to be mindful of these procedures.

Dr STEFAN HAJKOWICZ, Chief Research Consultant – Technology Strategy, Policy and Foresight, CSIRO Data61, sworn and examined

Dr PAUL TYLER, Data Privacy Team Leader, CSIRO Data61, sworn and examined

The CHAIR: Welcome. Thank you for making the time to give evidence and for the submission that we've received. I think the secretariat has informed you that you have the opportunity to make a brief introductory statement before we turn to questions, if I could start with Dr Tyler.

PAUL TYLER: I'll just let Stefan.

The CHAIR: Thank you.

STEFAN HAJKOWICZ: We are both scientists at CSIRO in the areas of digital technology and artificial intelligence. Paul has a stronger focus on privacy issues, and I work on digital strategy and policy problems. I wrote the national road map on artificial intelligence and the national ethics framework on artificial intelligence and have been working on strategy and policy issues, advising governments and companies about how to adopt and adapt to a world of artificial intelligence. That is the nature of the work we do. We work at CSIRO, which is about 5,500 staff, with research capability in all fields of science and technology. It is the national science agency. It has a mandate to do science to solve problems for industry and to improve the quality of the lives of Australians, and has been doing that for over a hundred years. Paul and I both work in Data61, which is the digital sciences part of CSIRO.

The CHAIR: Fantastic. Just so you know, we do have a couple of Committee members participating via Webex. That is Ms Munro, and I think later we'll have Dr Kaine. We will now turn to questions. We will start with the Government. Mr Murphy, would you like to kick it off?

The Hon. CAMERON MURPHY: Yes. I might kick off by saying thank you for coming along today and thank you very much for your submission. The questions I had were around the section of your submission where you talk about the interaction of privacy rights and AI. For example, you use as an example in your submission the difficulties in implementing things like a right to be forgotten—in effect, training AI, if it has spent months learning something, to unlearn how to do things. I'm wondering whether you can shed any more light on the processes around that and whether you think there is a way the State could, for example, regulate to ensure the privacy rights like that are protected as AI develops.

STEFAN HAJKOWICZ: I might turn to Paul because that's his area of expertise.

PAUL TYLER: Thanks, Stefan. From a research perspective, this is still an emerging area of research. Attempts to have particularly large AI models forget things invariably have caused issues where the model will look like it has forgotten something and then someone does something tricky with it, and it manages to remember. From a technical point of view—

The Hon. CAMERON MURPHY: Like the plague that just keeps coming back.

PAUL TYLER: Right, yes. You tell it to forget it and, like if you told a person to forget it, they may or may not forget it. Because of the way AI works, being a complex numerical thing, people can't dive into that model to extract it and remove that information. So, yes, it's an emerging problem, and one which has not really got a technical solution at this point.

The Hon. CAMERON MURPHY: One of the other things which isn't really in your submission but I was interested in is: What do you think should happen in terms of something like, for example, criminal or civil responsibility as a consequence of AI? Do you think that programmers should bear responsibility for the program they design, for example? What I mean by that is, if you use something like a self-driving car, if that then causes an accident and there is injury or perhaps a death as a result of that, who bears the responsibility for that? Is it ultimately the programmer that has faulty code or has failed to take something into account? That there's some responsibility for that AI outcome—do you have a view on that?

STEFAN HAJKOWICZ: When we were writing the ethics framework for Australia, there was a principle on accountability which we thought was important and should be put in there. Ultimately, AI can't be held to account for something because it's not a person, but the people who created it could be. There's the question of accountability around the safety and reliability of the system. But there's also potential for the misuse of the system, where someone might take a system which was intended for one thing and then apply it for something else, and the person who was misusing the system therefore becomes accountable. But there is much complexity in the time ahead about how that will actually play out.

In the press at the moment is the Air Canada incident, where a bot on part of Air Canada told customers something incorrect about airfares. The judgement was that it was still Air Canada that was responsible. They tried to blame the bot and say that it wasn't them. I think we'll learn a little bit more about that, but one of the challenges that I think the New South Wales Government would want to solve as it goes down the AI pathway is sorting out those accountability issues and understanding how they apply. We're at the early days of working through that.

The Hon. CAMERON MURPHY: Right. Dr Hajkowicz, in your submission I think you suggest that we should be looking to other jurisdictions that have already done some work in those areas, like the EU, for example. Is that where we should be looking when we try and formulate regulations?

STEFAN HAJKOWICZ: I think we should look with a lot of interest at the EU. The EU artificial intelligence legislation is the first in the world that we can see, and there has been a lot of research and effort to get to that. We don't carbon copy; some colleagues of mine would suggest that it almost could be done, but I think it needs a look at. We can take advantage of all the work they've done in thinking through the issues around artificial intelligence legislation. That's the main one. There's also the European data governance framework. Data governance and AI governance in legislation are very similar things in terms of the risk that's put around them.

The Hon. CAMERON MURPHY: Are there any other areas outside of the EU that you think are worth looking at?

STEFAN HAJKOWICZ: Across the globe, when we did our principles, we scanned to see what all the world was doing. If you look at the Australian Government principles, they are very close to pretty much all of the principles in all of the countries we looked at, including China. The thing is that there are different interpretations of what those principles mean. Writing the principles was pretty easy, but operationalising them and—like your question about accountability—figuring out exactly what that means in a specific context is challenging. That's the next part of the AI journey. So we see across the globe pretty high-level agreement on what the principles of ethical AI are.

We scanned the literature. We wrote algorithms that read it all for us and looked at what was out there. We found high levels of global agreement about what AI ethics is in terms of a set of principles, but how they're operationalised and they're done is a jurisdiction question and a specific case-by-case question that we are now starting to see a lot of work happening in. On one hand, you certainly don't want to let anything unethical happen but you don't want to dampen all of the productivity gains, the innovation and all of the good stuff that can happen with the technology as well.

The Hon. CAMERON MURPHY: On that note, do you see a world, for example, where something like industrial manslaughter laws extend to programmers of AI?

STEFAN HAJKOWICZ: I don't know the answer exactly, other than to say they're still accountable. If you've done anything and AI is involved as part of that, you're still accountable under the law. There is no change to that. I think it is interesting how much attention AI gets in terms of ethics and the legal, but it's not unique compared to any other technology or any other aspect of how we're living. The standards and cultural norms apply to it as they apply to everything else. I guess it does tilt the landscape a little bit in that it creates new capability to do different things and has a big unknown dimension to it.

The Hon. STEPHEN LAWRENCE: Thanks, gentlemen, for your submission. The first thing I wanted to ask about was whether you think that there are particular parts of the economy where small businesses operate, where those small businesses might be particularly under threat from AI, and whether industry policy might need to be adjusted in that respect.

STEFAN HAJKOWICZ: Our analysis suggests that Australia is falling into a pattern of being a downstream user of AI built elsewhere. This is great in that we're getting all these powerful tools at our fingertips that can do wonderful stuff. You can see it in the form of ChatGPT, Microsoft Copilot and Google Gemini. The capability of those tools is impressive, but I think what Uber did to taxis could start to play out in the knowledge economy and, absolutely, your question is on target in terms of wanting to think about things. Generative AI might change how things work. It could expose the knowledge economy, and a huge amount of New South Wales workers are in the knowledge sector. We are seeing really significant productivity uplift associated with these tools. The sort of disruption—if you look at the share price of Shutterstock or Getty Images, for example, as we have seen generative AI able to produce an image from natural language prompts. That has really changed that marketplace. There's been unambiguous drop in those share prices, which is associated with that disruption.

I think it is an interesting one. I think there are questions for Australia about the extent to which our own workers and our own companies might be exposed to capabilities from offshore, which are effectively delivered through an internet connection via highly advanced AI and generative AI capabilities that can do what they're

doing. If you're a fashion photographer, it is now a bit concerning for you because generative AI can create fashion shots anywhere in any clothes really cost-effectively. What does it mean for your business? There's an adaptation imperative. A lot of my work is in the space of adaptation as well as adoption. How do we adapt our companies into this new space?

At high levels I think governments need to be thinking a little around sovereign capability and how the marketplace will play out. In the United Kingdom, for example, there's a lot of attention on competition policy associated with the big generative AI models from Google, Microsoft and also OpenAI, in that they could just take over the whole market and have a monopoly forming. What if they change the prices on these models? What about the competition aspects of it? Those sorts of questions are getting examined.

The Hon. STEPHEN LAWRENCE: I was thinking, for example, of a small graphic design company that might have two or three people working in it. Is that a sector that, if it doesn't adapt so it can use AI to further its business, might be eliminated because other basic functions are now just available through AI?

STEFAN HAJKOWICZ: I think there is a possibility that is a plausible scenario, but it can also adapt. I think that's the challenge that a lot of New South Wales businesses will have—will be understanding what the risks look like and succeeding in making that adaptation. The other scenario is that it figures out how to use all of this technology and then can start serving clients in America or Europe. Just as effectively it can increase its workload because it's able to automate components of the workload. The adaptation issue is a risk but it's also an opportunity.

There's no turning off this AI technology development pathway. It's going to increase; it's going to get better. The speed of adaption—I think of AI strategy in terms of adaption and adoption. That company wants to adopt, work out which AI tools to pick up, then choose the right ones. The toolkit has gotten bigger and way more powerful but it's got to make the right sorts of decisions about which tools and how to apply them. Then it has the adaptation question, which may be partly doing things with AI but it might be doing things differently that are beyond the reach of AI as well.

The Hon. STEPHEN LAWRENCE: In terms of this question of the legal regulation of AI conduct or outcomes, is the difference between things that might be occurring in the society now and things that might occur through advanced AI that you might have outcomes and conduct in the future that is quite disconnected from the original decision to deploy the AI?

STEFAN HAJKOWICZ: Yes. The advent of foundation AI models, such as—what is behind ChatGPT is GPT-4, which is a large language model trained on vast swathes of data. Trillions of records—more than we've ever seen before, and that's what gives it its power. Foundation models like that are getting developed and we estimated there are around 125 such models, not just in the language space but in all sorts of fields. Foundation models have a characteristic that they are capable of being adjusted to wideranging downstream tasks, many of which the inventors did not even envisage. So that's the world that we are moving into. These tools are being developed and then a lot of downstream applications happen from that.

So there is risk and opportunity. There is a risk of misuse but there is also this incredible opportunity around productivity uplift and doing things faster, better and cheaper and finding solutions. We're seeing that now with generative AI. A lot of people this morning in New South Wales will be logging on to their computers and opening up a generative AI product to do their work differently, better and faster. This, I think, is a key ingredient to productivity uplift for the New South Wales economy.

The Hon. STEPHEN LAWRENCE: I'm not sure if you're familiar with the Backpage prosecution in the United States?

STEFAN HAJKOWICZ: No.

The Hon. STEPHEN LAWRENCE: That was a website-based newspaper, I suppose you would call it, or classified service. It transpired that while it was involved in the marketing and selling of quite innocuous and legal things, people were also using it for illegal things. They prosecuted Backpage successfully because the owners of Backpage, at some point, became on notice that it was being used for those things and continued to do it. So it seems to me there are maybe two things in terms of deployment of AI: There's what might be reasonably foreseeable when you deploy it or create it, and then there's what you come on notice of in the course of its use and potentially profit from.

STEFAN HAJKOWICZ: I think when we wrote that accountability principle in the national AI Ethics Framework, that's what—when you know and you have reasonable basis of understanding about the misuse or harm associated with something you're doing, you should stop doing it and fix it. That's what the notion of

accountability is. I understand the creators of some of these models will argue sometimes it's not possible. Then there's an argument about whether or not that's the case.

The Hon. STEPHEN LAWRENCE: I suppose it might be legally different as well when a particular use or program—I'm not quite sure what you'd call it—is disseminated on a free basis, so it's not a product that someone is paying for; it's a technology that someone has invented and then allowed out there.

STEFAN HAJKOWICZ: Yes. Probably one of the core risks, looking at generative AI going forward, is that the guardrails that we're seeing on the main systems that we use—someone breaks the law, forgets about that, doesn't do it. We see some more risky AI hit the—"marketplace" is the wrong word because you're saying it's coming out for free. We see emerge more dangerous AI with no guardrails on it. I think that's one of the risks for 2024-2025: We start to see a ChatGPT but without guardrails.

The Hon. STEPHEN LAWRENCE: Do you see potential risks in the national security space in terms of violent extremism, terrorism and that sort of conduct arising from AI technology?

STEFAN HAJKOWICZ: I think the risk envelope has changed, yes. It has elevated because of the use of generative AI to mimic someone, to trick you, to make it look like you're talking to someone, seeing someone, looking at their language that's written exactly like them but it's not them, and the ability to do that and deceive people en masse or a targeted attack on an individual. So that is something that we have to grapple with. It increases the risk of manipulation from a foreign entity, potentially, to achieve those sorts of outcomes. But we use these tools at our end as well. The race is on for us to be able to use the same tools to protect Australia at the same time. I think that's the reality of the world we're moving into: The risk envelope is changing, the threat surface is changing, different attack modes are becoming possible. I have to say, cybersecurity around AI isn't an area of really strong expertise of mine, other than to say that it's pretty clear that the risks that we're exposed to have changed, the attack surface has changed, the way that they can use these tools has changed, but at the same time these tools are all about how we protect ourselves.

The Hon. STEPHEN LAWRENCE: I'm not sure if this is your area of expertise, but are you familiar with the video that recorded things that occurred at the Opera House on 9 October last year that took a particular chant that became controversial, and then there was a police examination of the video that had been shared on social media and the police reached a particular conclusion about the video?

STEFAN HAJKOWICZ: I don't know, sorry. Paul?

PAUL TYLER: I think I'm vaguely aware of the event, but I don't know about the particular video.

The Hon. STEPHEN LAWRENCE: Part of the context of that was there was a caption or a subtitle, I should say, in the video that turned out not to be correct on the police analysis of it. I'm not sure if that was artificially generated, or automatically generated, or not. I'm interested in how you think as a society we can prepare the community to receive information in real time that may or may not be accurate or misleading in terms of AI deployment, and how we prepare the community to be sceptical at first glance of something and more aware that what is presented might not be necessarily real? Because this seems to be an AI issue in terms of fake pictures.

STEFAN HAJKOWICZ: The responsible AI activities being done by the Federal Government that were announced earlier this year about high-risk and low-risk AI—I think there are components in there about placing a watermark on an image or video which was generated. It's voluntary at the moment, but it could be something worth looking at as a mechanism to help inform the public. Obviously education, awareness and training—we want people to know and understand that what they're seeing now on the intranet could be completely fake.

The ability to just have that radar active is one thing, but then watermarking images or video that is AI generated, especially in significant consequences where it might be something to do with an election or it might be something like that—at the moment it's voluntary. If it's mandatory, at least whoever created that video has done something pretty bad if they haven't put the watermark on it. So that might be part of it. Then tools that are able to look at that material, the generative AI material, and diagnose whether it is real or not—research to improve the quality of those tools is important, but we're struggling at the moment. A lot of what is gen AI created we cannot detect.

The Hon. JACQUI MUNRO: Sorry I can't be there in person, but I really appreciate you coming in person. Thank you for your submission as well. I'm curious about how you see the role of State governments in this space. Obviously there is more of a focus on Federal regulation, but is there a role for State governments to play in regulation or should we actually just be responding to what the Feds are asking us to do?

STEFAN HAJKOWICZ: That's a really good question. As I read the New South Wales assurance framework, that's exactly what went through my mind. Some of the companies I work with in AI are looking at

does this apply to us here or is Victoria going to have another one or do we look at the Federal stuff? I think there will be a coordination issue. You don't want to create a complex regulatory environment for companies who are doing things with AI, which is pretty much all companies. So that's one of the challenges you need to work through. The things the New South Wales Government could be doing is, one, around all the productivity uplift. There's enormous scope to improve the efficiency, the quality, the safety and the quality outcomes from services in New South Wales via the use of AI. In so many spaces it can improve customer service, it can get better outcomes and it can improve decision-making if used well. So there's an enormous adoption challenge. My view is Australia has been a bit focused on the ethics and the risks associated with AI but less so the practical and productive use of the technology to do useful stuff.

Then, secondly, obviously the ethics, the regulation, is all important. Then the challenge you identify—you've got to find the right sort of balance. The basic problem with AI regulation is you do not want to dampen the productive and beneficial use of AI, which is huge. It's almost essential to make sure our economy stays competitive in the future. But you definitely don't want to let anything bad happen with it, so you've got to balance those two things: You've got to regulate and make sure it's all safe and used well but, at the same time, you don't want to stop all the innovation that's going to be happening. And you need a simple—

The Hon. JACQUI MUNRO: Building the mix between the EU approach and the American approach?

STEFAN HAJKOWICZ: That's right. Balancing those things is—

The Hon. JACQUI MUNRO: That's the sweet spot.

STEFAN HAJKOWICZ: That's the sweet spot. It's not easy to deliver on that, but that's the challenge of AI policy.

The Hon. JACQUI MUNRO: I don't want to ask to stray into areas that you don't feel familiar or comfortable with. I noticed in your submission that you didn't respond so much to the questions around democratic impact, but you did mention the quality outcomes from government. It will be useful at some point in this Committee to get to a point where we're a bit more specific in terms of what we're trying to find. I do agree that the use of AI by government departments is going to be very important. One thing that I've been interested in is how Taiwan is using what they call a "prosocial technology" in Polis, trying to essentially crowdsource feedback from the population about policies and culture, really. Do you know much about that and do you think that's a useful way of incorporating AI technology for better government outcomes?

STEFAN HAJKOWICZ: I don't know about that particular technology or how it's being applied, I'm sorry. I will look up with interest and read more about it. Overall, I would suspect, though, that there is a lot of scope to use AI technology to stay in touch with the community, to understand what they're thinking and saying and how to be responsive and deliver the best possible service outcomes to them and the best possible policies that respond to what they're after. AI can read a lot of information out there. In our work, in science and technology, we do a lot of interaction with stakeholders to try and understand how CSIRO technology and science can be most beneficial to the stakeholder community—industry and community and government.

We can use generative AI to gauge the views of people more effectively sometimes than we can by asking people, because we create AI agents that model the beef industry or model another industry and then we start asking those AI agents about what their concerns or priorities might look like, and the quality of the information is pretty good. I think that would be a microcosm of a very much larger package of capabilities the New South Wales Government could develop to ultimately deliver better services to real New South Wales citizens. I think that's the sort of thing we can see. I don't know what the Taiwan case is doing, but I that around the world we're seeing a lot of development to do those sorts of things with AI.

The Hon. JACQUI MUNRO: You mention the strength of our universities in New South Wales as one of our real shining points of difference perhaps nationally and internationally. I'm wondering how we effectively move from the academic into the commercial, the industrial or, I guess, the governance side of utilising that technology. It's good that we've got these universities, but how do we actually make sure we're capturing the skills and talents that are coming out of them?

The CHAIR: Can I just add to that really good question, Ms Munro. You do in your submission make the case around universities for investments. To Ms Munro's point, how do we capture that and what are the investments we need to be making in relation to our universities and our institutions to benefit from AI?

STEFAN HAJKOWICZ: The limited commercialisation of Australia's R and D knowledge and scientific breakthroughs is a 20-year problem—you're right, the exact question. The University of New South Wales did the core development of the solar cells that are used in 80 per cent of the world's solar panels—perhaps more—but the amount of solar panels made in Australia, our solar-panel-making industry, is almost non-existent.

They're almost all imported. We can point to a lot of spaces where this might be occurring, where invention is not turning into commercially valuable products, but there is so much capability and knowledge in Australia.

We examined digital technology hotspots in Australia. We created maps to work out where our Silicon Valley is, where we find a massive concentration of capability. We called it the "Sydney arc". It starts in Redfern, Eveleigh; it goes up through North Sydney. But we found an enormous concentration in capability per square metre of skilled workers and skilled researchers in all things digital and a lot of AI as well. This is a global hotspot. But the translation of a lot of that into commercial products doesn't seem to be happening. If you look at the AI foundation models, they're not being built in Australia. One of my things is that we are falling into the space of being a downstream user. I think there are strong benefit-cost ratios for investing in AI model development in Australia for private and public sector, and we should be hunting more for those as well.

I guess another backdrop to it is that the countries that are doing well on AI and are leading the space put a lot more intensity into research and development overall. Australia would be—I'm going to get the numbers maybe a little bit wrong, but about 1.6 per cent of our GDP goes into R and D in Australia; the OECD average would be closer to 2.4 per cent, I think. Please forgive me for some slight, probably, errors in those numbers, but it's roughly in that ballpark. The US would be a lot more. Israel has a huge amount of its—it's near the top. We don't seem to have that R and D culture built in.

AI becomes useful when you accept experimentation. The companies that I work with that get benefit from AI, they are willing to try it once—it fails. They try it again—it fails. Three, four, and then it starts to work, and then they start to understand it. Over one project it's risky, but a portfolio of projects—you're almost guaranteed to see things get better. Where I see it fail for some companies is where they buy a product off the shelf and plug it in and hope it's going to do what it's supposed to do. It doesn't work the first time, and then they give up. That's not really the AI journey. AI is about experimentation. As you use it in the New South Wales Government, this is what it will look like. It will look like a lot of experimentation, and you'll have to accept a portfolio of a lot of things not working for quite a while, some things starting to work, and then you really start to learn how it works. In another 10 years time, you have transformed the efficiency and quality of New South Wales government services with it.

The Hon. JACQUI MUNRO: Nice. I will ask a question that leads on from that a little bit. Small business and business has been spoken about. It's obviously not going to be possible for many small businesses to go through that iterative experimentation process to improve their output and their labour productivity. Are there ways that the Government can help small business get to the right AI technologies for them? You say, obviously, globally that you have a pretty concentrated group of technology companies that are actually coming up with the technology that many, many people use. How do we make sure that businesses are able to access that technology and know what AI can do for them and not have to take so much risk or investment as, say, a government might?

STEFAN HAJKOWICZ: Your question is on an important matter, because I think productive use of AI by small businesses is key for national productivity uplift, and the whole big challenge with our economy, really, is productivity uplift. My work is increasingly—I am working with the National Artificial Intelligence Centre right now and developing a framework for AI adoption and adaptation for Australian companies, focusing on SMEs. We're currently working with the Melbourne Business School to talk to a lot of SMEs about what the AI adoption challenge really looks like for them, and we are learning a lot.

I think you do actually see micro-experimentation. In my own journey with ChatGPT, I suddenly found out how useful it was when it could generate Python scripts for me. This project about mapping Australia's digital technology hotspots—I didn't know how to get all my data into map form so it was visually really appealing. But within a flash, I was able to ask it the question about writing a Python script, a computer code, to do this, and it did it within a few moments. It cut out about three or four days worth of work for me by giving me the template, all of the libraries, all of the syntax and code correct that I could then start to work with.

Small businesses are doing that at the moment. A lot of small businesses in New South Wales in the knowledge sector are finding similar things just by micro-experimentation. Then there's larger experimentation and more risky things that are done where you're actually training up a machine-learning model with masses of data, where you're building your own capabilities, and then I think government can help by taking on board some of these risks. I have been working in the space of foundation models and policy on foundation models. I think that there is this possibility that we look at GTP-4, the foundation model that supports ChatGPT.

We could build one for health care. We could build a model that reads all of the health records. We've got to do it with privacy and confidentiality at the top of our minds, and we've got to do it really carefully, but it could improve diagnosis. So the next time you're with a GP, it's running behind the scenes to provide an extra layer of safety around the diagnosis. It really matters because a huge number of Australians are misdiagnosed

every year. That would be one area, I think, you could—the health records that we have are valuable. There's a lot of health data. We can use that to train up foundation models.

Weather forecasting would be another space. We've got masses of weather data. The United Kingdom has built a supercomputer and AI facility to tackle meteorological forecasting with a foundation model that, effectively, does it. I think if you look at traffic, I think if you look at so many different spaces—this is where the New South Wales Government could pick up spaces where foundation models that support those small businesses that then can interact with them can really start to do things better. Otherwise, all the foundation models will come from offshore, and we will be making use from those ones. I think, in our work, we feel like the hunt should be on to look where benefit-cost ratios are good for public sector and private sector investment in foundation model construction to support small businesses.

Those are really risky ones. A lot of risk was involved in getting to something like ChatGPT, in that at the outset of a machine-learning problem, you really don't know whether it's going to work. So you need an entity big enough to be able to take on board that risk, and I think that's the other space that the New South Wales Government can help. The sort of work I'm doing is around decision frameworks and tools and information. We're building a website where you can put in what industry sector your company is in, and it will tell you different types of AI projects that could improve your productivity, and it's using generative AI to do this. These are information tools that help, and I think we want to put more of that sort of capability out there.

You want to look at schools as well, and kids coming through school and university. I think AI needs to feature into the curriculum more. We need awareness of machine-learning and what it's doing and how to use it and how to wield these tools powerfully. My view is that coding is getting more important, actually. The experience I had with Python, that's how I can talk to the AI. That's how I know what it's telling me and how I can use what it has given me to do some practical, useful thing. So that would be another space as well. The main one, I think, is starting to build AI ourselves. Where it makes sense, start building AI to support industry and starting to speed that up. Ethics and regulation are important, but that building bit is being underdone.

Ms ABIGAIL BOYD: Good morning to both of you. There's a lot I want to ask, but I want to start with energy usage. Obviously, that is a big factor here. What has happened in terms of research for projecting the type of increase in energy demand that we might see through AI and AI infrastructure? It's not just energy, but also the sort of pinch points in the network and the connectivity aspects. Has there been research done on that?

STEFAN HAJKOWICZ: Is this about the use of AI to improve the management of the electricity grid?

Ms ABIGAIL BOYD: Sorry, no. This is—

STEFAN HAJKOWICZ: Or is it about the energy consumption of AI itself?

Ms ABIGAIL BOYD: Correct, that one.

STEFAN HAJKOWICZ: This is noted as an issue. Training large language models uses heaps of electricity, which, if it's not off renewable, will generate heaps of carbon emissions. I think the community is becoming increasingly aware of this. Really, we are seeing some moves towards renewable energy cloud computing centres, and we should look more at that. I think users of models should increasingly demand that that model at least tells what the carbon emissions are associated with its creation. So there is a pathway by which we can take cloud computing and electricity usage.

Ms ABIGAIL BOYD: Has there been any regulation in other jurisdictions, where they have looked specifically at, for example, trying to ensure that developers maximise the energy efficiency or something around trying to—

STEFAN HAJKOWICZ: I'm not aware of any legislation yet. I think there are some attempts around the world to try to estimate what the carbon emissions look like. Before legislation, you'd want to have a close look at the extent of carbon emissions from this activity compared to all the other activities to determine the import.

Ms ABIGAIL BOYD: My concern is not so much on a comparison of carbon emissions, but more on are we prepared in terms of the capacity of our infrastructure and our networks?

STEFAN HAJKOWICZ: The electricity network?

Ms ABIGAIL BOYD: Yes. Has it been modelled? Are we able to look at that and have some sort of understanding of how we would manage if we did see unprecedented demand?

STEFAN HAJKOWICZ: In demand for electricity? I am not sure that it has been modelled. I haven't seen that. My mind is immediately going to the question of the extent to which artificial intelligence and all of the

data and stuff that goes with it will actually place a load on the grid. That's important. I wonder whether it would be as significant as transportation or heating and cooling costs. The biggest thing by far causing electricity consumption to rise in Australia is air conditioning. I'd be wondering where does it feature next to that before we really have a good look at it.

Ms ABIGAIL BOYD: But it sounds like we haven't modelled it.

STEFAN HAJKOWICZ: No, I don't think we have, to my knowledge. I might have that wrong, but I can't think of anything I've seen which shows me what that is. It wouldn't be nothing; it would be significant. But I just don't know what it would be.

Ms ABIGAIL BOYD: Particularly in industries where we are effectively replacing, I guess, humans doing hours of work with technology.

STEFAN HAJKOWICZ: Robotics.

Ms ABIGAIL BOYD: You can imagine that as AI becomes more and more successful and we have more infrastructure, particularly the huge amount of GPUs and everything else that is being used, that we could see quite a demand.

STEFAN HAJKOWICZ: I think the carbon emissions and electricity consumption associated with AI development needs a good look. I think the community itself is trying to do this, but I'm not aware of any sort of regulation. Are you aware of anything?

PAUL TYLER: No.

Ms ABIGAIL BOYD: In terms of the internet capabilities, the network, has that been modelled and looked at, particularly if we—

STEFAN HAJKOWICZ: The capacity of broadband to handle it?

Ms ABIGAIL BOYD: Yes.

STEFAN HAJKOWICZ: We are always wanting to see that improve. My work has shown that that's not really been a bottleneck for industry yet. When we did the road map we found that AI development was mostly about the skills and capability of humans, but they weren't limited by cloud computing or compute power and they weren't limited by internet connectivity and speed and reliability. That didn't come to us as a barrier. I think that's still the case, at least for the internet connectivity, but the compute power is something we are now starting to question a little bit.

To build these really big models, you need GPUs from—it's NVIDIA or AMD that make them and they are expensive. They are graphics processing units which are no longer used for graphics, but they are now used to handle the matrix algebra of AI, because they handle it very quickly. NVIDIA and AMD are two of the main makers of these, and there's a big queue up. Microsoft signalled to its investors one of its big risks going forward is the inability to access them. That is one of the challenges we face nationally. If we are going to build foundation models, we need these things. At the moment there is not enough of them. Otherwise you can't really do it.

Ms ABIGAIL BOYD: Interesting. I have a final quick question. What is your personal p(doom)?

STEFAN HAJKOWICZ: My personal what?

Ms ABIGAIL BOYD: P(doom)—probability of doom.

STEFAN HAJKOWICZ: If I look at all the things that are going to wipe us out, AI is very low on the list. We've got the capability to do that ourselves. Humans are still way more risky than AI. I think a lot of the stuff that we see in the media about the end of humanity is just not founded in anything that we can really see. There are risks, and they are significant, and we need to manage them. There is also a risk of not using it—there is also a lost opportunity. There is also the risk that the economy won't grow and that people will be limited in careers; that New South Wales companies will be disrupted if we don't. I think, by all means, we need to acknowledge there are big unknowns and keep our eye on what is happening in the space of AI. But I think the end of humanity concerns are pretty low.

Ms ABIGAIL BOYD: Have you got a number? Apparently everyone in AI has a number.

STEFAN HAJKOWICZ: Do they? Mine would be somewhere at the lower end of the scale. I would look to other things before I would look at AI as being the thing that extinguishes us.

The CHAIR: Dr Hajkowicz, before we conclude—I think we could probably talk to you for the rest of the day, really. You were talking about the need for training. Not everyone is going to be able to go out there and

do a computer engineering degree and become proficient in coding. You talked about the curriculum. Is there an opportunity for the State Government in particular to facilitate some of the training we need in a vocational sense, through TAFE? Should we be doing a cert IV in AI, or that type of thing?

STEFAN HAJKOWICZ: The beauty of AI is you don't need software engineering to get a lot of value out of it. It's micro-mini flash courses. You need to know a bit about something. My view is, yes, kids coming through year 12 should be pretty fluent in a programming language like Python, R, or pick another. Python is great because it's just so ubiquitous. Why I think that's important is, yes, these systems can write code for us, but they really get valuable—you can manipulate them and do things with them and use their inputs—if you also speak their language a bit. I know the view that we don't need to learn Python anymore because we can talk English to the computer now and it will do is partly true, but not really. Where it really gets useful is where it gives you something that you can verify, test and then apply in your own work. I think that's one thing that we will do.

If you look at a lot of kids leaving year 12, they probably don't have those skills. I think all of us in the room would benefit from that ability to use code to interact with these tools, so I think lifelong learning, as well—it just improves our ability to do our job if we have that sort of thing embedded into us. I think one of the tricks around it is not trying to get everyone through a software engineering degree, but creating those mini flash courses where there might be something in your job that you need to pick up that lets you wield these powerful AI tools at your fingertips, lets you figure out what that is, and then you take a couple of days off or you work a couple of extra hours a day to get that training and capability and then you can do the task.

The CHAIR: Unfortunately, that ends our session. It has gone very quickly. Once again, thank you very much for taking the time today to appear. Thank you very much to CSIRO for the comprehensive submission. It's very helpful indeed.

(The witnesses withdrew.)

Professor KIMBERLEE WEATHERALL, Chief Investigator and University of Sydney Node Leader of the ARC Centre of Excellence for Automated Decision-Making and Society; Professor in the School of Law, University of Sydney, affirmed and examined

Dr JOSÉ-MIGUEL BELLO Y VILLARINO, Research Fellow ARC Centre of Excellence for Automated Decision-Making and Society; Senior Research Fellow in the School of Law, University of Sydney, affirmed and examined

The CHAIR: Good morning and welcome. Thank you very much for taking the time to attend. Do you have an opening statement?

KIMBERLEE WEATHERALL: I do, actually, if you don't mind.

The CHAIR: You have done your homework. Let's hear it.

KIMBERLEE WEATHERALL: As a little bit of background, I'm a researcher with the ARC Centre of Excellence for Automated Decision-Making and Society, which is a multi-university cross-disciplinary research centre funded by the Australian Research Council. Between us, with Dr José-Miguel here, we have expertise in privacy, intellectual property, risk-based regulation, the use of AI in education and the regulation of AI and automation in government, both from an international and an Australian perspective. I convey apologies on behalf of our director at the centre, Professor Julian Thomas, who wanted to be here.

ADM+S has more than 80 active researchers whose work spans many of the issues set out in your terms of reference. One particular area highlighted in our submission is our work on digital inclusion and the digital divide, which is very relevant to your inquiry. Our submission shows that rural and remote communities in New South Wales do face restrictions in their access to and use of digital technology. We have researchers working on democratic impacts of AI, on environmental impacts of AI.

But another ADM+S project that's relevant to the Committee is a project working with the NSW Ombudsman's office and the New South Wales Government to create a world-first mapping of the use of automated decision-making systems across New South Wales State and local governments. That report has been tabled with the New South Wales Parliament today, and we've just handed around physical copies of the executive report—which is the short version, believe it or not. The project follows on from the NSW Ombudsman's work in their new machinery report, which pointed out the importance of ensuring that when government is using automation in its decision-making it has to be consistent with good administrative practice and the law. That report also highlighted that we didn't know much about how that technology is being used in New South Wales government.

That's what this report seeks to do. It seeks to answer that question. It represents a sustained effort to map ADM—automated decision-making—across State and local governments based on our surveys, a systematised review of official websites, annual reports and procurement records and some in-depth case studies. Alongside this report, the Ombudsman, we understand, will be publishing a compendium of ADM systems building on and supplementing our data. This report builds on leadership already shown by New South Wales—initiatives like the AI Assurance Framework—and puts New South Wales at the forefront globally of efforts to create more transparency around automation and AI in government. We can talk about the findings in more detail, but we do find that use of ADM systems is widespread, varied and increasing across government. There's automation across every portfolio in councils, rural and metropolitan. There's extensive use of structured decision-making and systems that pull together and link data for public servants. A third of the systems that were reported to us are either planned or in pilot, so it's increasing.

We understand this Committee is focused on AI. New South Wales government organisations are also thinking a lot about AI and making plans for it. Simpler forms of automation and data linkage are more widespread, but in an AI-everywhere world the incorporation of AI into some of these systems is sometimes only an upgrade away. I want to point out that we heard again and again in our research that government bodies and public servants want more guidance about when, where and how to automate. There's a significant amount of accumulated knowledge within government but that learning doesn't benefit anyone if it's not shared. For example, there's widespread use of sensors, computer vision and analysis by local councils, but unless local councils know that others are using it and what works and what doesn't, we risk repeating work or exposing ourselves to poor administration. I hope that gives you a bit of a summary of what we can talk about. In addition to the submission, there is the report—open to questions.

The CHAIR: I might start. Professor, with your legal expertise, when it comes to the immediate priorities in terms of reassessing or revising existing legal regimes, could you name two or three areas that we could or should be focusing on as a legislature?

KIMBERLEE WEATHERALL: Based on the work that we've just been doing around government use of AI and ADM, it's clear that there really isn't good legislation that tells government either that it can delegate decisions to be automated or how to do that well. So I would say that is a priority. In terms of other areas of legislation, our submission at page 8 identifies—

The CHAIR: Sorry, Professor, are there other jurisdictions who are doing that well? We heard previously that the EU has got a data framework, and it is legislation, but is it working? The model that they've adopted in the EU—is it working when you compare it to, say, the US? Are there jurisdictions who are doing it well?

KIMBERLEE WEATHERALL: In relation to the EU, it's worth noting that that's still being implemented. Literally the text is new. But if I were going to highlight a jurisdiction you should pay attention to, it's Canada actually, because Canada is a jurisdiction very much like Australia—very common legal background. But for a couple of years now they've had a system, a directive from government in effect, about automated decision-making use in government and that's been in effect for a couple of years. They're now in the process of thinking about expanding that and talking about the private sector and there's a bill that expands that, but they have a system around government. It's in place. It's not as complex as some of the systems I've seen globally, so I would be paying attention to them. Is there anything that you wanted to add?

JOSÉ-MIGUEL BELLO Y VILLARINO: Precisely that. It's a perfect question. Normally there are two models: the models that think that you can regulate AI across the board, across the economy, public and private, or the models that think the good place to start is the public sector because it's the one that has a particular relation with the citizen, particular commitments towards the citizen. The Canadian experience shows that that is possible because the systems are already in place. We are procuring staff. We are barely ever procuring staff. We have terms of reference. So inserting into those terms of reference specifications about the safety of the systems, how trustable it is—it's much easier. It's what Canada did.

They have impact assessments when they procure systems with automation and they classify it in levels of risk and then, depending on the classification of the level of risk, there are more commitments and further impact assessments. As Professor Weatherall noted, they have the experience. They have been doing this for three years. More than hundreds of systems have gone through that process. So they have learnt, they know what works and they will know what doesn't. They reformed it literally last year, October 2023, to adapt it to generative AI, so it can also adapt as the technology develops. That for me is an absolute priority and it's something manageable at the scale of New South Wales that can work upon the existing AI Assurance Framework.

The CHAIR: Is that operating at a Federal level?

KIMBERLEE WEATHERALL: Yes.

The CHAIR: That's Federal legislation, but are they doing that at a State level, a provincial level as well?

JOSÉ-MIGUEL BELLO Y VILLARINO: Quebec tried to adapt it but for the time being it's only at the level of the treasury department, which goes all across the board of the Federal Government but not at the State level. But Quebec has plans to implement it.

Ms ABIGAIL BOYD: If I could ask you a question about the report that you've just handed us, the use of facial recognition technology in the police and also in transport, I believe, in New South Wales government is something that has been occurring for a while. Are you able to share any insights into that? One of the questions I keep asking the police Minister about in estimates is in relation to the bias that's built into a lot of those systems. In the last estimates I asked about whether there had been any bias testing done on the systems that they use within the police, and the answer was no. Can you give me any insight into what you've discovered but also what we should be doing in terms of regulating in relation to bias?

KIMBERLEE WEATHERALL: To respond to the general point, I don't think we have—it's important to realise that the research that we did was a voluntary process with government. We sent surveys to everyone, including, I believe, the police, and we report what they were willing to share with us. So I'm not aware of information from the police that I can share with you. It's important also to distinguish between facial recognition and computer vision and image analysis, because the systems, for example, that Transport for NSW might use to detect mobile phones is more about detecting things, not faces or identification.

Ms ABIGAIL BOYD: This was more about—they were using on stations basically to try to identify disruptive behaviour on train stations. The one that the police use is Cognitech, which I understand has been found to have racial bias in studies in the US. What do we do now? The cat's sort of out of the bag. What do we do in terms of regulation to—

KIMBERLEE WEATHERALL: Well, if you're asking me in my capacity as a legal researcher who works on AI governance and someone who works in the area of privacy, I would say we're well past the time when we should be setting up legal frameworks to govern the use of biometric information. I think our regulation is inadequate on that score and I think models have been produced. For example, I know that the Human Technology Institute has been working specifically on these questions around facial recognition and obviously the Commonwealth Government is also working on questions around privacy and privacy reform. Again, we are well beyond the time when we should be thinking about how we regulate, particularly sensitive, information of that kind. I think it is important for the New South Wales Parliament to be paying attention to the use by government entities such as the police and I would like to see a form of regulation introduced.

I should also note that there's two pieces to the report and to the work that's been done here. We did a research report and that's based on information that was sent to us. I'm also aware that the NSW Ombudsman's office has worked further with various parts of government to update and expand that information and that they have plans to publish their own compendium, which builds on the data that we have produced. I did see in the hearings that you're speaking to them on Monday, I believe, so you'll be able to ask them whether they have further information beyond what we have talked about.

Another thing that comes out of the report that's worth mentioning is that beyond facial recognition there is a lot. We found a lot of references to computer vision and systems that are looking at detecting various things like the number of people in areas or the number of people at the beach, or road defects on the road, including uses of these systems by local councils, as well as at the State government level. One of the things we draw attention to in the report is that we'd like to see more guidance provided generally on these technologies because, leaving to one side—which I don't like to do—the very important question of biometric information and facial recognition, there is also the question of how we're using our cameras and how it is appropriate to use and analyse material from the cameras. I think there are entities out there across New South Wales Government—both levels—crying out for more guidance on those questions. There's a lot of knowledge in the New South Wales Government but it does need to be brought together.

Ms ABIGAIL BOYD: Reading through some of the submissions from other people, it does seem like there's a lot of research that has been either conducted or commissioned by big business. I noticed in one of the submissions a reference to Microsoft's estimates of how much the Australian economy would benefit et cetera. Is there a disproportionate amount of research in this area being conducted on behalf of or with a business-friendly approach? Do we need to be putting more money into public or more independent research in relation to some of the other aspects of AI?

KIMBERLEE WEATHERALL: Speaking as someone who's funded through the Australian Research Council and Australian public funding, I would emphasise that that kind of independent work that we're able to do as a government-funded research institute is really important. The project that we've been able to do with funding and support from the NSW Ombudsman is also really important work. Again, speaking as a university researcher, I'm not going to say that the government shouldn't invest in research. Actually, if you look at the figures—putting to one side ADM+S—the figures on Australian government investment in research aren't that encouraging. Of course I'd like to see more.

The Hon. JACQUI MUNRO: Thank you very much for coming today. I'm wondering about transparency and how we make decision-making processes or frameworks—for example, you mentioned using cameras in lots of different ways. How do we make those policies or decision-making frameworks or processes transparent in an effective way, without flooding people with information, but being clear about what's actually happening and how either their data is being used or just how governments are making decisions or infrastructure is being utilised?

KIMBERLEE WEATHERALL: Yes, interesting question. Again, I'm going to refer you to our report, which I know is a funny thing to say. We have spent the last 12 months trying to get a sense of how automation works across the different levels of New South Wales Government. It has a lot of challenges. I think what is produced out of this process is a really important step towards transparency. It brings together a lot of information in ways that can be analysed not only by us but by other specialist groups who are interested in particular questions around this. We have a number of comments we make about the benefits of transparency about the use of automation for government and for the general population. There are a lot of proposals on the table. A trend globally is towards I guess you'd call them AI or algorithm registers and ongoing publication of information. Again, I'm speaking about government, by government, about how it is using these systems. More generally, I do think—actually, I'm going to leave it there. Did you want to add anything?

JOSÉ-MIGUEL BELLO Y VILLARINO: Yes, probably, because in academic circles I have bit of a countercultural approach to this topic. I apologise if I sound a bit—

The Hon. JACQUI MUNRO: Love to hear that.

JOSÉ-MIGUEL BELLO Y VILLARINO: First of all, let me give you—again, referring a bit back to Canada, they have decided that not all systems require the same level of detail in terms of transparency. For most of the systems that a person would interact with, it's enough if they tell you how it works roughly, like the frequently asked questions. You don't need the inner workings of the system, as a user, to be able to extract something of value for yourself. In other cases, you may have to appeal the decision and you may want more information and that should be also facilitated to you. But in reality what we should be caring about primarily is: Do the systems work properly or not? That's where I am a bit countercultural.

If you cross customs when you are entering Australia and you get sniffed by one of the dogs, you care if the dog is reliable or not. How the dog came up with the conclusion of why they bark or why they didn't bark is less meaningful than the people who are using the dog. What does it mean? Is there any relevance in the way it is approaching it? But in the end you want to detect a substance that should not be coming into Australia and that is why you're using a dog. You don't want to go into the brain of the dog and say, "Why do you bark?" The level of transparency depends on the effects of the system and the use that you will do to that transparency. Is it working properly? That is a question which has nothing to do with transparency, if it is working properly. Do I want to appeal that decision? What do I need to know, because that decision was made against me? That's another question. So transparency works, in my view, in a scale. Asking for transparency in general—it's just adding paperwork and adding data that may not help us use AI better.

The Hon. JACQUI MUNRO: That's very interesting.

KIMBERLEE WEATHERALL: It does go back to the question that we were asked earlier about testing though. You'd want to know that the dogs are tested just like you'd want to know that there's been bias testing of facial recognition systems.

The Hon. JACQUI MUNRO: Yes, fair enough. Learning about what people care about—like, do people actually care about that transparency on a day-to-day basis versus experts being able to go into the data? I'm mindful—I think my internet connection is not amazing so I'll pass on.

The Hon. STEPHEN LAWRENCE: Thank you for the submission from the university. The first question I wanted to ask was in relation to the legal regulation of AI. If you were dealing with a scenario such as an AI application or program or tool being used for a criminal activity, what sort of legal test do you think might be applicable to the regulation of that sort of conduct?

I'm thinking, for example, if you had an AI application that was able to commit a fraud through an automatically generated phone call or something of that nature, with no human being actually doing the physical conduct, what types of legal tests do you think would be applicable to ensure that liability exists somewhere? We don't, obviously, make a person who sells a knife criminal for all the consequences of the use of a knife. So I'm just curious, is there a rationale to extend liability to the use of AI technologies? And, if there is a rationale, what sort of test would you apply? Would it be reasonable foreseeability or intention or knowledge, or something like that, or maybe a strict liability objective test?

KIMBERLEE WEATHERALL: That's a very difficult question to answer in—I was going to say, how long do you have? I'm not just being a law professor there and saying "It depends", but it depends. I think that's because not all outcomes or harms or even criminal provisions would be dealt with the same way. There are also two different kinds of questions that can arise in this context. One is the question about an individual human criminal out there in the world and using a tool—you know, individual liability. Then you have some other questions around the liability of companies that are setting up tools, whether they're using them or whether they're developing them, and what sort of liability we might want to attach to different companies at different points in a complex AI value chain. That in itself is quite a complex question that requires probably a more considered answer than I can give you off the cuff now. There's actually been some really good work, ongoing work but good recent work, thinking about some of these questions. I'd be happy to refer you to some of that literature. I think I would rather put an answer into words and kind of piece out the different kinds of questions that need to be answered, because it's not simple.

The Hon. STEPHEN LAWRENCE: The reason that I asked that was it seems at the moment that there's a lot of offences involving fraud and extortion and so forth, where the primary actors are overseas but the victim is here, for example. I'm wondering whether there's going to be a lot more of that with AI, where the person who's instrumentalising the AI to commit the offence is not within our jurisdiction but the people who've made the technology and potentially are profiting from it are likely to be, at least in a corporate sense, if not here then maybe in a country with comparable regulation and accountability.

JOSÉ-MIGUEL BELLO Y VILLARINO: I can take that one.

KIMBERLEE WEATHERALL: Sure, if you'd like.

JOSÉ-MIGUEL BELLO Y VILLARINO: I'm going to give you a quick answer to that one, because we had a seminar recently with the ambassador for critical technologies of Australia and the director and CEO of IDCARE, who are dealing with these problems precisely, which is stealing the data of Australian citizens but mainly the actors are overseas. Obviously what was quite blunt before, because you could see the scam miles away, now with AI it has got much more sophisticated because it can even speak in a way that sounds right to you as an individual. So the scale is much bigger. One of the things they highlighted is that Australia is starting to use something called the Magnitsky sanctions. Normally used for corruption or violations of human rights, it's considering using it against these foreign actors that you cannot reach with traditional methods of criminal law or things like this. You can ban entry, restrict the movements or coordinate with other countries to be able to seize assets. So there are steps in that direction which are trying to deal with this. But, obviously, in the current state of affairs, it's only going to get worse before it gets better.

The Hon. STEPHEN LAWRENCE: That's very interesting. In terms of ADM, in light of things like the robodebt scandal and the Post Office scandal in the UK that you may be familiar with, is there an argument, do you think, to regulate ADM so that certain types of outcomes can't be automatically generated, in terms of punitive ones, sanction-based ones or ones that impose a penalty? Is that a way of conceptualising the regulation?

JOSÉ-MIGUEL BELLO Y VILLARINO: The answer would be yes; you shouldn't do that unless you're sure it's going to work. I think the presumption against allowing things that are punitive, there is a strong case of saying as a default position, you need more testing for things that are going to be punitive—not a complete ban but putting the threshold so high that you're comfortable that if that ever happens, it's going to be fairly reliable, at least to a level of less mistakes than the best human can do it, something like that, to create that kind of a standard. Not the average human but look at the best—I don't know, I can think of in judgement trials, like court cases, like the same kind of situation. So you set a threshold that is high enough and the people developing the system needs to prove it so you can move into there. For many other things that are not punitive or are not going to have that direct effect, maybe that threshold should be lower and you should favour the automation and the saving of resources and the scales. Again, maybe the threshold for all of the ones is it should be better than what we are doing now because there will be already an improvement. So not a ban but a threshold really high.

KIMBERLEE WEATHERALL: I would also add one nuance to that, which is it's simple and appealing to draw a line between punitive and other things, and the line isn't that simple. If your system is deciding who gets benefits to support living expenses, for example, and you might say, "We're only going to automate the positive decisions," effectively you may be also automating the negative decisions unless you've got an equally accessible, efficient method to deal with the cases where you're denying that benefit, that there's an alternative. It's important not to think we can just draw a line and say the punitive stuff is a really high standard. You have to think about the stakes of the outcome, not just the particular way we framed that outcome.

The Hon. STEPHEN LAWRENCE: In terms of job losses across industries, are you in a position to give us a sense of how widespread the job elimination is going to be in particular sectors, as compared to previous transitions and previous technological shifts? And, as a follow-on from that, is there a way to incrementally adjust what's going on in the education space to anticipate that, or are we going to see a whole lot of people, for instance, in customer service, who might be taking phone calls or doing other tasks that might become automated—are they just going to go off a cliff at some point or is there a way to ameliorate that?

KIMBERLEE WEATHERALL: I'm going to say that's not my area of expertise, nor is it Dr Bello y Villarino's area of expertise. We do have some people who do some research in that space in the centre, and I can refer you to them. I believe Professor Jason Potts, who does some work in that space, is actually on your list for Monday, I noticed. So I'm sorry, but I'd rather base my answers on research that I know.

The Hon. STEPHEN LAWRENCE: Absolutely. Lastly from me before I hand over to Mr Murphy, I noticed in your submission that Australia is going backwards in science and maths compared to OECD averages or outcomes. I'm wondering, what do you think the implications are of that in terms of this transition to more and more AI jobs and so forth?

JOSÉ-MIGUEL BELLO Y VILLARINO: Sorry, I'm not sure that it's in our submission.

KIMBERLEE WEATHERALL: That doesn't sound like our submission, but I could be wrong.

The Hon. STEPHEN LAWRENCE: It was in the University of Sydney's submission.

KIMBERLEE WEATHERALL: I believe you have someone from the University of Sydney coming later before you.

JOSÉ-MIGUEL BELLO Y VILLARINO: This is submission 46.

KIMBERLEE WEATHERALL: This is ADM+S. Apologies.

The Hon. STEPHEN LAWRENCE: No, not at all. Thank you very much.

The Hon. CAMERON MURPHY: I had a quick question about automated decision-making but also the presumption under section 69 of the Evidence Act in relation to business records and the hearsay rule, where, in effect, we treat something that is a business record as though it is true until it's proven otherwise. With deepfakes in AI generating all types of things, from video footage to other types of records, do you think that's something that we need to revisit, where we can no longer have a presumption that because something is said to be CCTV or a recording of something or a business record, we treat it as true? Do we need to go through some other process to verify it first before accepted?

KIMBERLEE WEATHERALL: The short answer, I guess, is that we are all going to have to get wise about how we treat things like video evidence and the like. As Dr Bello y Villarino mentioned earlier, it is going—fakes are getting more convincing. And while there are research efforts strenuously underway to impose things like water marks and tracing of origin, the reality is this is going to become more challenging. I would have to take on notice the particular question about the Evidence Act. But, look, more—

The Hon. CAMERON MURPHY: I mean, really, generally, there's a presumption at the moment that if someone produces a document that's a business record, it's accepted to be true unless there's evidence to show otherwise. Do we need to really—

KIMBERLEE WEATHERALL: But that's the point, isn't it? If there's evidence to show otherwise and those presumptions can be more readily challenged, I would assume, on the basis now that it is possible to fake things more readily, and we will have to get better—

The Hon. CAMERON MURPHY: In light of that, do you think we need maybe to revisit that, where instead of just accepting it as true, there ought to be some verification process first that that document goes through before it's submitted into proceedings?

KIMBERLEE WEATHERALL: I think I would be reluctant, at least on the face of it—and this is an off-the-cuff answer. I would be reluctant to say we have to change the presumption generally because, as we've been pointing out throughout, in many cases it's going to depend on the stakes, right? It's going to depend on the stakes and the context. One of the things about a presumption in the law that can be challenged where it needs to be is that it can be flexible to that. A court can take into account, "Well, how important is this document? Do we need to look at it more deeply?" There's a certain level of flexibility in a presumption that means—as a scholar who works on AI and technology regulation generally, I am in most cases reluctant to try to start redrafting legislation in response to particular technological developments, because that is a loser's game, and it is a loser's game particularly now. It's a focus on the outcomes and whether the principles that we have in legislation can actually be applied in a way that makes sense with the new technology.

It's why, when we think about legal reform in response to AI, there is a question of actually going back and working out if the legislation is working. There's a case, for example, for going back to anti-discrimination law to see if it really works in this new context. There's a case to go back to legislation in New South Wales to see whether we have the appropriate laws in place for delegating to automated systems in appropriate circumstances. But making little tweaks to legislation—for example, setting in place legislation that actually specified how the verification were to occur, I think, would be a mistake because I think that would be too technology-specific and would come out of date very, very quickly.

Ms ABIGAIL BOYD: I've got a final question. I was just wondering if you could elaborate on the explicability issue when it comes to appeal processes with ADM. This is something that you raised in your submission on page 8 about these mechanisms that challenge AI-made admin decisions. Can you talk about explicability in that context?

KIMBERLEE WEATHERALL: Yes, or explainability. Again, it can be a complex question because it's hard to answer in the abstract because how much explanation you need in a context, as José-Miguel was outlining—

Ms ABIGAIL BOYD: Can you give me an example of a type of AI-made admin decision that we'd see from an administrative law perspective?

KIMBERLEE WEATHERALL: I suppose—to go back to an ADM but not AI example, one of the problems over the course of the history of robodebt was the absence of clear explanations to people about how those debts had been calculated and why and what data was being used to calculate those. In some cases, where you are going to want to challenge a decision—I mean, in terms of some of the systems that we are hearing about in the research we've done, where data is being linked from one source and another, you may only need to know

"Well, where did the data come from? What is the data on which this decision was based?" You may not need to know all the details of the algorithm or the technology or who made the technology. You may just need to know "Okay, the data has come from here and here", which may then give you the capacity to then work out if that data is accurate. If the data isn't accurate, then the decision, under administrative law, would have been made on an absence of the right basis, and then you can challenge it. I suppose it's providing the right—a level of information that will enable members of the public to understand the basis for the decision, which will then enable people to then make an appeal.

Ms ABIGAIL BOYD: So you don't just get a notice that says you owe X amount; you get a notice that says, "Based on the following data we have about you"?

KIMBERLEE WEATHERALL: Yes. "Based on a calculation of this data and this data, we suspect that you owe this money."

Ms ABIGAIL BOYD: Then the person can clearly see it and then work out whether they can appeal it based on whether or not it's correct.

KIMBERLEE WEATHERALL: Where the problem might lie.

Ms ABIGAIL BOYD: Yes. It seems quite logical and straightforward.

The CHAIR: Ms Munro, do you have another question?

The Hon. JACQUI MUNRO: Hopefully my internet is a little better. I was curious about, going back to my earlier question, why you each have different views on what transparency looks like or what's valuable to the public. I know that's quite a big question, but are you basing it on international experiences or surveys? How do you come to those different positions?

JOSÉ-MIGUEL BELLO Y VILLARINO: Perhaps I can explain my position, where it's coming from—mainly, a public servant, on loan to academia for a while. When I look at the legislation that has been in place in different jurisdictions—and this case is, practically speaking, Canada—it's how did they make something feasible? How do they make something meaningful that can actually make a difference? So I shifted from—this is what we, as academics, are debating: that this is probably what can make work in the real world in the operational terms, and how do you match both ends, what level of transparency is a reasonable one. That's probably where I developed my thinking in this space.

KIMBERLEE WEATHERALL: I suspect that the difference is more apparent than real, because, actually, if you look at the way we have outlined our thinking in the report, one thing we try to do is work out "Well, how would you make a level of transparency—how would you make it feasible so that the public in New South Wales has some sense of how these systems are being used and how it is affecting their lives?" Actually, we made an effort in the report to think about, "Well, are there ways that it could just be fed into existing reporting processes that already exist in government, where you could just add a few lines that would say, 'These are the systems we're using'?" The benefit of that is you're not building great elaborate edifices that no-one will look at, but you're providing the information which may benefit the public but also then can benefit researchers who can go in and draw attention to areas which may require—which may warrant a little bit more public discussion.

As we say in the executive report—in the report we're releasing today—there are a whole range of uses of these technologies. There's a whole lot of examples that I think are worth some public discussion, whether it's the use of sensing and imaging analysis in local councils or some of the uses of predictive analytics within the New South Wales Government. Raising those and thinking about them is worthwhile. I am a believer in making transparency work for people, and in a way that is feasible for government and doesn't add a massive administrative burden to no end. So I would refer you to the comments that we make in the executive report and the suggestions we make for how that might be done.

JOSÉ-MIGUEL BELLO Y VILLARINO: Concretely, in the annual reports that agencies have to file anyway, and into detail—so inserting there some information about the systems that are in place, that we see as a manageable, easy and low-cost way of creating more transparency.

The Hon. JACQUI MUNRO: Okay, great. I look forward to reading the report. Thank you.

The CHAIR: Thank you, Dr Bello y Villarino and Professor Weatherall. We very much appreciate you appearing today. That concludes this session. Once again, thank you for taking the time to come and inform the Committee, and also for your excellent submission. It is going to assist us greatly.

KIMBERLEE WEATHERALL: If you have any further questions to follow up, please just let us know.

The CHAIR: I am sure—there was some literature that you were referring to.

KIMBERLEE WEATHERALL: Yes, I will let that out.

The CHAIR: If you could forward that to the secretariat, that would be very much appreciated.

(The witnesses withdrew.)

Mr PETER DERBYSHIRE, Director, Policy and International Affairs, Australian Academy of Technological Sciences and Engineering, affirmed and examined

Professor IAN OPPERMANN, ATSE Fellow and Industry Professor at UTS, affirmed and examined

Mr BEN RICE, Head of Policy Advocacy, Tech Council of Australia, affirmed and examined

Ms ERIKA LY, Policy Manager, Tech Council of Australia, affirmed and examined

The CHAIR: Good morning, everyone. Welcome to this Committee's inquiry into AI. Do any of you have opening statements?

PETER DERBYSHIRE: Yes. Good afternoon, members, and thank you for inviting us to appear. We would first like to take a moment to acknowledge the traditional custodians of the land we are meeting on today, the Gadigal people, and reflect on their long history of achievement as Australia's first scientists and technologists. We pay our respects to Elders past and present. Artificial intelligence is not new. Google's maps and search results, TikTok's content algorithm and even your email spam filters are all forms of AI. But what is new is the rapid pace of change in generative AI and automated decision-making systems, and the increased accessibility of these technologies. Today AI can write a report for you, summarise data, diagnose a disease, create a picture, fake a voice and generate realistic videos up to a minute in length. This is a powerful set of tools and we need to make sure they are used responsibly.

Once, we thought about automation on the factory floor. Now, up to 46 per cent of all work activities are expected to be automated. In many cases, this will simply allow workers to be more productive and spend less time on menial tasks. Workers are expected to spend 60 per cent more time using technological skills by 2030, so we need to make sure that the workforce has these skills to embrace this opportunity and ensure the nation's competitiveness. The decision by the New South Wales Government to allow generative AI systems like ChatGPT in public schools from this year is a positive step. So, too, is the New South Wales Government's recently announced trial of a specially built educational AI system. Our students are already exposed to AI, so we must teach them how to use it responsibly.

The New South Wales Government has an opportunity to invest in AI systems that provide benefits to society, consolidating its position as an AI leader nationally: systems like robotic supported surgeries or AIs to assist with decision-making and approvals. But any AI system used by the New South Wales Government or government agencies needs to be accountable, have humans in the loop and be designed with a human rights focus. AI systems are often described as a black box and we cannot always look inside to see how they arrive at their outputs. It is therefore up to the people to monitor the data that it is trained on or the outputs of the system to ensure that these systems do not inherit current societal biases or reinforce disadvantage. This means the public service needs to be trained on how to use AI systems ethically, and the public should have full transparency about how AI is used in decision-making that affects them.

Acting to ethically support AI will see whole new industries and jobs enter the New South Wales economy. New services will pop up to support the finance sector, surgeries will be supported by automated systems and AI-enabled manufacturing will see reductions in waste and the building of new greener production techniques. New South Wales is already the nation's leader in AI. Seizing this opportunity will see it become not just a national leader but an international one.

BEN RICE: Thank you for the opportunity to provide evidence to this Committee's important inquiry. The Tech Council is Australia's peak body for the tech sector. We represent around 160 companies from a diverse cross-section of the tech industry, including companies developing and deploying AI technology. AI, as you've heard today, is one of the most transformative technologies of our time and holds the potential to drive advancements across every part of our economy and society. The potential it holds for Australia and New South Wales is immense. In a recent report, we forecast that generative AI could contribute between \$45 billion and \$115 billion in annual economic value for Australia by 2030. These benefits will come from productivity improvements of adopting AI in existing industries such as health care, retail, manufacturing and professional services. They will also come from the development of AI products and services that will create new jobs and businesses that were not previously possible.

New South Wales has strong foundations to be a leader in AI. The State has an expansive tech workforce, which is the largest in the country, strong AI research capabilities and some of the country's most innovative tech companies and startups. Our capacity to realise this major opportunity will require a clear strategy, integrated policy choices, targeted investment in skills, assets, adoption, and the growth of new companies and industries. AI is already being actively and widely deployed in a range of industries and settings from finance to transport and manufacturing. We're also starting to see the huge potential of AI to improve outcomes in medicine, health

and aged care. By supporting safe and responsible AI use and development, we can empower the people and communities of New South Wales to leverage the opportunities and benefits of AI.

While the potential upside is substantial, we do acknowledge that AI does come with risks that will need to be managed. Appropriate safeguards are needed to ensure that the technology is developed and deployed safely and responsibly. The Tech Council has made a number of recommendations on managing these risks and harms relating to AI. Crucially, we support governments taking a risk-based approach to the regulation of AI, ensuring that oversight of this nascent technology is targeted and appropriate. We also urge governments to build on existing laws and regulatory expertise. We do not need to rewrite the whole rule book for AI or other emerging technologies. Australia's model of core technology-neutral laws, industry-specific laws and standards, and expert regulators has worked well for decades. The TCA recognises and is supportive of existing New South Wales specific initiatives, including the New South Wales AI Ethics Policy, the New South Wales AI Assurance Framework—which is the first of its kind in Australia—guidance on AI definitions that draw from leading standards, generative AI and guidance on cyber.

We also commend the work undertaken by the New South Wales Department of Education on NSW EduChat, the pilot—and we learn from the AI-powered learning tools in classrooms to empower our next generation of learners and the workforce. Enabling trusted and responsible AI innovation requires more than just regulation. We need to continue building our tech talent pipeline and upskilling our workforce; increase investment in AI research, development and commercialisation; provide organisations with the right AI tools and assurance frameworks; and build digital literacy across the community. We are very pleased to see these issues reflected in the terms of reference for this inquiry and we commend the Committee on its work so far in this immensely important area. To conclude, AI offers significant potential benefits for New South Wales. There are risks, but past experience with the introduction of new technology shows that they can be overcome with thoughtful and careful implementation. Thank you, and we'd be very pleased to assist with any questions.

The CHAIR: Professor Oppermann, do you have any introductory comments?

IAN OPPERMAN: Yes. Nothing prepared—I'm just going to speak off the cuff if you don't mind. I was delighted to hear Ben's outlining of some of the initiatives that the New South Wales Government has led over the course of the years. I'm here at the invitation of the Academy of Technological Sciences and Engineering, but I held the position of chief data scientist in New South Wales from 2015 until the end of last year and actually led the development of the New South Wales AI Assurance Framework, versions one and two, and also led the development of the AI policy and strategy. As we were developing these different artefacts in order to understand, to appropriately communicate the uses of AI and to assure the uses of AI, there was a fundamental belief that AI was something that you did: It was a project that you carried out; it's something that you deliberately embedded into a tool, a system or a process. Over the course of the years, we applied the assurance framework to AI as was, prior to the release of generative AI, and looked at AI as an accelerant, an amplifier and something which adapts, and built in place ways of thinking about the uses of AI as a data-driven tool that had these characteristics.

Once generative AI came along, all of a sudden AI was everywhere. It was just built into the tools. You would get a software upgrade and all of a sudden you were using AI. Or you might be using a tool that increasingly became more adaptive or more sophisticated over time and, even if it doesn't meet the definition of AI that a scientist may offer you, is involved in either a part of an administrative decision-making process or some sort of selection or prioritisation. The most recent draft release of the New South Wales AI Assurance Framework was something which was to be accepted and adopted by all States and Territories in Australia. Some of you may have known about the AI summit that we held in Sydney in December last year, where the Commonwealth and various States and Territories reiterated their expectation that they would start to use either the framework as was or the principles underpinning the framework.

What I hope going forward is that the work of New South Wales in continued development of assurance frameworks, continued adoption of international standards, continuing clarification of use of those frameworks is something which is done and adopted within New South Wales but continues to be an effort which is considered nationally to ensure that we have a nationally consistent approach. The benefit, I believe, is that Australia could actually develop the reputation for being a place where AI is used responsibly, AI is understood in terms of its implications, and Australia could really develop the equivalent of a Swiss-watch approach to the development of AI-assured projects and AI-assured software products. I believe the potential is enormous. I believe the potential to do things well to increase productivity, to generate economic wealth and to improve the personalisation of services is almost unbounded. Potentially also so are the risks. The use of that assurance and the use of that responsible AI framework is essential to the future positive development of AI in New South Wales.

The Hon. STEPHEN LAWRENCE: Professor Oppermann, as a lawyer in my previous life, I came across some of the systems where predictive tools are used, where a program or maybe even sometimes just a

checklist is used to analyse data and then make a prediction about behaviour, which is not then used in the form of an automatic decision but used to inform a decision-maker. I'm wondering, in the context of data collection and then AI analysis of that data, what are some of the ethical concerns that arise in relation to, firstly, the collection and aggregation of the data but then its use in the decision-making process?

IAN OPPERMAN: Thank you for the question. The answer is there are many ethical considerations. To build a response around the New South Wales AI Assurance Framework, the way we thought about the use of AI was an algorithmic use of data and you apply that algorithmic output, the data product—an alert, an alarm, an insight, a decision—in a way which has to be appropriate and ethical. Working back from that point, you can think about potential harms, you can think about unintended consequences, you can think about potential discrimination.

If you are mindful of those potential harms or those potential adverse outcomes, you can then work backwards through what's the algorithm doing, what's the data doing, what are the consequences of using data of a particular quality, what are the consequences of including or excluding minority groups, and what are the consequences of including or excluding data of different quality. And then you reassemble that algorithmic use of data from the data part, the algorithm part and that amplification, acceleration, adaptation, potential amplification of harms, and then consider how appropriate is it to use this in either an administrative decision-making process, a selection process, a prioritisation process, or must it only be something which offers an insight for a human being to consider carefully as opposed to an alert or an alarm where a human being is expected to respond potentially in a heightened state of awareness.

So the assurance framework was built around that thinking: Look at what's being done with AI—that amplification, acceleration, adaptation, or more recently, generate, translate, interpolate; think about what the algorithm is doing and what you can do to assure the algorithm; think about what the data and data quality issues are and how the algorithm and data interact with each other; and come up with an assessment of how could I appropriately use this tool so that I don't misunderstand the potential harms or misunderstand the potential consequences.

The Hon. STEPHEN LAWRENCE: Something that has concerned me, and I think it's pretty well known to the law, is the so-called white-coat effect, which is that juries are more likely to take certain forms of scientific evidence as infallible. I was interested to hear some previous evidence about automatic decision-making that talked about using systems when you can be certain that the risk of error is lower than for a human, basically, and you should take the assumption that that's the most accurate human. But I'm concerned that if you have a predictive tool which might more accurately collate the data, that's not the same thing as a human being able to reliably act upon it. How do we deal with this tendency for us to defer to technology and defer to these things when we probably shouldn't, or there is a danger in doing that in an absolute sense?

IAN OPPERMAN: Again, there are elements of how much should you rely upon it. As we were developing the assurance framework we looked at uses of AI in education, in transport, but also in health and in policing. There are some very serious considerations about the over-reliance: Either you ignore it—machine goes ping and you ignore it—or you assume that it's always correct. Partly the issue is having the human in the loop at the right point at the right time with the right amount of time and resources to actually say, "No, I don't accept that; I need a second opinion," and for them to be truly empowered.

You will probably know about the mobile phone detection application that's been running for some time. You drive under a camera at an intersection and there's an algorithm that says, "Driver and passenger. Driver holding something. Driver holding chocolate bar or mobile phone? This looks like a driver holding a mobile phone." In an inappropriately established system, that may lead to an automatic fine. But what happens in New South Wales is it goes to a human being and the human being is empowered to say, "No, that's not," and the human being has enough time, resources, confidence and competence to say no.

Similarly with the New South Wales sepsis prediction—there is an algorithm that monitors patients in an emergency department. The consequences of that algorithm alerting—and it's not a ping; it's an alert where people are expected to respond. The consequences of it getting it wrong are quite significant. If it has a false positive then the consequences are cost, inconvenience and the people in the ward not paying as much attention to other patients. The cost of getting it wrong from a false negative is potentially death. If you look at that compared to how it's done at present, where you have a nurse doing rounds of the ward and taking vital signs on an episodic basis, who will go back to their desk and punch the numbers into a report—the algorithm may pop up and say, "That patient you saw 10 minutes ago—they're about to have a sepsis episode."

If the false positives and false negative rates are shown to be lower than, or at least equal to, the episodic system then, arguably, constant monitoring does add something, but the person who's responding to that alert must be aware of the limitations of the system and must be confident and competent enough to say, "No, I don't accept

that. I want to do something else with it." That's the trick. We talk about "human in the loop" very often, but a human in the loop who's effectively a rubber stamp is not realistically a human in the loop. Providing enough time and enough confidence, training and competence to say, "I will override," or "I will disagree," or "I will do a second evaluation," is the important part overall.

The Hon. STEPHEN LAWRENCE: Thank you, Professor. One question for the Tech Council, if I could. I know that in the past there's been what turned out to be alarmist predictions about job losses as a consequence of technological changes. Would you hazard some sort of view or estimate about how widespread job losses might be in the coming decades as a consequence of AI and maybe offer some insight into what industries might be particularly affected? One that came to mind for me was people answering phone calls and making phone calls and whether, for example, that might end up being automated or not, but I'm sure there are others I haven't thought of.

BEN RICE: Thanks very much for the question. I wouldn't go so far as to give you a number at this point but we are undertaking some research in this area and we'd be happy to provide that to the Committee once we've finalised that. The call centre one is an interesting example because I think what we see is AI being really good at augmenting jobs but not as good at completely replacing jobs. To take the call centre example, the way that some companies will start to use AI, particularly generative AI, in the coming years is through knowledge management and database management of a service provider. For example, a call centre operator might be using generative AI to generate more accurate and faster responses to callers calling up asking questions. That work will be augmented by things like chatbots. On the back end as well that'll be supported by internal systems that use technologies like generative AI to give answers really quickly and really accurately at a speed at which a human operator might not be able to do themselves.

The Hon. CAMERON MURPHY: I've got one question for the Tech Council. In your submission you talk about the important steps that New South Wales has taken but you identify that Australia lacks a model to help coordinate AI regulation and policy. What sort of model do you think should be in place? Do you have a view on that?

BEN RICE: We're currently working at a Federal level with the department of industry and science and the industry and science Minister on this exact question. In our response to that inquiry, what we've recommended is what we've called a hub-and-spoke model that would help to coordinate AI regulation and AI decision-making across government. That needs to be in line with both global standards—we really don't want to see Australia being an outlier compared to similar jurisdictions like the US and the UK—and that will also have a role in coordinating the various State-level work that's underway in AI as well. AI being this ubiquitous technology will certainly impact parts of the economy that are regulated at a Federal level but of course, as the Committee knows, there are a range of State-based issues that will be affected by AI. What we want to see is a coordinated model across the various levels of government.

ERIKA LY: One of the things that we also point out within our submission is that New South Wales as a State has also led. We've mentioned earlier the NSW AI Assurance Framework and ethics policy and we heard this morning from Kimberley and her colleague that we do have one of the States in Australia that lead in AI research. There is a leadership role here for New South Wales also to be able to contribute to some of those Federal processes happening at that level, and to inform some of the work that's going on behind that with the experts that we do have here in New South Wales.

The Hon. CAMERON MURPHY: In your submission where you say we should be looking towards international models, is there any one in particular? Is it Canada? Is it the EU? Is it Chile? Or is it somewhere else that we should be looking to as a basis for that framework, in your view?

BEN RICE: Rather than saying we want to look at the UK model and implement it here, or we want to take Canada's model and implement it here, I think our approach is to take a step back and to say that what we would like to do is look at those jurisdictions and how they're going about the regulation of AI. One of the really positive things about the UK model is its coordination across government—across all of the different regulators and agencies. We think that's a really powerful model to go down. What we really don't want to see is decision-making around AI regulations centralised and coming out of one part of government, whether that be Federal or State. We want to see that coordination model across different levels of government.

ERIKA LY: One of the things, as mentioned before, is Australia is still a fairly small market economy on the global stage. It's really crucial not just to align with international standards but also to be able to harmonise our laws internationally or else we might lose out in terms of our industry and not necessarily give the chance for our ecosystem here in Australia to really flourish. There's still a lot of work to be done in terms of figuring out the best practices and the best approaches that are happening globally. But I think it's important to also be engaged in

those conversations as we start thinking about the technical standards, the approaches that we'll take for that, and also for Australia to have a voice in that standard-setting process in the international fora.

The Hon. CAMERON MURPHY: There is a real risk, isn't there, if we were to go it alone and come up with our own standards, borrowing a bit from here and there, that if it ended up being, say, out of alignment and perhaps substandard compared to the EU, in the future we could be locked out of markets like that because we don't have the right framework in place? Isn't that right?

BEN RICE: I think that's a real risk. As Erika has just said, Australia is a relatively small market in this sense. We completely understand that there is a unique Australian context in which businesses are operating here and in which government regulation will apply here. But, absolutely, we wouldn't want to see Australia be an outlier in terms of the way we go about regulation.

IAN OPPERMANN: Would you mind if I just add something to that? I mentioned the NSW AI Assurance Framework. The adoption of that as a national approach was being done through the Digital and Data Ministers Meeting. It's Minister Dib who represents New South Wales and Minister Gallagher who's the chair. Minister Husic is aware of that activity but is not part of the Digital and Data Ministers Meeting. Driving a nationally consistent approach has been done through that activity so far. There is a point in time right now where there is a risk that it all gets translated slightly differently to different jurisdictions. The assurance framework is based on international standards, or at least was cognisant of the development of the international standards that were being developed, and they're ISO/IEC standards. It's the 42000 series that's the underpinning.

The 42001 AI system management has just recently been adopted by Australia. If we're aware of those standards and we follow them, then we are internationally compatible, and that's a really market-opening opportunity. But those standards are a little bit like, "This is how the engine works. This is how a steering wheel works." What we decide is which side of the road we drive on and which side of the car the steering wheel is on. Those are the sorts of things we need to get nationally consistent in Australia so that we have a common market and interoperability, and all of the underpinning technical standards are still aligned with the international community, so we have a chance to put an Australian-responsible context on top of the hard facts of the standards.

The CHAIR: Before I pass to my colleagues, I have a question on the AI Assurance Framework. If at all, how is the assurance framework embedded in our agencies? Is it on an ad hoc basis or is it being led by Premier and Cabinet? How is it embedded? The second part is: What is the oversight of it? How do we know it's working? What's the feedback mechanism there?

IAN OPPERMANN: It's all changed recently, but the history of it is that the AI Strategy was released in 2020, along with the AI Ethics Policy—two components that foreshadowed the development of an assurance framework. The assurance framework was developed through trial and testing during the course of 2021, working with what was then the AI Advisory Committee, which I led, where we had experts from industry and representatives of government, but experts in technology, in standards, in ethics and in AI more generally. We released the first version of the assurance framework in December 2021. It went through the parliamentary process and was endorsed by Cabinet in 2021. It was made mandatory to apply across government as of March 2022.

The AI Advisory Committee took a step up to become an AI Review Committee and operated all the way through 2022 and 2023. When generative AI came along, we started the process of developing a version two of the assurance framework. The role of the review committee was to raise awareness but also to look at complex cases. Increasingly, it was understood that many, many parts of New South Wales Government were using AI or automation. You would've heard from Professor Weatherall earlier this morning of the audit that was carried out on behalf of the Ombudsman. These were complementary pieces of work. The idea was that increasingly agencies would develop their own internal capability.

The assurance framework is a self-assessment. That self-assessed assurance framework would come to the review committee if it was a large project, \$5 million or more; if it was funded by a particular source of funding, what was known as the Digital Restart Fund; or if that self-assessment revealed high levels of risk. It would come to the review committee, who would iterate and work through mitigations and ways of assuring that the use of that AI would be something that was brought back into the realm of appropriate risk, or structures were built around it to limit its application. That worked with Health, Police, Education and Justice. The NSW Education use of generative AI came through that process; the facial verification came through that process; some early Police projects came through that process; and the sepsis prediction and wound management from Health came through that process.

The CHAIR: Professor, you said "was". Is it still operating?

IAN OPPERMANN: As of the end of 2023, the role of Chief Data Scientist was retired, so I moved on to doing something else. The review committee still exists. It still is established, but I am now no longer sure

whether it's operating. NSW Health has decided that they will do a variation on the theme with the draft of the assurance framework. Health have said that they need something a little bit different. My information is up until the end of December 2023. I hope that that group continues. It was a very, very good group of people. They have some excellent people. The assurance framework was ready to go to be released. It was being shared through the Digital and Data Ministers Meeting and through the Commonwealth States and Territories, and New Zealand. Minister Dib said he would take the final version through to Cabinet again for endorsement in order to essentially ratify that this was the update, this was the version two of the earlier version of the assurance framework.

The Hon. ROBERT BORSAK: Professor, should that review process be normalised as part of, say, what the Auditor-General would do in their rotational audits?

IAN OPPERMANN: It's a good question. If you mean the audit that Professor Weatherall led or the review of individual—

The Hon. ROBERT BORSAK: I'm talking about the review process. You're talking about it coming through, being reviewed. Even if that review process is still running, to my mind it would be more logical to have, since all government organisations are audited in some way by the Auditor-General, this large area of new technology, AI, also develop a protocol for audit, and that it be included as part of an overall audit process that happens regularly and in some detail.

IAN OPPERMANN: It's a very good question. I might answer it by suggesting that the review or the assurance of AI is different to other sorts of audits, based on the nature of the sophistication about how data and algorithms are used. The way the AI Review Committee operated was, in conjunction with the Privacy Commissioner and the Information Commissioner, we would look at projects and come up with recommendations which relate to the use of data and algorithms, as opposed to what was being done. The Privacy Commissioner would then look at those same projects if they came through that same mechanism and look at it from a privacy perspective. Our working assumption is that we're in a moment in time where this is still a specialisation. Over time, it should not be a specialisation; it should be just—

The Hon. ROBERT BORSAK: That even supports what I'm saying.

IAN OPPERMANN: But the argument we would put forward is that it's a moment in time. So it may well be that that becomes a normal process for normal reviews and it's less exceptional over time.

The Hon. JACQUI MUNRO: Thanks so much for coming and for your submissions. I'm curious, from the Tech Council perspective, what your members are telling you about how they feel in a State and Federal environment in Australia and New South Wales. Are they worried about uncertainty in terms of regulation at the moment or are they feeling that there's enough freedom? What's the feedback that you're getting on the use and potential use of AI?

BEN RICE: There are a couple of parts to that question. Our membership is quite broad. We have relatively small Australian scale-up companies all the way through to some of the bigger multinational tech companies. Almost all of them are operating in a global environment. Even the smaller companies are operating in overseas markets, for the most part. What that means is that they are incredibly conscious of the work that is going on overseas in jurisdictions like the EU and the UK. There was a White House executive order from President Biden towards the end of last year. There is a range of work going on internationally in this area, as well as the work that's going on at a Federal and State level here in Australia as well. Businesses generally prefer certainty in a regulatory environment, as is the case but, for the most part, they are quite conscious that this is an emerging area and that governments are all now rapidly upskilling themselves to deal with these issues as they're emerging.

The Hon. JACQUI MUNRO: Are there particular industries from your membership that are highlighting concerns with use of AI? Obviously, they're involved in technology broadly, but you've got lots of different sectors underneath that. Are there any sectors that are really highlighting concerns to you?

BEN RICE: I wouldn't say so much highlighting concerns. What we've found is businesses coming to us and saying, "Government is talking about new regulation. We're already regulated quite heavily." We have some medical technology members who have been using AI in med tech for some time. The TGA has a number of frameworks in place to govern the use of AI in medical technology. It is not simply the case that a company can start up and create a medical device using AI and get it onto the market without a rigorous process that they go through.

That's the feedback that we're hearing—that there a number of these regulatory systems that exist already that our member companies are complying with. The laws that exist now don't simply not apply just because AI is being used in the way that that product is being developed. So, for example, we have competition law and

consumer protection laws that all apply regardless of whether the technology is AI driven or not. The same goes for things like defamation law, things like privacy. There are Federal laws against using carriage services to harass and intimidate people that apply regardless of whether AI is used in that process. So that's the feedback that we're getting.

ERIKA LY: I have nothing further to add to that one.

The Hon. JACQUI MUNRO: Thank you. Professor Oppermann, I'm curious about the removal of the chief data scientist role and, essentially, the hole that that leaves in what has been a pretty nation-leading initiative in highlighting the importance of, obviously, the work that you've done before. I'm wondering if you can expand on what is lost in not having that role.

IAN OPPERMANN: What was abundantly clear is that AI is being used everywhere in government and, ultimately, having a centralised office or having a centralised activity means that it will be overwhelmed. If that was the requirement, everything would come through that office or that central part. So there needs to be a general upskilling of capability within government, and part of the work we were doing last year was actually linking it into what's called ICT assurance, so information communication technology software assurance—again, part of that process of making something extraordinary into something quite ordinary. There was also an effort to connect through to procurement. There was very, very good work done, and I had hoped that the draft of the assurance framework would actually be officially declared, version two, but then the Cabinet process essentially took that a little further.

What has been lost, I think, is having a central expert group, and we really did have quite an extraordinary group of people. We had the former human rights commissioner; we had the chief technology officer of Microsoft Australia and Zealand; we had the head of the standards group, of doing the AI standards; we had distinguished professors in AI; and we had data ethicists. We had an incredible group of people who dedicated quite substantial amounts of time to dig into really, really significant and subtle issues. We also had some leading legal voices around the table. So I think that has been lost. But, ultimately, what has to happen is the capability must be uplifted in all parts of government. The philosophy we had was we will make this extraordinary technology as ordinary as possible; we will remove the need for specialisation. But I think New South Wales has lost that small expert group who could deal with really complex and subtle cases.

The Hon. JACQUI MUNRO: How long do you think it takes before the extraordinary is made ordinary in this context?

IAN OPPERMANN: That's a very good question.

The CHAIR: Apologies for the lighting. We don't know what's going on. I think the building doesn't like your evidence.

BEN RICE: It happens at the Oscars, doesn't it? They start flicking the lights on and off.

IAN OPPERMANN: If I'm speaking too long, you can just tell me.

The Hon. JACQUI MUNRO: The music starts playing!

The Hon. ROBERT BORSAK: The AI is on watch.

IAN OPPERMANN: Indeed. It's such a fast-moving field that I'm not sure we'll ever actually get there. We had come to terms with sophisticated AI, pre generative AI being released. The "amplify, accelerate, adapt" was the mental framework we had for thinking about uses of AI and then decomposing it: If you're concerned about what's being done, take the AI out. If you're still concerned, take the data out. If you're still concerned, it's what you're doing. It's the amplify, accelerate, adapt.

When generative AI came along, we had to rethink the framework, and that's to allow not only those things but the generate, translate, interpolate. If technology keeps advancing, which it will, and if it keeps accelerating, which it's likely to, then it's very, very likely we'll never quite get to ordinary. Part of what we tried to do was ensure that we stayed in a principles-based framework but increasingly adapted to get closer to the technology with things like an assurance framework, as opposed to changing regulation or changing policy, but we would touch the fast-moving technology with faster moving, more responsive things like an assurance framework.

So I think New South Wales is very well set up to normalise existing AI technology. The procurement folks have taken it seriously. They understand where they need to apply the assurance framework. The large ICT projects understand it and where they need to apply it. So we got there. The trick will be subtle cases, and it was the subtleties that we were learning about with facial verification or with some of the health uses of AI. There's

always going to be a moving frontier, but I think New South Wales closed the gap quite substantially, at least with current generations of technology. But technology keeps moving.

The Hon. JACQUI MUNRO: Do you think that there is a useful interaction between politics and our work as representatives and the decisions that are being made at a departmental bureaucratic level—that you've got people who are not elected and not accountable to the public in the same way making decisions that are based in an ethical framework? Where does the role of politicians lie in making those kinds of decisions and communicating that to the public?

IAN OPPERMANN: That is the definition of a career-limiting question. Thank you for asking that.

The Hon. JACQUI MUNRO: My apologies! It depends who's in government.

IAN OPPERMANN: That's, of course, true. Let me offer you just a couple of perspectives on that. I was the adviser to Minister Dominello when he was in the seat and then to Minister Dib at the Digital and Data Ministers Meeting. There were two of us, one doing digital and one doing data. Early on I would constantly talk about standards, and I mean of the sort that Ben was just talking about—international standards around data and digital. My colleagues asked me to stop doing it because it's clearly not relevant to the conversation. But without awareness that those standards exist, it's like building houses without understanding that there are standards for building houses, or building electricity systems without understanding that they exist. So it's important for everyone to know that rules and frameworks and standards do exist and to encourage people to—for someone to develop expertise in that space. So that's important.

Secondly, there's a belief that data—and AI is a use of data; not to dismiss the value of AI, but it's just a use of data in many, many respects. But for people to think that data or use of AI is just like using an Excel spreadsheet or updating an app on a phone—it's trivially simple to do that. But to imagine the same uses of data and AI to prioritise delivery of services, to decide whether someone is having a sepsis episode—without putting the engineering effort, the investment, the time and the architectural effort into it, to assume it's just an Excel spreadsheet, really does disservice and does disservice to the point where it's not useable. You can't trust it and you can't rely on it for anything important other than simple generation of insights. And that is not well understood anywhere within government.

The last quick insight I would offer is that I spent a lot of time working with Minister Dominello. He went through an educational process over the course of the seven or so years that we worked together, but his colleagues never embraced the value of the data and what he was able to do. They liked Service NSW, they liked the vouchers, they liked the response for COVID but never truly took the time to understand what was necessary to make that happen. Without that awareness, we run the risk of diminishing the value and understanding of what you can do with data and making decisions which don't reflect the reality of what we need to be a modern society, which is actually driven by data.

The Hon. JACQUI MUNRO: I have one final quick question. Sorry, Mr Derbyshire, it's not to you, but to the Tech Council about the interaction of your members with government in providing services that relate to AI and whether there is any feedback on government being a good partner in procurement. Does there need to be further clarity, for example, in tenders? Are departments missing where we could be using AI and business could be better offering a service more efficiently or more effectively?

BEN RICE: My really quick response to that would be that this is one of the reasons why we've really recommended strongly a sort of coordination model and a hub and spoke model across government—to make sure there is coordination across agencies who are using and deploying AI. There are a range of different rules that apply to Federal departments and the way that they use AI technologies. We would obviously love to see that harmonised or some sort of structure behind that decision.

Ms ABIGAIL BOYD: Good morning to all of you. I will start with you, Mr Rice. There is a bit of research that is quoted in your submission and also in a couple of others in relation to the value that AI is forecast to add, and I'm unclear as to whether that's to productivity or something else. Can you give me some more detail about how that figure was derived?

BEN RICE: Yes. I will flag this by saying I am not an economist—for my sins, I am a lawyer, and so you will have to bear with me in that regard. But, essentially, we went through and modelled—we give a figure, a range, in the submission where we talk about AI contributing \$45 billion to \$115 billion in economic value.

Ms ABIGAIL BOYD: Is that to GDP?

BEN RICE: I believe so, yes, but I can clarify that if that would be helpful. This is for generative AI in particular. The reason we've done a range model here is to basically forecast three different scenarios for the adoption of AI in the Australian economy: a slow, medium and fast adoption. Obviously, a lot of that will be

dependent on regulatory settings across Federal and State governments, with a fast adoption, we believe, contributing that higher figure of \$115 billion, just in terms of generative AI, that comes from enhanced productivity—things like co-piloting work and augmenting routine tasks that we do every day.

Ms ABIGAIL BOYD: Are you able to provide a copy of that?

BEN RICE: Absolutely.

Ms ABIGAIL BOYD: That would be really good. Was that peer-reviewed at all? Has there been any process—

BEN RICE: I would have to come back to you on the peer review process.

Ms ABIGAIL BOYD: Are you aware of any other studies that have also looked into that, that give any other kind of number? It seems to be the only one that has been mentioned anywhere.

BEN RICE: What I'd say about that is that we really, really welcome increased funding for foundational research in AI development as a technology. Separately, we would also really welcome more of this economic modelling. We can't make proper decisions without having good data in front of us, and we would really like to see that happen, both from private businesses but also from government as well.

Ms ABIGAIL BOYD: Professor Oppermann, because things are moving so fast, there is a concern that a lot of AI advances will be gobbled up by business or gobbled up by—or maybe not even gobbled up but directed towards business uses as opposed to necessarily socially productive uses. Sometimes they will cross over, but not all the time. Should we be doing more to ensure that we are investing in those sorts of socially productive uses—the AlphaFold sort of vibe?

IAN OPPERMANN: AlphaFold—yes is the simple answer. If I could just expand on that slightly, there are many ways of using AI and they range from doing back-office process automation to doing front-office customer-facing, citizen-facing areas. Of course, there are many, many ways that government faces citizens. The back-office parts are really obvious uses of AI, so process automation—some tedious process that someone has typically been running through, or processes which don't quite join up and there is a person who has to take a column from here and stick it over there, and a column from there. There is a lot of that sort of stuff, which could be beneficial to the delivery of government services: Take out cost, take out tedium, free people up to do more productive things.

But then there is the question about how could you use it in citizen-facing ways? There are many. They range from, of course, things like seatbelt detection or whether or not someone is holding a mobile phone, which is really a sensitive use of AI for a particular purpose. Arguably it makes the roads safer—to have acknowledgement and understanding of who is wearing seatbelts and who is not and being able to change the behaviour over time. But then, of course, there are a range of other services that could be just made more joined up. If you're an ageing person in New South Wales, if you tried to find all the policies and regulation that relate to people who are ageing, it's actually impossible to do. But you could use AI to help people navigate through complex systems and by default effectively join them up in a more life journey sort of way. Birth of child, death of a loved one, ageing, ageing with a disability—all of those things could be used to make life easier for people who have to engage with government, and also make government service delivery more personalised.

And then there are things like AlphaFold, which are really cutting-edge uses of—really, it's at the research end of things, but that can help with drug discovery and all sorts of things like that. There's always the customer-facing "this is your next appointment" versus customer-facing "I'm going to predict you get sepsis or not". They're quite different levels of sensitivity and implication. But it's almost every problem that government deals with—almost every service would be improved by use of data, and sometimes that use of data is AI.

Ms ABIGAIL BOYD: In the context of talking about drug production and pharmaceuticals and things, our laws and regulations, and system of rights and ownership have been based on a company or a group of people having put a huge amount of time and effort into something, and then we grant patents and we grant other things over the product of that work. If we're looking at a system where effectively AI is being run to just spit out solution after solution after solution, that sort of disrupts everything. Is there an argument, then, for lowering patent times? Has that been looked at in other jurisdictions?

IAN OPPERMANN: First of all, I should tell you that my wife is a barrister with a background in intellectual property and copyright. I have been told many times I'm wrong about what I'm about to tell you. But I think creativity and formalisation of creativity is in jeopardy simply because an algorithm can do the things that you're talking about. You can take something which looks different enough to be something else, to be innovative, and then you can do a thousand variations on this dimension or a thousand variations on that dimension, and so

on and so forth, and come up with a billion billion variations and then optimise those that it thinks could actually be—excuse me for saying "those that it thinks", but those that seem to be even more creative and inventive.

I think creativity is going to be under serious pressure. Just like we have seen algorithms being able to mimic empathy, I think being able to mimic creativity is something that will be under pressure. Will the laws change? Not in a hurry. Copyright, for example, has been around for a very long time. But people like IP Australia are really concerned about being able to determine novelty based on the rising tide of noise of things that potentially could be created. Beethoven's tenth symphony was completed by an algorithm. Beethoven has been dead a very long time, but the symphony has been completed. I think it's worth all of us thinking seriously about how we would frame economic benefit rights in a modern world.

Ms ABIGAIL BOYD: That's right, and how we socialise that benefit of AI, because there's no doubt that the world can benefit from AI. But, again, if business gobbles it all up and we're not socialising the benefit, then that concerns me. I guess it comes also to the age-old problem of, as we automate processes, it is the businesses who are benefitting rather than the workers.

IAN OPPERMANN: I think government cannot afford to be a passenger in the AI conversation. Government must be a sophisticated user and consumer of AI. It's almost like saying, "Government must learn how to use computers." You can't not understand how to use a computer. You can't be a twenty-first century government without an understanding of AI. It doesn't need to be expert, but it does need to be minimum viable understanding in order to effectively operate the business of government, the delivery of services, match expectations of citizens, and also make sure that government is upskilled enough to be an intelligent consumer of those same services.

The CHAIR: In the five minutes we've got left, could I get a contribution from all of you—I'll start with Mr Derbyshire over there—to the Tech Council's submission, which made the case that there are significant shortages in the tech workforce. Some 330,000, a huge number, are already employed in New South Wales tech industries. Why are there shortages and what can we do to address them? If we could get a contribution from all of you on that issue.

PETER DERBYSHIRE: The shortages—it's a complex problem. I think it boils down in many cases to a combination of having an ageing population, so that we have less people being trained in numerous areas of the workforce that are needed to be trained in, whether it's education, health care or technology, in this particular case. But ATSE's focus has always been also on the importance of making sure the entire workforce has some level of training in this. Some of the discussions we've had today about what's needed in the public service really boils down to not needing huge numbers of experts but making sure that everyone in the public service has an understanding or a basic understanding of how this technology, how AI works and how it is used and how it uses data and how to use data effectively. It's a combination of not just needing the 600,000 people that the Tech Council has called for across Australia but also making sure that we have a baseline understanding for all people coming through into the workforce of how this technology works, and making sure that we are providing schools and teachers the opportunities to be able to effectively teach these technologies.

The CHAIR: Professor, do you have a view?

IAN OPPERMANN: I'll take two points. I absolutely agree with the point about that minimum viable understanding that's necessary. We all use electricity. Very few of us are electrical engineers but we all know how to plug in a light and we also know not to put a fork into a toaster. We have that minimum viable understanding of electricity. We need the same with data and AI. Then second point is that I chair something for CSIRO called Generation STEM, trying to understand why we don't have the pipeline of STEM graduates—science, technology, engineering and maths. If you ask around any room, you'll get many, many different views as to why it's the case. The honest answer is: We don't know. Putting effort into actually understanding why we don't get completion rates higher in STEM and how we ensure that everyone sees that STEM is relevant to their lives and relevant to their future, I think there's a challenge there that we really need to address as a nation.

BEN RICE: I think ensuring that our secondary education systems, both vocational and higher education systems, are empowered to quickly and rapidly upskill and reskill workers currently in the workforce will be hugely critical to this issue. One of the biggest avenues into the tech workforce, we know, will come from reskilling and uptraining. We need to make sure that industry is working really closely hand in hand with education providers and that education providers are reaching out to industry as well.

The CHAIR: Are there any jurisdictions doing a good job—other States, other countries—in this effort?

BEN RICE: I think New South Wales actually is a great example, through the New South Wales Institute of Applied Technology, which is a great example of an institution that is really cognisant of the need to quickly upskill workers with the right skills and competencies that industry needs.

The CHAIR: We very much appreciate all your submissions and your work in this area, especially Professor Oppermann over the years. We appreciate you taking the time to do the work, make the submission and appear today. It's been very valuable to all of us and to the Committee.

(The witnesses withdrew.)

Ms WENDY BLACK, Head of Policy, Business Council of Australia, sworn and examined

Mr CHRIS LOUIE, Director, Digital, Cyber and Future Industries, Business Council of Australia, affirmed and examined

The CHAIR: Thank you very much for appearing today and thank you for your submission. Do you have an opening statement or presentation that you'd like to make at all?

WENDY BLACK: Yes, Chair, I do, just a short one. Thank you for the opportunity to provide evidence to the Committee on artificial intelligence in New South Wales. The Business Council of Australia represents around 130 of Australia's largest companies. The use and deployment of artificial intelligence will be critical to our future prosperity and competitiveness. Estimates vary but by 2030 generative AI could add from \$115 billion a year alone. If you look at AI more broadly, in 2019 it was estimated at almost \$600 billion a year. This is why the business community has been so focused on making sure we get the approach to AI right. Indeed, AI is already widely used by many of the businesses that are members of the Business Council of Australia. It is used in automated vehicles, in Australia's resources sector, to deliver more efficient manufacturing, to detect and protect against scams, and to develop new medicines and medical devices.

Businesses are very focused on how AI can continue to deliver better products and services to the community. But they are also thinking hard on how to do this in a way that does not lead to harms for Australians. This has meant looking closely at areas like their data governance and their privacy policies, among other things. Governments at both a State and Federal level will also have a key role to play in supporting the use of AI. Australia and New South Wales have benefitted from technology-neutral laws of general application. This has meant most issues that may arise from AI are already captured. It has provided the framework in which businesses can innovate and grow. New South Wales has a strong technology sector, as our colleagues at the Tech Council have outlined already, and it must back in these strengths. This is how the State can secure its future prosperity and create secure jobs for workers. As we go forward, new issues may arise but, in the majority of cases, these can be addressed under existing laws. If new laws must be put in place, these must be risk-based, proportionate and focused on outcomes, not solely the technology. A bad outcome will always be a bad outcome regardless of whether AI is involved or not.

Finally, we know AI will change the nature of some jobs. However, this is not a feature unique to artificial intelligence. All jobs continuously change. The key will be making sure that workers are ready for these changes, including through reskilling and retraining. That is why it will be critical for businesses and governments to work together. We know the change is coming; what we need to do is help businesses and workers prepare for this. This means training on how to use AI tools and getting people ready to take advantage of the new opportunities. This will be the best way to ensure that everyone in New South Wales is ready to take advantage of what AI will bring. We are now happy to take questions from the Committee.

The Hon. CHRIS RATH: Thank you so much for your evidence today and for appearing before us. How do you think that the Government and the business community can work together in terms of trying to bring the community with us on AI? Obviously there's a lot of concern out there about everything from job losses to security threats to the fear of the unknown. How do you think we could maybe partner together to try and alleviate some of those very real or perceived concerns?

WENDY BLACK: I think that's a very good point because fear of the unknown is always a concern, whether it's AI or anything else. Generally, people like to know what's going on and I think AI has moved so quickly, particularly in the last 12 months, that people are concerned about where it's heading. That's why you've got this Committee going which is, I think, an excellent role for the Parliament to play—to look at something that is important to the community. From the Business Council perspective, making sure that businesses who are developing some of these areas are engaged and educating the community—we ought to go on this journey together. As I said in the opening statement, working with our workforce is really important—that they come on this journey—and, in fact, so is highlighting where AI is used at the moment and where it may go in the future.

One of the things that does need to be clarified is: What is AI? For many people, AI is used already. It's really just the way that we use data. It's meant great things. It's made jobs better and easier. We've got more information. Separating out some of the narrow AI from what is, say, generative AI, is a really important part. So there's an education about what it is, what it does and explaining those differences from a worker's perspective and seeing what that change will be.

One example from one of our members, Dulux—they've instituted AI into their processes. That's meant that for that workforce that they've got better jobs. They've got higher paying jobs and no-one has lost their job. It's meant that, as a business, they've remained competitive in an area which is really hard for Australia to compete

in—with other countries in the region which may have a lower cost base. It's using AI to show that there is an opportunity for Australia to use that technology to remain competitive, to have higher paying jobs and safer jobs. Jobs have changed continuously. We've got to see where those opportunities are, and I think taking people on that journey with us is really important—and providing that education.

The Hon. CHRIS RATH: And also collaboration on regulations. I noted in your submission in terms of your recommendations you've said, "Any new regulations must be risk-based, and adhere to well-established regulatory principles", and you talk about government, the business community and academia working together on regulation. Obviously we want to get the balance right, where it can't be so light touch that people are scared about AI, but also we could go too far the other way, where we have too much regulation so that you essentially kill off AI performing its true functions and unlocking its real benefits. Could you talk a bit more about that, in terms of where you see the regulatory landscape heading and what you think we should do and what we shouldn't do as the New South Wales Parliament?

WENDY BLACK: I'll make some comments and I might hand over to Chris as well. I think the role that needs to be played by the States—and New South Wales has certainly taken a lead on this and I think that needs to be commended because they've really looked at this in detail and set up the assurance framework—because, really, this is not going to be something that just happens in New South Wales. It's happening globally. What is really good to see is that you're actually getting global action and Australia has signed up to that process as well. So it's something that we have to do together and collaborate on. It would be good if New South Wales worked with its other State colleagues to come together for a joint approach, but importantly with the Commonwealth, and then look at what is happening internationally.

We certainly commend the Commonwealth for the work that it has done so far in setting up an expert advisory group—the BCA is an observer on that—and that is working together in a collaborative way to see what are the ways that we can regulate. An important part that needs to be remembered here is that, in many of these areas, laws already exist. Criminal behaviour and bad behaviour is still regulated under the current criminal offence laws. What we do need to do is say, "Let's look at the laws we currently have." I know there's a major review federally on the Privacy Act at the moment. How does that then fit with AI? Talk to business, the academics, the developers and each level of government and say, "How does that come together? Are there gaps? What do we need to do differently? Is there a broader framework that needs to be done?"

CHRIS LOUIE: Just a couple of very brief additions to that. This is a very fast-moving area and it's an area where everyone is learning as they're going and there really hasn't been time for that accretion of knowledge within parliaments, government departments and policy agencies but also within regulators. What we actually need to see is looking at different ways of trying to bring in some of that knowledge and expertise. A lot of businesses have deep expertise within this area, and really we can't just have regulators going and saying, "We'll just go and make up new regulatory guidance on how our laws apply", and try and shoehorn that into AI. It really needs to be a conversation that is had with the businesses who are either developing it, applying it or are on the receiving end of it to understand how it is actually playing out in reality because the biggest risk would be to have laws or regulations or guidance that doesn't actually fit for how it's happening in the real world. Then you have businesses really not able to work within that context and be able to do the right thing.

The Hon. CAMERON MURPHY: In terms of regulation, are you advocating that existing regulations are appropriate in other areas like privacy and whatnot and there shouldn't be anything in relation to AI? Is that your position, in essence?

WENDY BLACK: No, that's not. What I'm trying to say is that there are a lot of laws already. This goes to the education piece. We have to make sure that people don't fear that there's nothing out there because there are laws already. I think there's a piece to be done with the current regulators. I know that's certainly being done at the Federal level. How do we make sure the current regulators realise how the laws they have now apply? How should they apply? Then identifying where there may be gaps and how they could be strengthened—so it's to give reassurance to the community that there are protections in place already but maybe there are some gaps and how we address those.

The Hon. CAMERON MURPHY: In relation to that, some of the evidence we heard earlier where we've got other jurisdictions, like the EU, Canada, North America, that are putting in place specific AI frameworks—one of the risks may be that if we don't regulate in this area, it may pose issues for businesses in Australia that want to do business in those jurisdictions if we don't have a regulatory framework here that meets their minimums. Do you have a view about that?

WENDY BLACK: Whilst, as I said, there are laws in place already, there certainly is a need to look at, is there an overarching piece that needs to be done—absolutely. We've looked very closely at what the EU has done, the UK, Canada, the US, where we can best be in line with those countries, because AI will be developed

everywhere. We want to be able to both harness the opportunities for the great research that is done in Australia, in New South Wales, that can then go into the international marketplace. But, vice versa, we don't want to cut ourselves off from being able to use those new technologies as well. That's the work that's being done at the moment.

The Hon. CAMERON MURPHY: But, for example, if we adopted a framework that was similar to the one in the EU, that's being developed now, you'd be supportive of that, would you?

WENDY BLACK: I think we've got to work through—and this is the work that the Federal Government is doing at the moment. As I said, we've been very positive about the approach that they've taken, which is a collaborative one of putting an expert group together and engaging with business as well. As they're the ones developing a lot of this and seeing how it's going to be applied, to understand that is really important. What we'd say is it has to be a model that is fit for Australia, but mindful of what the EU is doing, absolutely. I think the EU model suits the EU. It is always best to have legislation that's best for Australia. Then to look from a risk-based approach at what is the outcome that's going to happen from the AI, rather than just being technology-based. Then probably more sector specific, being quite detailed to those particular areas: How is it going to work in health? How should it work in education? I think that's an important approach to take. Did you want to add?

The Hon. CAMERON MURPHY: So what you're really advocating is that we develop our own approach, which may not necessarily align with international frameworks, is that right?

WENDY BLACK: I would say Australia should always do what's in the best interest of Australia and Australian businesses. To do that, I think alignment internationally is actually really important. Going back to my earlier point, you don't want to miss out on the AI technology that's being developed elsewhere, because most of it, despite some great advances—I think of some of the great things in health like Cochlear is doing. That stuff is going out to the world. You want those companies to be able to go out to the world. The majority of it is going to still happen overseas and we want to take advantage of that.

The Hon. CAMERON MURPHY: So you're having a bet each way and saying you want to do some of it but not others, depending on the industry.

CHRIS LOUIE: Just to clarify here as well, and this maybe is reiterating some of the evidence that was given earlier: The experience overseas is still very much in its formative stage. The EU's approach to AI is only just kind of—the ink is, quite honestly, probably still not dry on that. Canada, again, is very similar. It's still very much in the early stages. So it is challenging to say that we should follow too far down these paths, when the implications of what some of these may have are still not yet clear in terms of what it will mean for the business investment environment and for how businesses can offer services that will help Australians and citizens. It's still a very early stage, so it's a bit of a challenging one, I guess, to say.

The Hon. CAMERON MURPHY: But if we just sit back and wait, don't we risk then being left behind and Australian businesses missing out on those huge export markets because we don't have a framework in place? Isn't that a big risk?

WENDY BLACK: If you look at our submission, we're not saying sit back and wait—absolutely not. That's why we think some of the work that New South Wales has already done in this space is very good because it's actually learnt in on this issue. What we're saying is you also need to look at—there are two parts, as our submission says: Look at what we have in place and make sure we use what we already have, but also then look at what is happening internationally. The EU one still has a couple of phases to go before it's finally in place. Canada—there looks to be some positives around that. This is the process that's actually underway federally at the moment. We would say that, ideally, you would be looking to have it done at a national level rather than every State having different regulations. That is where you're going to end up with complexity and that will deter investment in AI and miss out on some of those economic opportunities. We absolutely think we need to be doing it. That work is underway at the moment. The approach we would say is take a risk-based approach to the regulation to be put in place, and that's where we're coming from.

The CHAIR: Before I pass to my colleague Ms Munro, I'm going to ask the question: The assertion—and I accept it—by many is that AI will lead to a massive uplift in productivity. We get the settings right, we implement and we have the workforce. What's the Business Council's view on what we do with that productivity dividend? How do we make sure that business benefits from that productivity, with efficiency and savings and the rest, but that the workforce and society does too? One of our submissions said we should move to a four-day working week.

The Hon. CAMERON MURPHY: One-day.

The CHAIR: A one-day working week—you're already on that, Mr Murphy. With other technologies that have come, like email, which was supposed to make things easier but now just means that our work follows us home—how do we make sure that there's an equitable division of the productivity dividend from this AI?

WENDY BLACK: I would probably refer you to the NSW Productivity Commission submission, which really highlights that if you increase productivity, everyone does benefit. I don't have off the top of my head what their figures were, but they were quite substantive ones about what the benefit would be per household across New South Wales from using AI. The points you raised, though, about how do we all benefit in making sure we get things like work balance right—emails: Yes, my inbox, there are always too many in there for what I would like. But what has it enabled us to do? We can work from home. That has been quite a positive, the technology that we can now have. There are benefits to that.

It doesn't mean we've always got those balances right and we have to work through those things, but it does give us greater flexibility. The first point is productivity, from an economics perspective, will be shared across the board anyway, but then at the workplace level, as I said—I use that Dulux example again—those people have got better jobs. Technology has generally meant that our jobs are safer than they were before, where you have a better working environment than we had before. It doesn't mean there aren't going to be challenges and things that we have to work through, but overall most of these changes have been to people's benefit.

The Hon. JACQUI MUNRO: I'm curious about whether your members are employing dedicated AI experts in their businesses or whether they're preferring to upskill the existing workforce that they have.

WENDY BLACK: I think there's both. It depends on the type of business. You would probably be surprised that across businesses that are not seen as tech businesses, they are absolutely using AI, more so probably in the narrow definition of, "How do we use data better to service customers to get better outcomes for customers?" and therefore lift their business. But the upskilling is a really important issue more broadly, and I know there's already been discussion about this. The skills gap that we have in Australia is a serious one, and we need to do a better job in that regard. That goes right through our education system. I know one of the witnesses earlier said people should be able to do code when they leave year 12—absolutely.

The Hon. JACQUI MUNRO: Is it something that your members raise with you, that they're worried that they can't get enough staff equipped with these kinds of skills to service their businesses?

WENDY BLACK: Absolutely. It is one of the challenges, that we have a skills gap here. That is something that we have done a lot of work, particularly at the Federal level, on: How do we upskill? How do we get things like micro-credentials? Using AI actually can be a much simpler way of upskilling people. You do not have to be an engineer to use it. Therefore, in fact, I think—if done right—it has got huge opportunities for people in the way that they may change their work to be better jobs without actually having to go off and get a degree. I think there are some real positives there. Do you want to add anything, Chris?

CHRIS LOUIE: Again, a very small addition to that. I think there's also very much a focus on how to make sure that existing workforce are able to use AI within their existing jobs to enhance their existing jobs. So there is that real effort to recruit in AI experts and trying to find the best people both within relevant jurisdictions but also bring them in internationally. But then for people who are working in jobs that aren't notionally AI jobs but are going to be affected by AI, the questions are then how to help these people get this; or if there are people who are, for example, coming back from parental leave, how do we—if they're coming back in, are they the right kind of people to be skilled into AI jobs, if I can put it that way, or things like cyber, which is another area of huge workforce gap and there's a real need.

The Hon. JACQUI MUNRO: Do you think that your members are holding off on investing in AI technologies in their own businesses for fear of regulatory uncertainty? Or are people just forging ahead and hoping for the best or just working with what they've got at the moment, and if they have to change their systems in the future because some additional regulatory environment occurs, they just deal with it then?

WENDY BLACK: Certainly regulatory stability and certainty is key to getting investment decisions. This is why we do think that it is important that the Government, at all levels, are leaning into this issue and, as I said, looking both at what laws exist now and what we may need to do so that there is stability around that, because that will increase their investments. Are they doing it now? Yes, they are, because they can see what the opportunities are to lift productivity. I use the Dulux example and the value that that has been for that workforce. But that's across the board. So in the resources sector, things like driverless vehicles, why have those things been put in place? To lift safety so it's actually safer for workers on those big sites. I think there are many of them and it's not where you would think. In a tech company, it's about how they apply it into their business.

The Hon. JACQUI MUNRO: Do business owners feel that there is stability at the moment or is there a sense of instability?

WENDY BLACK: I think they're still looking for more stability and direction. I think that's why the work that is being done federally is really important, that the Government has been consulting along that way. I think everyone realises this is a fast-moving area. It's very broad. Just look at the terms of reference that your committee has, which shows how broad and wide the considerations have to be around AI. It's not settled globally either. Countries are putting in new legislation, but they're coming up with slightly different models as well and they're only just being implemented. It is moving. More stability is great, but this is an area of significant change.

Ms ABIGAIL BOYD: As automation is used by businesses to increase productivity and profitability, should we be regulating for there to be an automatic uplift in wages and conditions?

WENDY BLACK: As I said, if you just look at the NSW Productivity Commission report, productivity flows through benefits to everyone.

Ms ABIGAIL BOYD: I guess the question is: Will businesses do this if we don't regulate it?

WENDY BLACK: As the benefits flow through?

Ms ABIGAIL BOYD: Yes.

WENDY BLACK: Absolutely. The way to actually lift wages in a business is to be a productive, profitable business. That's the only way that that can happen. I'll use the Dulux example. They automated the process and those workers are now being paid more because their business is more productive. They've become more skilled. They've all been upskilled and that is an absolute example of where that actually has occurred.

Ms ABIGAIL BOYD: But then we haven't seen it in, for example, supermarkets and other places in Australia that have been increasingly automated. We have seen those wages stay pretty flat. Do we need to regulate to ensure that we actually are passing on those benefits?

WENDY BLACK: I think the latest data shows that there has been a significant increase in wages.

Ms ABIGAIL BOYD: Just recently, yes.

WENDY BLACK: But there is still more to be done. Australia is lagging on labour productivity. In fact, this is one of the great challenges that we have at the moment. If we can use this to kick-start our productivity agenda in Australia, that will absolutely flow through.

Ms ABIGAIL BOYD: I don't have the data here with me, but there is plenty of evidence to show that over the past 20 years we've had an increasing amount of productivity gains going to capital and less going to labour. I'm asking the Business Council if they think that we should be regulating for higher wages, and I think I know what the answer is.

WENDY BLACK: I would be interested in what data you have, because I think that's not necessarily the data that we would agree with. Our view is that if you have a profitable business that lifts productivity, that's when wages will absolutely increase. That's the way a good business will operate.

The CHAIR: The time for questions and answers has concluded. Once again, thank you both for coming today, and thank you for your interest and for your submission to our inquiry. It has been very informative. Once again, we very much appreciate it. We will reconvene at 1.45 p.m.

(The witnesses withdrew.)

(Luncheon adjournment)

Mr ASHLEY COOPER, Policy Director – Agricultural Industries, NSW Farmers, sworn and examined

Mr ADRIAN ROLES, Executive Councillor, NSW Farmers, sworn and examined

Ms LOUISE McGRATH, Head of Industry Development and Policy, Australian Industry Group, before the Committee via videoconference, affirmed and examined

Mr RON CURRY, Chief Executive Officer, Interactive Games and Entertainment Association, before the Committee via videoconference, affirmed and examined

Mr CHARLES HOANG, Director of Public Policy and Government Relations, Interactive Games and Entertainment Association, before the Committee via videoconference, sworn and examined

The CHAIR: Welcome to our next witnesses, both online and person. Some of our previous witnesses have given a brief introductory statement. Do the NSW Farmers have an opening statement to provide?

ADRIAN ROLES: Yes. Firstly, thank you for the opportunity to appear in this hearing today and to provide the evidence. As stated previously, I am representing NSW Farmers, and I am a farmer in southern New South Wales as well as a member of the Executive Council of NSW Farmers. The ag tech sector—that includes artificial intelligence—is an area where I can provide some expertise, and I sometimes consult on this matter, to many businesses as well, across the agricultural industry. I'm happy to provide a list, if you require that information, following the inquiry.

As a State farming organisation, NSW Farmers represents the interests of its farmer members in New South Wales across all agricultural commodities, and advocates for a profitable and sustainable New South Wales farming sector. The integration of AI into the New South Wales agricultural industry is a developing area of ag technology that the NSW Farmers Association is very excited about, and we're also very optimistic about. AI presents significant opportunities for the agricultural sector, but we need to get the regulation and privacy around data right. If we get this right, it has the opportunity to benefit everyone in New South Wales.

Agriculture in New South Wales is already beginning to utilise technology such as blockchain, artificial intelligence, big data and Internet of Things to help increase our productivity. Examples of these applications extend to precision agriculture, crop monitoring, predictive analytics, automation of machinery and robotics, crop and livestock management, supply chain optimisation, farm management software, weed and pest control risk management, and weather prediction, just to name a few. The increased uptake of these technologies in New South Wales will also bring opportunities for the regions and the communities that are required to support these new agricultural industries.

As farming continues to use more digital technology, new skills and new jobs will be required to support farming systems and equipment. To realise the benefit of these innovations—for example, productivity gains or food security and positive environmental outcomes—we need to ensure that the farm data is protected. This is absolutely critical because the data generated by a farm or related to its operations belongs to the farmers. We need to consider farm data as a farmer's commodity and classify it appropriately in the future. This is important because it will enable the adoption of advancements in technology without stifling these innovations. The regulatory environment and legal frameworks that govern this space need to be fit for purpose now and into the future.

We've made a number of recommendations and comments in the NSW Farmers written submission around the challenges and opportunities for AI and ag tech more broadly for the producers in New South Wales. These relate to the importance of the following considerations in this context: innovation and productivity, connectivity and infrastructure, labour and skills, digital literacy and training, regulations and data privacy and security. In particular, I would like to highlight the importance of the appropriate checks and balances being in place concerning the utilisation of AI in agriculture, coupled with the necessary safeguards for agricultural data. It is imperative that farmers retain control over the data produced on their farms and in their businesses.

Agricultural data needs special consideration because, unlike data from other sources such as social media, a farmer incurs most of the cost at every stage, from purchasing the hardware and equipment that generates the data to paying for the software and all the analytic services that go behind it to utilise that data set, which will include AI. The other thing is the data is actually a product of the farmer's intellectual product and knowledge that they've used to grow that crop or produce that animal. Unfortunately, due to unfair contracts in the form of end user licence agreements, farmers currently lack control over their data. This lack of control leads to a diminished trust and hesitancy amongst New South Wales farmers to embrace and utilise new technologies essential for the integration of AI into the New South Wales agricultural sector. Thank you.

The CHAIR: Thank you very much for that. Ms McGrath, do you have an opening statement that you would like to make?

LOUISE McGRATH: Yes, I do. Thank you. The Australian Industry Group, or AI Group, is the peak national employer organisation representing traditional, innovative and emerging industry sectors. We've been acting on behalf of businesses across Australia for over 150 years. AI Group is genuinely representative of Australian industry. Together with partner organisations, we represent the interests of more than 60,000 businesses employing more than one million staff. Our members are large and small businesses in sectors including manufacturing, construction, engineering, transport logistics, labour hire, mining services, waste services, the defence industry, retail, aged care, civil airlines and ICT.

In our conversations with our members around artificial intelligence, or AI, it's clear that AI will improve productivity, unlock human capital and lift our international competitiveness. It will bolster resilience in our supply chains, providing cushioning in the face of potential geostrategic shocks. Where traditional supply chains plan and react to disruptions, digitalised supply chains predict and prescribe actions to take. Our members are making strong investments in both technology and staff training this year, but it's important to remember that businesses don't actually buy AI. Instead, they buy a solution to their problem that simply happens, potentially, to have AI in it alongside other digital technologies and hardware.

This year, 41 per cent of our members predicted that they would be investing in staff development and training. That won't all be related to AI. In fact, when it comes to technology investments, members are telling us that cybersecurity-related investments remain the highest priority. Skill needs in advanced manufacturing and infrastructure are broad and deep. The sectors need a mix of technical, generic and leadership skills. However, the role of AI will increase the importance of digital within the skills mix. Skills in the digital space will be essential for not just specialist roles but the entire industrial workforce. Beyond business and leadership skills, the successful implementation of AI also requires two distinct but interrelated skills bases. One is the technical skills involved in the design, implementation and integration of digital industry technologies, and the second is operational skills required in workforces that make use of the tech, which will enable its safe and efficient utilisation.

A good analogy is our road system. It requires us to train both engineers, auto mechanics and car drivers, and so, too, the adoption of AI requires us to train the builders of AI and the users of robotics and AI. Both are equally important but builders often get prioritised in policy while the users receive less attention. Let's not forget about all that underutilised data, as has already been alluded to. We need to upskill all kinds of workers to understand what the data is telling them and how to use it to predict and prepare for the future. But fundamental to maximising take-up of AI will be lifting the capability of leaders and managers to develop and execute new business strategies that ensure Australian businesses are equal to, if not exceeding, their international peers.

The CHAIR: Mr Curry or Mr Hoang, do you have a brief introductory statement to make?

CHARLES HOANG: Yes, we do. I'll cover for Ron, because Mr Curry's voice is not that great, but I'm sure he's happy to chime in as required throughout this session. Thank you to the Committee for providing IGEA, or the Interactive Games and Entertainment Association, with the opportunity to submit to this inquiry and inviting us to appear at this hearing, especially the Committee's flexibility today to accommodate us at short notice to attend virtually. IGEA is the industry association representing and advocating for the video games industry in Australia. Our members include developers, publishers and distributors of video games, as well as the makers of the most popular gaming platforms, consoles and devices. The video games industry in Australia provides entertainment to players and families, and makes significant contributions to the economy. Our research shows that four in five Australians play video games. Further, video games are worth around \$4.2 billion annually in Australia while Australian-made games brought in \$346 million in largely export revenue last year.

Although the industry has historically been overlooked, it's one full of potential and innovation. Video games represent one of the earlier cases of AI and the industry continues to develop and harness AI for safety, entertainment and training. We consider AI use in video games is specific and low risk. As such, AI in video games is not a context that we believe would benefit from or require additional regulatory oversight. If specific AI regulation is to be established, we recommend it should be industry-focused and evidence-based before considering regulation. The video games industry is critical to innovation, including embracing the opportunities presented by AI. Like many other sectors, the video games industry uses AI to benefit its workforce and workflows. Further, AI and video games can combine to provide an innovative opportunity in service delivery in government settings through serious games.

With State-based support, industry-specific regulatory considerations and the uptake of serious games, we believe New South Wales can fully utilise the State's video games industry to become world-class leaders and innovators in AI. Beyond New South Wales, ensuring national and international coherence, including with the

Federal Government's work on AI, would also benefit the industry. AI is certainly a hot topic both in Australia and around the world, especially around generative AI, which touches upon different areas of policy and regulatory reform. In closing my remarks, our industry looks forward to the opportunities that AI presents as a potential enabler to drive industry growth, especially in our video games sector, which has been the case over the last several decades. As with any emerging technology, there are also challenges. We continue to be engaged in multiple areas of policy reform and working with governments and other stakeholders. Thank you again to the Committee for the opportunity to attend this hearing and we welcome any questions you may have.

The Hon. CHRIS RATH: To NSW Farmers, I want to ask about foreign investment. You mentioned it in your submission at recommendation 2. How important is it that we get foreign investment into Australia in terms of making our agricultural sector more competitive and upskilling in terms of having new technology in the agricultural market? I assume, if it was only relying on domestic investment, then we wouldn't have a particularly competitive or technologically advanced agricultural sector and certainly wouldn't have unleashed the benefits of AI without it.

ADRIAN ROLES: Yes, you are correct. When it comes to development of technology to the agricultural sector, it does require generally large sums of money, which the Australian agricultural sector is not big enough to foster by itself. Quite often a lot of the technology—for example, the technology we're using is embedded in machinery or in sensors that have all been designed and based overseas. We do need to have incentives to ensure that New South Wales farmers have access to these new types of technologies to ensure that we can increase productivity but also to make sure we can keep competitive advantages on the global scale as well. New South Wales has a role in being able to foster adoption. To me that has to come from the coalface, so behind farm gate. It's about actually encouraging farmers or providing the conditions that encourage farmers to adopt these new technologies. If we have the market demands coming from behind the farm gate, those companies will come to Australia and New South Wales.

The Hon. CHRIS RATH: What about skills as well? We hear a lot about the skills shortage that we have in Australia. Obviously there is a role for skilled migration coming in to plug some of those gaps. We just probably wouldn't have the right level of expertise here in terms of AI domestically. But then there's also a piece in terms of trying to train up our existing workforce; people already in the sector. I was wondering if you could expand a bit more on whether you think we're fit for purpose in terms of our skills base in AI in the agricultural sector.

ADRIAN ROLES: Once again, thank you for the question. There is definitely a skills shortage in the agricultural sector of New South Wales, as there is across many industries. Many farmers are now looking to automation to help solve that. In the horticultural industry, they're now using robotics in packing sheds. There are automatic tractors. The dairy industry has been using automation for a long time. Yes, that is one of the ways these technologies will help with it. To answer the other parts of your question, NSW Farmers actually sees this as a great opportunity to actually take unskilled or semiskilled labour and take the opportunity to upskill them, so we can actually increase the efficiency and the depth in the agricultural industry in that sector.

To the third part of your question, I also see this as a very good opportunity for the agricultural sector or NSW Farmers to work with the New South Wales Government to encourage people who may have never considered agriculture as a career path before. We are going to require data scientists, data engineers. We need to be getting people who have considered other industries before agriculture to consider us, because there is a very good career pathway moving forward for those types of skillsets. The other thing, too, is the people that have been coming into the industry and space get very excited about it, because agriculture is one of the ways that we're going to address some of the major issues facing the world as a whole, not just New South Wales, with food security and environmental things. If we were smart about this, this is a real good opportunity for everyone to benefit in the agriculture sector from the adoption of these new technologies. But, once again, it's only if we put everything in place to allow it to happen.

The Hon. CHRIS RATH: In terms of AI and tech in the agricultural sector worldwide, I think Australia is doing fairly well but where do you look to in awe? Where do you look to around the world and think, "They're really doing a great job"? I know Israel has got a lot of great stuff in terms of biotech, for instance, but where do you look to and think, "This jurisdiction in terms of using AI in agriculture are really smashing it"?

ADRIAN ROLES: To break that question down, it's a bit like being teenagers. Everyone thinks that everyone else is out doing exciting things and all the rest of it, but really no-one is out there doing it yet. Everyone thinks that everyone knows how to do that exciting thing, but no-one is doing it yet. That's a bit like where the tech sector is in ag. There are certain sectors—to use the Israeli example in biotechnology, some of the stuff they're doing with some of the technology based around genomes and mitochondrial splicing is very exciting. It's got huge potential. They've also done huge leaps and bounds in the irrigation as well. But, saying that, we also excel

in certain parts of the industry over here with some of our gene research as well. As for anyone that is actually doing it and smashing it out of the ballpark, because it's reasonably new, people are still trying to work out how to do it. I think some of the Australian farmers are actually leading the way—for example, the SwarmFarm guys with their automation process.

The Hon. CHRIS RATH: I'll turn now to AI Group. I asked this question of the BCA this morning as well. How do we bring the community with us in terms of—there are obviously huge benefits to Australia in terms of productivity that we can get from AI but the downside is that sometimes there's a fear of the unknown and a fear of potential job losses as well. How do you think that government and business can work together to alleviate some of those concerns out there on AI?

LOUISE McGRATH: I know that obviously the Federal Government is doing its own regulatory framework. National AI Centre are also working with industry to develop what they call a safety standard, but it's around responsible use of AI. While we don't want onerous regulation that will stifle innovation, we are aware that some parameters and some guidelines will be very useful for businesses to develop confidence in the take-up of this technology. I had one member from the construction sector say to me, "Well, we used to think asbestos was safe", so they don't want to be the ones holding asbestos or AI if it goes wrong.

In terms of how we help companies and our members in particular take up this technology, because it has the potential to be disruptive as well as beneficial our recommendation to them is that they should use it for the first time—if they're not already using it in a lot of robotics and other tools; as I said, it's an old technology—in terms of generative AI, they should use it not in a production environment straight off. So to many of our members, our recommendation is to use it in safety. We've been holding sessions for safety officers to practise with real tools that are available—it's not in theory; it's using augmented reality, the metaverse and sensors—so that the safety officers can have some confidence.

As a way of getting AI into an organisation, it's a great strategy because, number one, safety is something that Australian companies do well and are focused on. They already have the business case, so if they can see a way to make that more effective they will take it up, particularly for boards; and, number two, if it is a safety program that's being rolled out across the organisation, it means the whole organisation—all staff—will have exposure to it, which will grow confidence and reduce some of the fear of the unknown. That's just one example. For us, working with CSIRO and the National AI Centre, we are running workshops monthly all around the country with real demonstration of the use of AI in businesses.

The Hon. CHRIS RATH: To Interactive Games and Entertainment Association, in your submission in terms of regulatory consideration you're probably arguing for a more light-touch approach on regulation. What do you think the risks would be if the State or Federal governments were too heavy-handed with regulation? How do you think that would impact your industry?

RON CURRY: I think there are two parts to that question. The first one is around consistency in legislation. We know that each State—although New South Wales is ahead here—is looking at what the regulation might look like and various Commonwealth Government departments are looking at what that regulation is. The first thing is how do we get consistency in that. Because consistency is surety and that's what business is after. The thing that we're worried about is Australia is really set from a number of other levers that various State governments and the Commonwealth Government have put together around growing our industry. We want to ensure that the industry has, for want of a better word, a level playing field with the rest of the world, and that if there are AI obligations, that they're consistent globally as well, and that we don't isolate Australia as a place where games can be made and instead taken offshore. It's fairly easy to move where games are made and we would hate for Australia to lose that because of an over-regulated burden on the industry.

The CHAIR: To the NSW Farmers, in recommendation 4 you talk about the connectivity challenges faced by the agricultural sector. I take your point that mobile phone connectivity can be patchy. But satellite technology—and I'm thinking of things like Starlink or nbn Sky Muster—offers speeds far in excess of cable coverage in the cities, so how specifically are you asking the New South Wales Government to improve regional connectivity?

ADRIAN ROLES: It's about developing the infrastructure that allows the regional communities to access those services. So I agree with you, Mr Chair, that in the medium to long term a lot of the connectivity issues for regional New South Wales will be addressed by these new technologies coming. But as it stands at the moment we have the opportunity to backhaul, especially when it comes to talking about AI, not just about making phone calls. I'll park that because that's a whole other issue that gets people from the country very emotional.

We're talking about this from the aspect of allowing the data set required to move in real time to allow for the use of hard algorithmic or AI to be used on farm for productivity and gain. For example, the Farms of the

Future program was a great start in allowing the adoption of on-farm connectivity that would allow data from centres in situ to flow back, backhaul up through those solutions that you just mentioned, Chair, that would allow the adoption of the technology. In summary, if you could help provide the circumstances to allow on-farm adoption, that would be greatly appreciated from NSW Farmers.

The CHAIR: Fantastic. One thing that I'm particularly interested in and excited by is the combination of AI and robotics and automation. I follow all these Instagram pages where you've got robots picking grapes and robots—little drones—hunting, spraying, all the rest of it. My neighbour sprays his corn crop now with a robot— with a drone. What are the opportunities in productivity there—which is a pretty obvious question—and what about labour on the other side of it? With automation there's a lot of jobs in labouring in agriculture. What impact will that have on those types of jobs—unskilled labour—and how far away is that?

ASHLEY COOPER: Thank you, Chair, for the question. The Australian Farm Institute has estimated that the GDP of Australian agriculture could increase by about 25 per cent, or \$20.3 billion, with the improvements that Adrian has outlined and we've sought, as per recommendation 4 in our submission. It's a very valid consideration that we have discussed in different committees of the association. I think in any move forwards in recognising the advantages and productivity gains and efficiencies et cetera that AI and other agricultural technology can bring, there's the very real need to balance that with the fabric of regional Australia, and that is the workers that you've identified.

As most people in this room would be aware, there are significant issues in attracting workforce to Australian agriculture—it's not limited to New South Wales—and some of that comes to the tasks that are asked of these people too. Where there is an opportunity for technology to undertake the work that might be mundane or jobs that people won't want to do, it is going to be of benefit for not only that farmer being able to get that work done but obviously for agriculture as an industry and then for Australia as a nation with the outcomes that we can produce through our food and fibre.

ADRIAN ROLES: To reinforce from my perspective, by utilising technologies like that that remove the mundane tasks that come along with a lot of the labouring side of agriculture, it gives us the opportunity as farmers—as business people as well—to actually use that labour unit to help us grow productivity. It's also an opportunity for us to invest in our own workforce so we get them upskilled.

The CHAIR: My next question is to IGEA. You state in your submission that the video games industry is a pioneer in the development and application of AI, and often several years ahead of its time. Can you tell us about some of the emerging technologies and what we might be seeing spread to other industries in coming years?

RON CURRY: We've been using AI for a number of decades now. It's kind of inherent in how a game is made. It's also inherent in how a game is played and how a player will interact with the game. Some of the things, if we start moving forward, looking at taking those learnings—I think Charles mentioned in his opening, around serious games. From a training point of view, we're able to create games or learning modules where the AI can represent—one of the ones we were talking about was around counselling. It can present someone with a particular problem and a counsellor can learn how to engage with a client or a patient by the use of the AI function within the game. We've seen the same in training for surgery. Doctors are using video games to train in surgery and the AI can simulate what's happening in a surgical situation.

Similarly, the Victorian police for a while have been using AI simulations to train police officers in the use of tasers, because the AI can adapt to changes in a person's ability in a situation. It can look historically about how a police officer may have reacted and change the scenario. So there are many ways from a training point of view. It's also quite adaptive in tourism. We've seen the use of AI when you mix it with things like augmented reality. For example, you can go down to Parliament House today and you can overlay what Parliament House would've looked like or what Parliament would've looked like in the 1800s or the early 1900s, and the AI will adjust to what you're doing and the questions you're asking—what may be happening in Parliament at the time. They're a few examples.

The CHAIR: Thank you very much. That's very useful.

The Hon. CAMERON MURPHY: I might follow on from there and ask the Interactive Games and Entertainment Association a couple of questions about that. I don't know whether you're familiar with some of the other submissions to the inquiry, particularly from the Copyright Council and the Writers' Guild, but some of the things that they point out are that often exactly that type of content that you've just described in answer to Mr Buckingham's questions are based on copyrighted materials that form the input to the AI. Do you compensate those creative people that produce that copyright content for it during the learning process? Do you think there's an appropriate model to ensure that copyright owners are provided with appropriate compensation for the use of their material in the learning process of AI, as well as whatever comes out of the output?

RON CURRY: If I miss answering any part of that, please repeat it. I will start by saying we are a content industry. We create. We're copyright owners, whether that's in the code, the music, the art or the movement. A video game is a whole series of content creation and IP in and of itself. So, yes, we're very aware of people exploiting IP. A number of the datasets that we use for creating a game, if that's what you're talking about, actually come from within house. We create our own dataset from our own businesses, because the DNA of the product we're building is reflective of the business who's creating it. So in some ways the data we're using is our own, or our member companies' data. As far as how we compensate people for the use of AI and datasets, that's something that we're currently working through with the Commonwealth Government. Charles sits on—I think it's the copyright round table. But we don't have an answer for that yet. I don't think anybody does, but it's certainly something we're very cognisant of.

The Hon. CAMERON MURPHY: I wanted to ask the Australian Industry Group a similar question. A lot of the generative AI that's being used is going to be based on copyright material. If you're, for example, replacing mundane tasks in a workplace or you're generating content like pictures or artificial video to explain things, should businesses be compensating those copyright owners because they're the ones using their material both in the learning but also as an output from that AI?

LOUISE McGRATH: Yes. Our policy position on AI is that it is the action that needs to be regulated rather than the technology. Stealing copyright is not permitted if you're using a pencil, so it shouldn't be permitted if you're using AI. Similarly, if the behaviour we want to regulate is line of customers, our position is that whether you're standing on a snake oil box or using fax machines or using AI, it's the behaviour that should be regulated. It should not actually be around the technology.

The Hon. CAMERON MURPHY: How do you do that if it's a large number of inputs? A program like ChatGPT is just trawling through the internet and all of this copyrighted material in order to then generate some output, which is based on that. How do you actually determine who should be compensated and how much they should get as part of that? Do you have a view on that?

LOUISE McGRATH: There are some experiences with old technology that we can draw on to determine this. I'm not a copyright expert, so I don't want to go beyond my expertise. But if we think about photography, it's not the manufacturer of the camera, it's not the subject of the photo, it's not even the person who styled the photo who owns the copyright of the photo; it's the person who clicks the button. When we think about regulating AI in relation to copyright, there are already creative enterprises where many of these questions have been answered.

The Hon. SCOTT FARLOW: I've got a couple of questions for Farmers. On page four of your submission—this is in the second-last paragraph—you talk about the range of digital technology being used in ag to assist farmers to make decisions and gain efficiencies, including water management and maintenance alerts, remote monitoring, management of livestock, crops and soil et cetera. I was wondering if you could give us some real-life examples of what that looks like on a modern farm?

ADRIAN ROLES: Thank you very much for the question. Some real-life examples are from, say, the broadacre cropping sector. A lot of these technologies that we mention, like GPS guidance and auto section control, have been wholeheartedly adopted across the industry. Back in the early days, when I first got my teeth into this in the early 2000s, we used to talk about the people who had all this technology. Now we talk about the people who don't use this technology. The flow-on effect from that means that there have been lots and lots of datasets that have been generated that are generally sitting on farm, sitting in service somewhere, that farmers don't either have access to or they don't understand, or understand how to utilise.

In regard to AI, a lot of that dataset will be required to drive them to increase those efficiencies. So that's broadacre livestock, but for livestock, for example, the use of Australian Breeding Values and those indices have allowed for the increase of productivity. I know in my personal flock, the use of the EID tags has allowed for it to actually increase by 30 per cent in genetic potential, so that drove my direct link profitability. By adopting these technologies, there is a real flow-through to the bottom line. That's allowed us to expand our business, employ more people and have more money flow into rural industries and rural communities.

The Hon. SCOTT FARLOW: Then you talk about in the same paragraph:

Future applications will continue to realise data-driven technologies that use automation, robotics, artificial intelligence, machine learning and satellite and remote sensing capabilities.

I was wondering if you could paint a picture in terms of what is that going to look like in this farm in the future? What sort of things are we actually talking about? Are we talking about automated cropping, where the header is driving itself and the seeds are being planted themselves? What are we actually talking about physically? What's it going to look like?

ADRIAN ROLES: For example—and my apologies to Mr Rath too, because I just thought of a great example from overseas where they've been using these technologies.

The Hon. CHRIS RATH: Please, tell us.

ADRIAN ROLES: Hands Free Hectare. Is anyone aware of that? It's out of Harper Adams University over in the UK. Effectively, they have been growing a crop. The initial start of it was they grew one hectare of barley where no-one set foot in, the whole season. So it's sowed through automation, harvest through automation, all-in crop sensing. Agronomic work was actually made on those decisions. So it is coming, and it does have the potential for us as an industry to not start looking at our management zones by flock or herd or rows of trees or rows of vines, but actually start treating it by individual plant management.

What that now is meant to do is it has a huge opportunity for farmers to increase their productivity, reduce their environmental footprint, because the principles of—I don't know if this Committee is aware—the five Rs: right place, right time, right way, right directive, right product. It allows us to utilise those because we'll actually have the datasets to make those informed decisions. So we will eventually one day get down to where we're doing individual plant management in wheat crops. But for that to be happen, we're going to have to have the infrastructure behind it. We have to have the data protections in place to ensure that the data that is coming off farms are used for the right purposes and not used against or weaponised against New South Wales farmers. We have to, effectively, allow for the right environment to do it. To answer the question before, we're also going to need to have some really smart people consider agriculture as a future career.

The Hon. STEPHEN LAWRENCE: What does "individual plant management" mean?

ADRIAN ROLES: There was a trial over in Germany, where they did it on radishes, which isn't really representative to a lot of Australian growers, but it sells the concept. They had this little robot called Botti. In literally a 50-metre by 20-metre plot of radishes, this robot would trundle along. It was really lightweight, so it reduced soil compaction issues. For Australian agriculture, that's really important, as we require—we're basically water harvesters, which is how we work, because we are on the driest continent in the world. It allowed water infiltration. That was one upsell.

The second thing it did is it actually had its own ion-based sensors. So it would test for nitrogen levels in the plant, phosphorous levels in the plant, macro- and micronutrients. And then what would happen is it had its own onboard reservoirs of fertilisers. From that information that it collected from its onboard sensors, it would then create its own special fertiliser blend specifically for that plant so it could optimise the plant's growth, based along sound agronomic principles.

The Hon. STEPHEN LAWRENCE: Is this a machine that moves along the ground?

ADRIAN ROLES: Yes. Effectively, you could call it a drone or a ground-based drone.

The Hon. STEPHEN LAWRENCE: Okay.

The CHAIR: How did the radishes go?

ADRIAN ROLES: I must admit, apparently they grew quite successfully. I really do admit it was blue-sky R and D type stuff, which is another thing we need to foster in New South Wales to ensure that these technologies will be updated, because without that blue-sky fundamental research, it's a real step for a commercial reality to be adopted on farm without some sound agronomic business information behind it, which will come from the blue-sky research.

The Hon. STEPHEN LAWRENCE: Just by way of example, what do you think an Australian wheat farm is going to look like in 40 or 50 years, in terms of everything through from planting to cropping and so forth?

ADRIAN ROLES: There will still be a lot of human involvement. To this day, still the best dataset that I think we will have, in the Australian and New South Wales agricultural sense, will sit between the farmer's ears, because, unfortunately, we are growing in a very complex biological system that has too many variables. So there is a lot of gut feel, which I'll prefer to call wisdom, that farmers bring. But to answer your question, it is going to require a lot of automation. I actually see the future of some of the agronomic and service industry becoming more specialised, because AI and plant recognition and some of the camera ID technologies coming in the industry will remove a lot of that in-field grunt work. Also, too, because of the scale of the operations, no human can physically get across all of it. So it will actually allow for a lot more insight and better management on ground, using these technologies. But we still will be putting seed in the ground, and we'll still be praying for it to rain.

The Hon. STEPHEN LAWRENCE: I think you could be right about the human knowledge, because when the last drought broke in Dubbo, there was an old farmer who was in the newspaper and predicted the month

and the year and also what part of the month it was going to break in, based on his farm records from 1870 onwards.

ASHLEY COOPER: I'll just add to that. That's why we've also focused on the need for that investment in vocational education and skills and training, because people will always be at the core to the industry. No matter how automated things might become, if we reach that potential—the shift in jobs, for example—it's just going to change the skills that are needed to support the industry, depending on what the future may look like on that wheat farm.

The Hon. STEPHEN LAWRENCE: Just on that, you talk in recommendation 6, I think it was, of the need for the State Government to invest in infrastructure. So you talk about how the State Government should prioritise funding to stimulate industry and regional digital capacity and capability through investments in infrastructure and skills. What sort of infrastructure do you have in mind exactly?

ASHLEY COOPER: We've taken a broad sense in some regards, because we're not the experts on what the specific infrastructure is. If I could give a personal example, I'm very fortunate that I get to work from home in Gunning, halfway between Goulbourn and Yass. Whilst I've got fairly decent 4G mobile reception, my internet isn't that great, whereas I travel down 10 minutes down the road and they've got really good internet but the mobile connectivity isn't that fantastic. As Mr Roles was saying earlier, it's about everything that we'll need is needed to enable all sorts of different technology, whether it's—it was years ago that the livestock sector was looking at those low-orbit satellites for livestock traceability, and here we are now, I think I see more Starlink antennas on roofs around me than I did a year ago. So things are moving outside of agriculture that can be used in agriculture, but we need the technology and the infrastructure to support it across regional New South Wales.

ADRIAN ROLES: To add to that—I know it's probably a Federal issue, but it would be great if the New South Wales Government could assist—we need a denser array of weather networks. A lot of this stuff that we're going to do—as we mentioned before, we're all moisture farmers. Some of the key infrastructure is actually ensuring that we actually can get the dense enough datasets to allow the adoption of AI.

The Hon. STEPHEN LAWRENCE: Do you think that AI is going to assist the sector to deal with the effects of climate change, particularly in those areas on the fringe where the predictions seem to suggest that the viability of some of the farms in those areas might be affected in coming decades?

ADRIAN ROLES: We are potentially seeing declines in agronomic productivity already. I do think that these technologies, through the automation or the rest of it, as I mentioned in the individual plant management or individual animal management, will allow us to do more with less. So it's not about actually growing more crop, at times, as a farmer. I don't want to grow the most crop; I want to grow the most profitable crop. If I can grow the most profitable crop and maximise the returns out of my business, it will allow me to have a sustainable business. If I have a sustainable business, I can then afford to invest in good environmental outcomes and investment in staff, and that flows through the whole wider community. To answer your question, yes, I do see it being a major driver of agricultural productivity.

ASHLEY COOPER: I would just add, too, that I think it goes back to what you said earlier, that the brain will always be needed at some point in the decision-making. It's about supporting decisions to be made on farm, even if that is changing the commodity being produced—for example, if it's able to give better insights as to what that future climatic environment might be to make those decisions earlier so that it's not a reactive decision, perhaps. There's all of those potentials, I think—the potential for this technology to support the community more broadly, whether it's increased weather information coming into the bureau, for example. Who knows? The opportunities are endless here.

ADRIAN ROLES: I should have added, too, it has got a huge potential for risk mitigation when it comes to drought and even frost management, because we are going to see an increased variable environment moving forward.

The Hon. STEPHEN LAWRENCE: Is there an issue here in terms of the tension between the family farm and big ag in terms of the more productivity is driven by technology—this particular infrastructure that people are going to have to invest in—is it going to create issues about sustainability of smaller farms as opposed to big ag?

ADRIAN ROLES: So a couple of things to unpack in that question. Thank you for asking it. Unfortunately, adoption of these technologies is not going to be a rising tide that rises all boats. It will favour the people who are already doing things very well. To be honest, if you're already making bad agronomic decisions and you adopt these technologies, they will only exacerbate those agronomic decisions and you get worse outcomes. Same if there is a division between family farms. I'm off a family farm. We've been using a lot of these technologies for well over 20 years. We would say quite often we have been on the bleeding edge of these

technologies, not the leading edge. But, in what I do in some of this consulting work, I actually see that family farms tend to be nimble and will adopt technologies that have good return on investment before even the large corporates will.

The Hon. STEPHEN LAWRENCE: It's funny you say that, because I went to the wheat awards in Dubbo a few years back and I think all the winners were family farms.

ADRIAN ROLES: And the day of thinking of Dad and Dave as the family farm is not the case. Most family farms now are of considerable scale. They are multimillion-dollar businesses—and that is one of the things. Some of these family farms would be big enough to list on the Stock Exchange. And, because they have that really short business structure from the farm gate to the decision-maker, they can react and take advantage of these technologies before it can even come up—before some of the corporates can adapt. But, saying that, when the corporates adopt—because I have done work with corporates—they adopt wholeheartedly, and they can roll it out across their footprint very quickly because they have the economy of scale and generally the financial backing to achieve it.

ASHLEY COOPER: With some of the recommendations we've made, we've spoken about the need for that research and development in ag tech to be able to reach extension or commercialisation. But, to that point, it has to be affordable as well. I just wanted to add, too, because some of these businesses are not corporates in that context of corporate Australia as such—they are family farms—the need for the right security around the data that is being collected and clarification of ownership and all of those issues, which were touched on previously by one of the attendees also, are essential. That regulatory framework has to be correct.

The Hon. Dr SARAH KAINE: Thank you all for appearing today. I have a question which is a follow-up, I think, from the question that Mr Lawrence just asked, for Farmers and for the AI Group. You started to talk about this, I think, in your previous answer with regards to farms. Are you able to tell us about the proportions of your membership—where they fall on an adoption continuum? How far are different parts of your membership along that journey? Are the majority at that end looking at what's next, or how is that spread?

ADRIAN ROLES: Across the New South Wales farming community, we've got a breadth. If you look at the adoption curve of these technologies, what I like to jokingly call the PlayStation generation—the young farmers coming on farms now—they are not afraid. They embrace this technology. They're not afraid to get stuck in and they're not concerned if it breaks. Whereas, at the other end of the spectrum, we have people who will never adopt it. The only reason they'll adopt it is because they'll leave the industry. To answer your question, there is a broad breadth, but it is a growing sector of the industry that is adopting it, for example, in our farming business.

It also comes down to the personal interest of the person. My father virtually cannot drive his phone, but when the datasets for the young mating ewes come in, it's amazing how quickly he can drive the iPad to actually start interrogating those datasets. To me, that is also an opportunity we have as an industry to actually start bringing those people in and upskilling them through the potential of training courses and all the rest of it to actually make sure it's adopted wider.

The Hon. Dr SARAH KAINE: On that, I know one of the recommendations is about training and upskilling. I know you said you're not necessarily experts in what that looks like, but I do wonder if you had any thoughts, because you're talking about upskilling people who might have been doing this for a very long time, not just your new people coming through that you want to attract. Have you had thoughts on the best ways of doing that?

ASHLEY COOPER: I think it's a variety of ways. I think it's about digital literacy, whether it's courses that people can either attend in person or it's online training that's delivered. But we also very much, in our advocacy on certain tools—they call it the electronic National Vendor Declaration, for example, which is the livestock movement document. The usability of the interfaces has to be easy for people of various digital literacy skill sets. So, again, the specifics there. I think what's fit for some is not necessarily fit for others. We'd welcome a broad range of options to deliver that upskilling and digital literacy improvements.

ADRIAN ROLES: And just to add, too, the farming industry are not greatly concerned about accredited training. If there is the ability to have non-accredited training in this space, it would be greatly appreciated. I have, in the last 18 months, been asked to review some units. It's units of competency, I'm led to believe. It's outside my area of expertise, but the ones I was presented with were not fit for purpose. That's just solely because this is a very fast-moving space, and I would imagine it would take considerable time to develop those competencies. I would encourage the consideration for any training programs to be non-accredited so they can move with the times, so to speak.

ASHLEY COOPER: And I think, if I could just add very quickly, the agricultural space is one that has demonstrated, particularly historically, when you think about the R&D days and the field days and the DPI-led

extension of their research and development, that it's sort of a known model. It's about getting people together and helping them learn and making sure that extension is part of that commercialisation and realisation.

The Hon. JACQUI MUNRO: My question is for Mr Curry. Obviously, lots of people in Australia are gamers. I think you have released a report saying that four out of five Australians identified as gamers, which is huge. If you've got 17 million people using games and a lot of those games are using AI, is there any awareness of how many gamers understand that they're engaging with AI systems?

RON CURRY: I guess the answer to that simply would be no. I guess, as with any entertainment product we're engaged with, we don't tend to look under the hood. When you say AI, I guess it depends on what you mean for that. Games for the last 15 years or so have used AI as the basis of creating the game. I guess, even if they are aware of it, it's just what they've always known, and in some ways it's what they've always expected. They expect the game to function in a particular way, and it functions in that way because of artificial intelligence. Does that answer the question you're asking?

The Hon. JACQUI MUNRO: I guess my question or inquiry is really about whether people are worried that they're engaging with technology that might have AI built into it somehow, and they either know or they don't know. Or people find out and then say, "Oh my gosh, I've been engaging with this technology that has all these assumptions about how I play a game, or why I'm a particular type of gamer, and that's changing my experience." Or do people not really care very much, because the output is interesting enough or useful enough for them to not need to worry about what's under the hood?

RON CURRY: I think, for many years, they haven't worried about what's under the hood because it's nebulous. It's had no impact. It creates a rock while they're moving, or it sends a spaceship in the right direction, or it creates a landscape within a game. I think where we're landing today is generative AI, and that's what has captured the attention of consumers. They may then look at a game and go, "Okay, what does this mean now? What does my input into this game mean for generative AI?" I think we are getting to a new stage where people's appreciation of AI is changing, but also, I guess, we need to educate them on what is the difference. What does the AI that creates a rock look like as opposed to what does taking your data and creating a dataset look like. There are kind of two ends to that.

The Hon. JACQUI MUNRO: I guess it's now more about, "Well, if I put in these inputs about myself, does that change my experience of a game compared to somebody else," for example, and whether people are expressing concern about that or interest in that.

RON CURRY: True, there is that. There are continuous levels and layers there. Also, the games are using AI as a way of looking for hate speech, for example—so whilst they may consider some parts of AI perhaps intrusive or want to know more about it, other parts of it are trying to make the experience a safer experience.

The CHAIR: Thank you, Mr Curry, and thank you all. The time for questions and answers is over, but we thank you for your attendance here today or online and for your submissions and your work in advocating for your industries. It's very much appreciated and it greatly assists us in our inquiry and deliberations. Thank you all again. We will be in touch in due course.

(The witnesses withdrew.)

Ms EILEEN CAMILLERI, Chief Executive Officer, Australian Copyright Council, sworn and examined

Mr NICHOLAS PICKARD, Executive Director, Public Affairs and Government Relations, APRA AMCOS, affirmed and examined

Ms CLAIRE PULLEN, Group Chief Executive Officer, Australian Writers' Guild, affirmed and examined

The CHAIR: Good afternoon and welcome to the Portfolio Committee No. 1 inquiry into artificial intelligence. Thank you for coming along today. I assume you have some introductory remarks you'd like to make. We'll start with you, Ms Camilleri.

EILEEN CAMILLERI: I'd first like to acknowledge the Gadigal people of the Eora nation as the traditional custodians of country where we are and pay my respects to all Elders, past and present, as Australia's first cultural creators. The Australian Copyright Council is the national independent not-for-profit non-government organisation dedicated to promoting the understanding of copyright law and its application and, importantly, to fostering collaboration between content creators and consumers. As representatives of the peak bodies for professional artists and content creators working in Australia's creative industries, the council supports initiatives that strengthen and broaden the copyright framework for the proper respect and remuneration of the use of creators' works.

We note that this inquiry comes at a time of other investigations into AI, most notably the Attorney-General's ministerial copyright round tables in 2023. As I understand that you know, AI was a part of those discussions and led to the establishment of the copyright and AI reference group, which held its first meeting on 22 February. Whilst I don't propose to go through copyright law in detail, I'm happy to answer any questions now or after—in brief. I know that a couple of panellists earlier today have touched on it, but copyright law is based on Commonwealth legislation, the 1968 Copyright Act and, of course, the judicial interpretations of that Act. Copyright is the only non-registerable intellectual property right in Australia. So copyright grants a bundle of rights to the creator or owner of copyright material including, for today's purposes, the reproduction and the communication to the public right—that is, the right to make copies and the right to make copyright material available online or, what we say, electronically transmit via the Internet, for example.

A copyright owner also has the right to grant permission to others or license others to exercise those same rights. Charging fees or receiving royalties for the exchange of that permission is the most common way that copyright owners derive income for their creative material. This copyright law framework, which includes moral rights, aims to balance the legitimate interests of creators and copyright owners with the benefits that the public may gain through access to that cultural material. It does this through provisions that outline circumstances where, for public policy reasons, the use of copyright material without permission from the copyright owner will not be an infringement. These provisions take the form of sector-specific, statutory licensing schemes, for example, or exceptions that serve specific public policy purposes, for example, reporting the news.

This framework is integral to the discussion of generative AI, both in relation to the inputs—that is, the training and development of AI models—and to the outputs, and both the availability of copyright protection for those outputs but also the risk of copyright infringement of those outputs. The potential for AI technology misuse cannot be understated for creators, including for First Nations creatives, whose cultural and intellectual property is particularly vulnerable and not appropriately protected. The large-scale use of copyright material without transparency and permission denies all creators the right to choose if, where and when their copyright material is used and receive remuneration for use of that work. The council is very grateful to be asked to be part of this process. Thanks very much.

CLAIRE PULLEN: Warami, everyone. I'm not going to repeat the "who we are" section of our submission but I will say that the majority of my members live on the eastern seaboard and New South Wales is our largest State by membership. There are three principal areas of concern for us in the area of generative AI, but before I get to them I just want to note that how we talk about this is sometimes part of the problem. Having had the benefit of watching this morning's proceedings, we say "artificial intelligence" but we're not talking about intelligence here. We are talking about algorithmic decision mimicry, where the quality of the mimicry is determined by how closely it resembles the decision a human would have made and, in the case of my members and other creative workers, that human is a creative worker who has honed their practice in some cases over decades.

In the same way, when we say "training" or "learning", which was near ubiquitous through all the proceedings I've been able to watch today, what we're referring to is the theft of someone's intellectual property. In proceedings overseas there have already been concessions from AI companies that their models rely on and, indeed, only work if they are able to take the copyrighted work of artists without consent or remuneration or any

consideration given to the moral rights of those creators. Our first area of concern is how actively infringement is being pursued now, including in Australia and by at least one Australian tech startup that we are aware of. Any time an AI company or proponent of AI talks to you about LLMs—the large language models—it is critical for you to know that those models rely on stealing my members' work and the work of authors and writers all over the world.

We've been able to confirm this infringement has already taken place here and overseas for some of my members. This infringement also potentially exposes the New South Wales Government to secondary liability. Create NSW funds various offsets, incentives and programs for the development and production of screen and arts in New South Wales, including games. Wherever the New South Wales Government's money is in play where a creative project utilises generative AI, the New South Wales Government is potentially exposed to that risk, as are any downstream businesses in that transaction chain. This runs directly counter to the aspirations of our creative industries expressed by the New South Wales Government in our cultural communities plan.

Reducing the creative workforce is the second of our concerns. As some of you will be aware, how AI can be used in writing for screen was perhaps the most bitterly fought issue by our sister guild, the Writers Guild of America, in their recent series of actions overseas. You would all be aware of what the loss of income to any class of workers and the closure of the small businesses and the production sector could mean for the New South Wales economy. We also run the risk of AI products generating classification-breaking and harmful products in games, where the game essentially plays you right back, all the while monitoring your biofeedback and feeding that into how it learns to exploit other users. I'll conclude now with what is perhaps the most important issue, and one where the New South Wales Government can and should act now—that is, to ensure that First Nations creatives and stories are never ingested into AI models and to impose penalties where they are. Generative AI platforms present a particular threat to First Nations cultural assets and community custodianship models of story.

NICHOLAS PICKARD: I'm here on behalf of APRA AMCOS, Australia's largest music industry body, which has been representing the rights of songwriters, composers and music publishers from across Australasia since 1926. We hold significant concerns regarding the risks generative AI presents to artists, rights holders and creators through the use of their content as inputs to and outputs of generative AI. For clarity, we are not so much concerned by the technology itself. Technology empowering human expression is of course nothing new. We also acknowledge that there's a number of other useful and important purposes to which AI more generally is currently being applied. Our concern, as my fellow panellists have outlined, regarding generative AI arises from the lack of transparency that most generative AI has so far demonstrated in terms of acknowledging the content which has been scraped, mined, listened to, trained on or, to use another word, copied in order to create their outputs. To the extent that this conduct involves the reproduction of copyright-protected material without permission, it is potentially unlawful and certainly unethical.

The potential impact of this type of unauthorised use of the work of Australian creators, including First Nations creators, is deeply problematic. The New South Wales Government should seek to develop a definition of "cultural risk" that reflects the concerns of the local creative sectors, including Aboriginal and Torres Strait Islander creators. These sectors are uniquely vulnerable to AI risk and will be heavily impacted by the unchecked development and use of generative AI models. APRA AMCOS submits that the New South Wales Government's Artificial Intelligence Assurance Framework does not, in its current form, adequately incorporate the concerns of artists, creators and copyright owners. As well as this, the principles mandated through the Artificial Intelligence Ethics Policy as outlined in the framework also do not appear to include issues of cultural impact or transparency around the use of copyright material. The New South Wales Government should develop a definition of "transparency" in AI that benchmarks how AI developers and users provide sufficient information in respect of the original creative works that have been used to generate AI content.

Creators pour their hearts and souls into their work, investing countless hours refining their craft, yet they face the reality of seeing their creations exploited by AI platforms without acknowledgement or compensation. One of the biggest generative AI platforms, OpenAI, has already admitted to a UK parliamentary inquiry that it would be impossible to train today's leading AI models without using copyrighted materials. We would like to see these five questions asked of platforms ingesting creative material to train their computer system: What datasets did you use? Did you seek permission from the owners of the copyrighted material included? Where did the ingestion and the use of the datasets occur? Do you have servers in Australia? What direct licences have you taken out for the use of copying creative material? I thank the Committee for setting up this inquiry. It's an important first step to ensure that New South Wales and Australia have policies that promote transparency, accountability and fair compensation for creators within the new AI ecosystem.

Ms ABIGAIL BOYD: Good afternoon to all of you. Thank you for your submissions. They were really interesting to read. One of the things that leapt out to me—and maybe this is a question for you, Ms Pullen—is that in the process of AI replacing, for example, the notes being taken at writers' meetings and that kind of thing,

we're losing that exposure to the process at an early stage of a writer's career or whatever. Can you talk us through—are there other examples of that? What would we replace that with? How would we provide that exposure and experience if AI is doing all of that?

CLAIRE PULLEN: The first thing I would say is that we don't have to let that happen. If you're taking a screen product, for example, there is a difference between letting an AI take the notes in a production meeting, where you're talking about budgets and scheduling—I use AI on my board meetings, for example. That's different to using AI in the writers' room. We're yet to come up with a way to create writers, whether or not they've gone through a tertiary training process, other than you start out as a note taker in a writers' room. Then you might get asked to draw up a screen breakdown. The next step is that you might write a scene, for which you get a co-writing credit with another writer. Then you might get asked to co-write the episode with another writer. You go on from there, through that work, to one day being a showrunner, where you're running the creative process for the entire season or series of whatever you're up to. There is only that way to train writers. It's an apprenticeship system. It's not called that—we don't have the yearly check-ins the way that you do as an apprentice—but that's the only way to learn how to do it.

There is a place for AI in our industry. Some of my members are really excited about not having to do budgets and not having to do shooting schedules. They don't want to do their diaries anymore, and I deeply empathise with this feeling. But when we're talking about creative work in our sector, this is a solution in search of a problem. There is no shortage of writers, librettists, dramaturges or playwrights who are eager to see their works made and to practise their craft. In that way, I just want to sound a note of caution. We're all three of us on the Attorney-General's copyright round table.

We've been part of that process and went to the meeting that Eileen referred to. Quite often the pushback we get as a group is that we're Luddites—"You're not interested in learning about new things and you don't want to change an industry that's going to be as out of date as the horse and cart compared to cars." That is very much not the case, but it's also—we're not saving lives here. We're not designing new antibiotics. We're not designing new vaccines. It's important to have a bit of nuance about what we're actually trying to achieve. Nick will be able to speak more to composers and musicians, but for writers, that's how you produce a great writer: They do it for 20 years.

Ms ABIGAIL BOYD: So we've got, as you say, this kind of lumping-in of everybody together. On the one hand, we were talking earlier about how you can have programs now that can churn out the formula for different drugs and prescriptions and things. It would be great if that could be done in a way where we could socialise the benefits to people in a really easy way and maybe reduce the patent time et cetera. But then, on the other hand, we have these creative industries where, as you say, it's not like people are sitting there and saying, "Well, that writing's just not good enough. I wish we had a computer do it." There's no real problem there, but we have this sort of imposition of AI in the process where it's not always wanted. Is it purely profit driven?

CLAIRE PULLEN: The cynic in me wants to say yes. Writers and other creatives are a large expense when you're talking about a production budget that is micromanaged to within an inch of its life. It's expensive to make TV. You're talking multiples of millions of dollars per episode to make television in Australia, and these are projects that go on for six or eight or 10 years, in some cases, to get things made. Without wanting to suggest that there's a hierarchy of workers, within a TV production you can't make it without a writer, but you could make it without a third production accountant, for example.

The risk here is that it's really easy to automate parts of the writing process because the AI companies have already engaged in the copyright infringement and created the large language models. When I referred earlier to infringement already happening, my members' plays are in the Books3 database of *The Atlantic's* "Journalism Exposed", and there's an Australian startup that are quite open about the fact that they've already been inputting scripts without consent or remuneration. I think part of it is profit driven and also curiosity to see what it could deliver. But, ultimately, what it comes down to is that good creative takes time and time is money in a very expensive business.

Ms ABIGAIL BOYD: Outside of writing, are there other places or other examples of where the exposure as a trainee or an intern to a creative industry is taken away or threatened by AI? Does it happen in music composition and other—

NICHOLAS PICKARD: Yes. At the moment, because of the ease of ingestion of text and image, what we're seeing first is the proliferation of works that are text and image based. Music is the next cab off the rank, in many regards. It's a little bit more complex, but it's essentially the same. Our colleagues in France and Germany—so our equivalent society, Sacem and GEMA—did some research analysis over the last six months. It was a really big survey of their members and some financial analysis that showed that, in a worst-case scenario, revenues for

the music industry will be cut by 27 per cent within four years. That's over a quarter hit on the music industry because of AI.

The Hon. CAMERON MURPHY: I might just ask a few questions. Let's start with the easy ones. Will AI be the death of all creative industry?

CLAIRE PULLEN: Go on, Eileen.

EILEEN CAMILLERI: Where do you start with that, though! Thinking like a copyright lawyer, it's always taking it back to first principles. Building on what Claire said, we want to have good writing. The Copyright Act also protects bad writing. The bar for originality, which is one of the requirements for protection under the Act, is really low. But you've got to go back to that thing, "Are we going to stand by and allow for the wholesale stealing of work?" That's what it comes down to. Before we even get to a discussion of the transparency of what these LLMs have used, we've got to say, "Well, how did they get to use it in the first place? Was permission sought? Were licensing arrangements in place?" So I think to make a broad statement like that is that we can't let it go. We can't sit by and just let that happen.

The Hon. CAMERON MURPHY: Does it get to a point where it becomes virtually impossible to, in effect, unpick that mass stealing of copyrighted material that has formed the input to the LLMs?

EILEEN CAMILLERI: I appeared before another inquiry earlier this year into the use of AI in Australian education, and one of the committee members likened it to a cake and the ingredients that you throw into a cake. Once you throw in the ingredients, you can't pull the ingredient out to make the cake. This is the same thing. Once you've thrown in Claire's members' screenplays and Nick's members' music composition, you can't pull that out, so we're looking at—

The Hon. CAMERON MURPHY: Do we need to come up with new models in order to compensate people?

EILEEN CAMILLERI: We need to make good—

The Hon. CAMERON MURPHY: I'm thinking maybe dealing with the end users. If a business is deploying AI that has used that input in order to create the AI, should that be the mechanism that's then—the revenue from that is used to compensate the creative people that, in effect, devised it?

EILEEN CAMILLERI: I can't speak with my CEO of the Copyright Council hat on, but what I will say is, within the existing copyright framework, we've already got a regime of something called authorisation liability so that if an organisation has participated in copyright infringement where they should have or could have done something to stop it, they are already exposed under the Copyright Act.

The Hon. CAMERON MURPHY: One of the other things I want to ask about, that seems to be often forgotten when we are dealing with these large language models, is the issue about moral rights where, in effect, it's completely eliminating the creative person that has designed their work, whether it's music or a screenplay or a novel, from exercising any of their moral rights about the context in which it's used later.

CLAIRE PULLEN: And exercising their right against false attribution, which is one of the key issues. Nick and I were speaking about this earlier. One of the things that we're often told by proponents of AI is that your members are going to be fine because they will go from doing what they do now to being superb prompt engineers; it will be their job to come up with the best possible combination of inputs into the AI to get the product out the other side. I think some of the research that Nick has done has shown that the prompts are all copyrighted materials. There is a difference between what you'll get when you say, "I would like to see a *Strictly Ballroom* TV series," as opposed to, "I would like to see a TV series about dancing set in Sydney." Baz and Craig, over the years that they've developed their particular film oeuvre, have a look and feel to the work that they produce, which is central to their identity as artists.

I think where we need to land is a system that accepts that some artists will be happy with simply licensing and receiving remuneration—consistent with the licensing schemes we have now. Copyright is so much tied up with licensing. I run an entire business about it, so does Nick. But there will be artists who will say, "No, I don't want my work ingested because I have spent my entire creative practice developing that look and feel, and my moral right is around how that's used in derivative fashions." This is particularly important, for example, if you think of a First Nations creator like Leah Purcell, who's developing her own songline through the ballad of Molly Johnson. I think she should be able to say that that isn't ingested into AI. I'm usually accused of being a cynic, and I struggle with this idea that you can't know where the inputs came from and that these AI companies are all going, "Look, it's all in there now, so we can't do anything about it." Well, if it got in, it can get out. If you haven't warranted at some point that you aren't stealing other people's work, then perhaps you should have thought about that before you started.

The Hon. CAMERON MURPHY: The last area I want to ask about is in relation to traditional knowledge, so things that aren't covered necessarily by copyright because there is not individual ownership but there's community ownership over a traditional medicine or a story or something of that nature. What sort of a role do you think New South Wales as a State has in order to put in place regulations that may protect—perhaps as a form of heritage—that knowledge so that it's not exploited by AI?

EILEEN CAMILLERI: To the extent that New South Wales can set best-practice guidelines, I think that would be amazing. There are, as you'd be aware, a number of initiatives going on at a Commonwealth level, particularly through the Office for the Arts, to look at standalone protection for Indigenous cultural and intellectual property. But, apart from that, and following on from what Claire said, remuneration in those cases will not be enough; it just won't be relevant because the complexity of issues—and I can't speak to those—is vast. There is a gap in what the Copyright Act can do for that kind of work, because the Copyright Act does not protect material that's over 70 years old. To your question, I think some best practice guidelines—that New South Wales will be well placed to do that. There are other States that have set in place guidelines about the use of cultural knowledge, so that would be great.

The Hon. CAMERON MURPHY: One of the problems that has been identified in some of the submissions is that it might be AI that is overseas that is exploiting that traditional knowledge, in effect. What can we do about that in Australia, really, if it's beyond our jurisdiction where it's being created? Is there another formula or mechanism by which we can provide people with those moral rights or compensation if that's required?

EILEEN CAMILLERI: Sorry, I'm talking a lot. The Attorney-General, apart from the round tables last year, was also looking at copyright enforcement. Within the copyright enforcement space there are mechanisms underway by the attorney to look at perhaps the establishment of another quasi-judicial entity to look after smaller claims outside the existing legal frameworks and perhaps, as part of that process, cultural knowledge can be looked at. I understand IP Australia is also looking at it, in addition to the Office for the Arts.

CLAIRE PULLEN: Just to build on Eileen's answer, the issues of territoriality are not necessarily confined to First Nations creators. There is the difficulty of—now that I know my members' plays are in the Books3 database—how I sue a company that's based in America but their finances are not there and then their corporate headquarters is somewhere else. There are issues around that. But for screen and certainly for a number of the creative products that my members are engaged on, whether or not government money is put towards the project, whether it's through an incentive or a grant or through development funding, can often be the difference between whether the thing is made or not. That's where the New South Wales Government can exercise leverage with a relative degree of ease, simply by saying, "We will not give you this grant, you will not be funded to do this development, we will not give you this offset if you have used generative AI in this process. You must warrant that you haven't done it, and we may, down the track, if we become aware that you've done it, ask you for the money back."

The Hon. JACQUI MUNRO: Can I just jump in on this? I'm curious about the role of blockchain technology and NFTs here, because it seems as though there are different issues where you have—the first one is attributing work to an artist in the first place, and, arguably, you could have some sort of blockchain system to secure a person's output with their identity. There is the issue of that information or that output being used without their consent, but then there's also the issue of saying, "Actually, as an artist I am happy to be paid for that content to be used in certain ways," and you've got the opportunity, essentially, to have choice among artists. Are NFTs being used? What's the uptake of that kind of technology to secure that ownership?

CLAIRE PULLEN: There isn't an equivalent technology for LLMs, for the written word. I am aware of—I think it's a crowd-sourced project going on, where artists who don't want their visual art used as part of an engine can add a code. It's called poisoning the model and it means their work is no longer legible for being ingested by the AI. That still puts the onus on the artist to do the work and say, "No, I don't want this used." I don't necessarily know that the technology will develop as quickly as the ways around it, given the commercial imperatives on both sides for the written word, which is where my members are occupied.

The Hon. JACQUI MUNRO: Sorry. Are your members exploring opportunities to make that possible? If you've got files that are associated with a particular written word or you have code embedded for a particular website that has that content, are those ways of securing written content being explored?

CLAIRE PULLEN: Without knowing where the content has come from, it's really difficult for my members to follow the chain of where it's ended up. It is something we have looked at, particularly when we saw the reportage around how you can poison the visual AI models. That was something my members were really interested in. But unless you know what the ingestion point was—and a lot of this will be overseas—it's really difficult to exercise that control.

The Hon. JACQUI MUNRO: So there is a difference between the work that has already been obviously used, but then it's also about how we create a system in which it doesn't happen into the future.

CLAIRE PULLEN: That's right.

The Hon. JACQUI MUNRO: I guess, from that point of view, it's about exploring new technologies like blockchain which could be associated with particular works to try to secure them or to secure ownership, let's say.

CLAIRE PULLEN: Starting with transparency, knowing where the thing came from.

The Hon. JACQUI MUNRO: Yes, definitely. Sorry to interrupt.

The Hon. STEPHEN LAWRENCE: No, not at all.

The CHAIR: I have a question to Mr Pickard. In the context of the New South Wales Government's Artificial Intelligence Assurance Framework, you were mildly critical of that in your submission for making no mention of cultural harm. You urged the New South Wales Government to develop a definition of cultural risk. Beyond the issue of copyright, could you talk to the areas that need to be explored?

NICHOLAS PICKARD: I think that really needs to be done in consultation with the sector, ideally through Create NSW. They really need to get on the forefront of this work and start consulting very quickly with a whole heap of user groups. That would be not just musicians, songwriters and composers but also the visual artists, authors, poets and screenwriters. That consultation would really bring to light what needs to be outlined in that. I think there needs to be quite a specific consultation very much along the lines of what the Federal Government is doing around the development of ICIP—Indigenous cultural and intellectual property—frameworks as well. I think that would really be the first starting point for the Government, to get those references in there.

The CHAIR: That's how you would get the reference, but what are the cultural risks? What are some of the other cultural risks beyond copyright that we should be mindful of?

NICHOLAS PICKARD: Beyond copyright—copyright is really just the legal function that supports the return on investment for a creator's storytelling, for their artistic creation. So I think that the harm really goes to the very stories that are being created in this State about those experiences, about those stories and about the types of music, literature and film that we want our kids and the next generation of kids to be consuming. The cultural harm goes to the very core of our community and the stories we tell each other and the stories we tell ourselves.

The CHAIR: To explore that further, it's not just the type but it's actually the quantum too. We could see a lot of people walk away who are already in the industry or people considering—like the next generation—saying, "Well, I've been replaced. I'm not going to go down that pathway of 20 years of honing a craft or learning an instrument." Is that a real risk, when you talk about cultural risk?

NICHOLAS PICKARD: To the point that Claire made earlier around the years that it takes a practitioner to hone their craft, there's a sort of cheeky statement that platforms and the tech sector will use about saying, "This will make your job easier." It's not about ease; it's about perfecting a craft and a creative craft.

EILEEN CAMILLERI: It's already very difficult to make a living as a creative artist.

CLAIRE PULLEN: We also want to make sure—and I think this is particularly important in a big State like New South Wales where the Government funds creative work and explicitly talks about the importance of having a cultural sector that thrives and tells us stories about ourselves, that we want those stories to be good. We want them to be diverse. We want them to be commercially successful overseas, like with *Bluey* where you've got, essentially, an Australian cultural juggernaut taking over the world. We have the capacity to do that, and it's important to make sure that we don't reward or—despite what Eileen says about the Copyright Act—protect bad stories too much.

One of the pitches that was given to us around a particular AI product was that it'll input all the scripts, say, for a long-running Australian soap opera. If you hit a knotty point in the story where you don't know what this character has done, who they've been married to in the past and which family drama they've been involved in, it'll help you to figure out that knotty point in the story. The writer I was on the call with, one of my elected leaders, said, "That's called being a writer." That's the craft: It's the figuring out of those things. You don't want to make it easier to make mediocre, bland, boring stories because, frankly, the canon it is ingesting is not particularly diverse. It hasn't really kept up with the explosion of new forms of storytelling. We want to make sure that, when we're using AI in our industry, it supports the creative work rather than supplants it.

The Hon. JACQUI MUNRO: Can I ask another thing that follows on from that? I totally agree that artists should have ownership over their work—no doubt about that—and the value of culture and of the creative process is what makes us human. If we're protecting certain cultural artifacts or outputs—say, for example, Aboriginal artists' work—is there a risk that it's protected so much that it then doesn't become disseminated and part of a wider cultural movement? For example, if you've got lots of inputs that are coming from a particular background or whatever, and you're missing the experience of certain people, we risk dividing it further.

CLAIRE PULLEN: From my point of view, there are sort of two speeds to this. There's what we do now about the problem we have right now, in that there is already fake AI Aboriginal art in circulation, in the same way that we know that our members' work has already been infringed. There's what has to happen now, frankly, to ensure compliance with the law. What's happening is unlawful; it's just not being policed in a particular way. There are points at which the New South Wales Government can intervene fairly easily to at least put a stop to some of that and then go through the processes that Eileen and Nick are talking about, which is saying, "Right, how do we actually make this work? What's the long-term framework? How do we make sure this serves us as a tool rather than runs us as a subject?"

The Hon. STEPHEN LAWRENCE: Just on the legal infringement question, is there any doubt that to input a creative work into an AI technology infringes copyright in circumstances where it's used to inform or potentially inform an answer that the AI might spit out, even if the text isn't reproduced in any particular way?

EILEEN CAMILLERI: The short answer is yes, if you're inputting copyright material, whether it be text, music or visual art. We're just talking about works, because there are films and sound recordings and published editions, which are separate. If you're uploading it, that's a reproduction. The upload is a reproduction. There's some conjecture in the middle part of where these systems store it in the way in which we understand. Irrespective of whether that's a yes or no, we say a reproduction has occurred at that point.

With the output, the difficulty becomes that there's a risk that the output infringes copyright if there has been what's called the substantial reproduction of an existing work, but also the situation where, as a user of an AI model, you can be liable for infringement of the output even though you don't own copyright in it. The output of material generated by AI doesn't have copyright subsist in it because it hasn't been made by a human author, which is another requirement for protection in Australia and in most other jurisdictions. Claire talked about the complexity of prompts. There may be situations over time where the skill involved in putting together a prompt into an AI platform results in a way in which copyright might subsist in the output, but it's a bit like using a pen or using Microsoft Word. I say it's a continuum, so how much of that human effort has gone into inputting into AI will dictate whether copyright protects the output.

The Hon. STEPHEN LAWRENCE: When you talk in your submission about AI training, is that the process of input?

EILEEN CAMILLERI: Yes, inputting. That's the wholesale reproduction of material.

The Hon. STEPHEN LAWRENCE: Is the Copyright Council of the view that that is reproduction?

EILEEN CAMILLERI: Yes.

The Hon. STEPHEN LAWRENCE: Has that also been confirmed in court cases or anything like that?

EILEEN CAMILLERI: No, there's been no litigation. Well, there has been no litigation in the AI space, but there have been a number of cases about the ingestion of phone books and the like. There have been analogous cases for the use of existing copyright material in that process. From a technical point of view—yes, I might just leave it at that there.

The Hon. STEPHEN LAWRENCE: It's a bit like a library, isn't it? If you put a book in a library, that would infringe copyright, wouldn't it?

EILEEN CAMILLERI: No, because—are you talking about a hard-copy book?

The Hon. STEPHEN LAWRENCE: Yes. If you put a hard copy in, would it infringe copyright?

EILEEN CAMILLERI: No, because there's no copyright act that has happened there because I'm giving it. If I photocopy a book and give it to the library, that's a copyright infringement. But if you think about—if I put a photo on display on the wall here, that's not an infringement of copyright because there's no copyright act. But if I've reproduced it, communicated it—so that's electronically transmitted or uploaded it to the internet—performed it or adapted it, then they are the copyright rights. That's the bundle of rights that I was talking about. For these purposes it's the reproduction right—so the right to copy—and the communication right which are enlivened.

The Hon. STEPHEN LAWRENCE: Do you think it's beyond doubt that that uploading act would constitute reproduction for the purpose of the Copyright Act in circumstances where it's not then reproduced to a significant degree in the output, where it somehow informs what might be spat out but it's not reproduced word for word? Is that beyond doubt?

EILEEN CAMILLERI: I'm not a technical expert, but I cannot see how a work getting into an LLM could happen without reproduction.

The Hon. STEPHEN LAWRENCE: Yes, because you're sort of copying it completely.

EILEEN CAMILLERI: Or even substantially. To copy the entirety of a work—as I said earlier, there are some exceptions in the Copyright Act which permit the use of substantial parts of copyright works, but I'm yet to see an example where I'd feel comfortable saying, "Yeah, that's cool; that's a fair dealing."

The Hon. STEPHEN LAWRENCE: The other question I have is for the Writers' Guild. It is in relation to this issue of certain creative jobs being replaced by AI and what you say in your submission about how that shouldn't occur. What form of legal regulation do you think could achieve that?

CLAIRE PULLEN: I think there's a question which mostly will end up being a copyright one, frankly, which is a Commonwealth issue that we're all dealing with at the moment. In terms of how we can address it now in New South Wales, I come back to the fact that so much cultural production that my members are engaged in in New South Wales is marginal and solely depends on government funding, whether it's a grant or a development fund being given or an offset. I don't necessarily know that it needs to rise to the level of statutory regulation, as opposed to a policy decision that says, "We won't replace these jobs."

The Hon. STEPHEN LAWRENCE: Okay, I understand now.

Ms ABIGAIL BOYD: One of the concerns of the Hollywood Screen Writers Guild was in relation to the idea that studios could use AI to churn out concepts and then writers would be used to edit or write but, because they weren't the owners of the concept, they weren't getting any of the royalties; they were only then getting a wage. Is that something that you are also concerned about, in Australia?

CLAIRE PULLEN: There's some complexity in comparing the two countries in that screenwriters in the US don't have copyright. Their rights to their work have been extinguished at a statutory level, which is why something that looks very much like royalties but is not and is therefore called residuals is a key part of every bargain that the WGA engage in. They've essentially replicated a system for payment for subsequent use, which we do through copyright and call royalties and they do through an industrial bargain and call residuals. From the conversations I've had with my sister guilds, it is possibly the most contentious issue. The reason that the negotiations were so drawn out is what role that material would play in the writing process, because your payments are linked to your credits and how much you contribute to each script determines what your credit is, which is why when you watch an episode of a television show, the last two minutes are just a list of names with very technical-sounding things next to them to describe who's the note taker, who did the scene breakdown.

It's important to make sure that the relationship between the person and the work is maintained, regardless of what the legal system is. The US is very different to us in terms of that, but they wind up somewhere similar to us, which is payment for future use of the work and for reproduction of it. The status of AI work now in the US is that there can be no writing credit attached to it and therefore no remuneration. AI works, whatever is produced, is given the status of research notes in the same way that you might look at Wikipedia, for example, and that's how it's put into the process. There's no remuneration, there's no change to the status of the writer, and there's no change to the relationship of the future use of the work because there was AI.

Ms ABIGAIL BOYD: Do they have to disclose where they've used AI?

CLAIRE PULLEN: Yes, they do.

Ms ABIGAIL BOYD: But we don't have that here yet?

CLAIRE PULLEN: No.

Ms ABIGAIL BOYD: I noticed right at the end you mentioned gambling and in-app purchases. I've been thinking about that as well and this idea that already those companies really do rely on a sort of AI. They rely on AI in order to work out and predict how to make you spend more money. The recommendation there about prohibiting AI—is that something that we've seen in other jurisdictions? It's compelling to look at it, but how realistic is that?

CLAIRE PULLEN: Realism is, fortunately, in this area not something I have to grapple with; that's more for the Committee. It's a recommendation that came out of our elected leaders who work specifically in

games, and we are seeing our game narrative designers are dealing with this more than other sectors of our membership, and it's much more acute in terms of the fact—the game plays you back, is how my advisory council member puts this. I don't know if you've seen the reportage around AIs that do what's called hallucinating, where they have inputs that are coming from places that the creators aren't sure where they're coming from, and you start getting hate speech or wildly defamatory—it goes off in a direction that no-one predicted.

Because of the nature of the relationships between the inputs, the algorithm and then what you get, essentially you can't maintain a game within a classification when there's AI involved. If someone, for example—I'm not in any way commentating any of the teenagers in my life—is using particular language and the game is listening to that and is monitoring their pupil size, checking in on all the feedback that you can get from having a camera, it will play you out of your classification; it will play you up from, say, a 15-plus to an R18 game. There is the capacity to do that.

I was listening to the earlier evidence before we came to the table, and I think a good place to start would be that you'd tell people that you're monitoring when you're playing that kind of game: that they are going to have their biofeedback monitored, that their in-app purchasing is going to be looked at. The risk here is that we end up being a follower in our industries. You know that really annoying alert you get when you go to a website and it asks you to choose your cookies option: Do you want to accept all cookies? Do you want to personalise it? That has come out of the European Union, because they have made their decisions around how websites should be allowed to track us when we are engaging on the internet. If we don't make a decision now, we're just going to end up with someone else's that might not be suited to our industry or our particular settings and the opportunities here.

The Hon. JACQUI MUNRO: I don't know if this is a cheeky question, but are you finding that there is a market for AI-created content, as in B2C, not so much B2B? Are consumers saying, "Yeah, that's cool. I love an AI-generated song"?

NICHOLAS PICKARD: I'm not quite aware of how that market might be developing. What I do know is some of our members are working with AI, and they're making that choice to work with AI development companies and maybe doing some sort of sandbox where they are providing their material, but it's done on the basis of acknowledgement, permission—

The Hon. JACQUI MUNRO: Consent, yes.

NICHOLAS PICKARD: —and remuneration. That's the very simple tenets that we are seeking, really. Artists will work with technology; they will find ways. There might be some cool stuff that would come out of that music generation that people will like. But it's another step altogether to—I know the Chair is a big fan of AC/DC. Do we want a system where a user would just go, "I just want AC/DC songs pumped out", and they don't have any say on whether their material, their intellectual property, can be used in that manner?

The Hon. JACQUI MUNRO: Do you think that people have a preference for AI-assisted involved material being badged as such, like, for example, the script writing?

CLAIRE PULLEN: I think there are two issues here. One is the question of creative control. For example, there are artists who won't allow their music to be on Spotify or other music streaming platforms because they have views about the ways in which they want their work to be consumed. So if you want to listen to a particular album, you can't do it; you've actually got to buy it. So the first issue is creative control. Whether or not the audience wants to engage it, at the first instance you have to deal with the question of transparency, consent and, ultimately, licensing around how the artist wants their work to be dealt with. On the second issue—look, there's no sign that the global appetite for content is diminishing. We just watch more TV. We engage with TikTok. We keep consuming—

The Hon. JACQUI MUNRO: Yes, we have three devices going at once.

CLAIRE PULLEN: We actually had to put some of ours down, and it feels a bit odd to only have one in front of us. Whether or not the content is good is what is going to attract an audience. There will be novelty to it. Remember when autotune came out and you couldn't listen to anything that didn't have autotune on it? But once you get the first issue of the artist's control of their work and the remuneration from it right, whether or not it's any good will be answered by the market. But without the artist having control over their livelihood and the future of their career, you can't end up with a good product.

The Hon. STEPHEN LAWRENCE: It seems like there's widescale noncompliance with copyright law going on. In terms of enforcement, particularly for creatives who aren't going to necessarily have the financial resources to take action and vindicate their rights, have you got—

NICHOLAS PICKARD: That's almost all of them.

The Hon. STEPHEN LAWRENCE: Well, yes. Have you got any suggestions about legal innovations that might make it easier to take enforcement action, whether individually or collectively—so questions of standing, funding, representative bodies and that sort of stuff?

EILEEN CAMILLERI: That is the very thing that the Attorney-General is looking at as part of the process of the examination of enforcement. I understand that he was over in New York, meeting with the head of the copyright office to have a look at how their copyright small claims court is working over there. So there is some move to do something like that. But I guess it's about the collectives, as is the Authors Guild in the United States. It's the large collectives doing that. The council gives advice to hundreds of individual creators every year, and where they have a right of action it's always whether they're able to enforce it, which is the threshold issue, as you say.

The Hon. STEPHEN LAWRENCE: I think it maybe highlights the importance of the legal provisions that allow for class actions and representative actions where firms can actually receive money for undertaking them on behalf of people who can then opt in, potentially.

CLAIRE PULLEN: It is something we have actively canvassed as part of the copyright round table. I will also just put in a plug, as someone who's spent a non-zero portion of her life appearing in a no-cost, essentially small claims jurisdiction around unfair contracts, that something like that in New South Wales would be very useful to have back. But there are mechanisms available now, frankly, that State governments aren't using. I'll give you an example of a First Nations creative that is a member of my guild. There was a particular amount of development funding given for the production of an episode of television—to produce the script to see if it would then go on to be made—and it required a First Nations writer because part of the story was a First Nations story. There was a particular amount of money, \$15,000, that was specifically for that writer. The producer who received the funding commissioned the script from my member and then didn't pay them, shut down their production company, phoenixed—started another one—and is still getting Government grants.

Because our arts funding is point in time—the assessment is done at a particular instance—when we raised this issue with the funding agency, the response we got back was, "Well, at the time, the contract between the production company and the screen agency said that they would get it, so there's nothing to see here." We have agencies who do funding but don't do any monitoring of what happens with it, so that seems to me a really practical way you could start—is actually have some enforcement there.

The Hon. STEPHEN LAWRENCE: Yes, an interesting one.

The CHAIR: We have one last question from Dr Kaine, online.

The Hon. Dr SARAH KAINE: I think that's probably a good segue into the question I have. We run a lot of inquiries, and one of the inquiries I'm chairing, coming up, is about procurement and how governments spend their money, so that's a very interesting example you've just given. I wanted to ask if you had or could provide—and I'm happy for you to send it through on notice—any examples of other jurisdictions who use that sort of economic power of the government in the way that you've suggested throughout this hearing today. If you've got any examples, that would be really helpful.

CLAIRE PULLEN: I'll have a look. Thanks very much; that's a good question.

The CHAIR: That concludes the time for questions and answers. Thank you very much for making the time today to come and give evidence. It was very powerful and we will consider it in depth. Thank you also for your excellent submissions and the work you do for your organisations.

(The witnesses withdrew.)

(Short adjournment)

Mr PETER ACHTERSTRAAT, AM, NSW Productivity Commissioner, NSW Productivity Commission, sworn and examined

Dr MATTHEW COSTA, Director, Productivity Reform, NSW Treasury, sworn and examined

The CHAIR: Welcome Mr Achterstraat and Dr Costa. Do you have an opening introductory statement to make?

PETER ACHTERSTRAAT: Just a few sentences, Chair. First of all, thank you so much for the opportunity to say a few words in relation to our submission. I look forward to any questions or dialogue after that. We've just had the 200-year anniversary of the Audit Office outside. I know that when I started as a young accountant in the 1970s I thought I was a pretty good accountant. I could add up pretty well; I was pretty fast. I wasn't good, necessarily, at reading and writing, because of my background, but I was good at numbers. I really enjoyed it and I thought, "I'm going to make a really good accountant."

Then, towards the end of the '70s and the start of the '80s, a thing came along called Lotus 1-2-3 and, from a history point of view, it's a spreadsheet that adds up for you. Excel now does that sort of thing, but when Lotus 1-2-3 came in at the start of my career, as it were, I thought, "Hang on, I'm going to lose my job. The only thing I'm good at is adding up. I'm going to be in big trouble. This new Lotus 1-2-3 is going to do it instead." So myself and all the others about my age who were good at maths but not much else—well, the others were good at a lot of things. I thought, what are we going to do?

As it turned out, over the years we got other jobs—more analytical jobs, more business analytical jobs—and the role of the junior accountants changed completely. Instead of just sitting there adding up numbers we were able to look at those numbers and do predictions and look at trends, and it was fabulous. I guess the lesson—what I'm trying to say, Chair, is that when this new technology came in, we were all quite worried to start with, but the benefit of it, not just to the organisation I worked for but also to the enriched role that I had, was phenomenal. I see similar in relation to artificial intelligence if we look after it the right way in relation to privacy, security and accuracy.

Our submission is based on a report we did about 14 months ago called *Adaptive NSW*. We've done analysis through that report which shows that, if New South Wales adopts the technologies to the full over the next 10 years or so, it can improve gross State product by up to 3 per cent. One of the other issues that we do cover is the concern about job losses. Historically, there's not a lot of evidence to suggest that there are wholesale job losses when new technologies come in.

In fact, our report shows that, while a few per cent of traditional jobs may change, some existing traditional jobs will grow and there'll be new jobs coming up. We're very positive. Our report *Adaptive NSW* and our submission are very positive about AI. We are conscious of the framework that needs to go around it, and we are supportive of the assurance map that the Department of Customer Service has got, with the five pillars—community benefit, fairness, privacy, transparency and accountability. I might leave it at that, Chair. If people have got any questions, myself and Dr Costa will be happy to take them.

The Hon. CHRIS RATH: Thank you for that opening statement, Mr Achterstraat, and thank you both for being here today. I completely agree with you: I think if this was a parliamentary inquiry at the turn of the century, we would probably be talking about candlestick makers and blacksmiths and how they might lose their jobs, but maybe they'll have new opportunities in the future with electricity and the motor vehicle. I couldn't agree with you more on any of that. What do you think the risks are to the New South Wales economy if we don't get the regulatory settings right—that is, if there's too much regulation and it stifles the rollout of AI? What would that mean for the economy?

PETER ACHTERSTRAAT: There are two parts to that question. We'll talk about what are the risks if we don't do it properly, and then I'll touch on—it's not part of your question, but it's subliminal—how do we make sure we can get the community behind us on this? To the first part about it: What are the risks if we don't do it? We're going to go backwards. It's a bit like an escalator. If you're trying to walk up an escalator the wrong way and you don't take any steps, you're going to go backwards. It's the same with technology. If we don't, along with the rest of the world, adapt to and adopt the technologies which are already there—I'm not talking about necessarily R&D ourselves; we can do that. But if we don't adopt the stuff overseas then we're going to go backwards and our productivity—which at the moment is negative—will be worse. There'll be a real productivity drop, which will cause standard of living and quality of life to fall. That's the general concept if you have it.

The regulatory arrangements around it—we do have to balance both sides of it. We have to balance. We want to encourage people to use AI but, at the same time, we've got to make sure that the regulation is appropriate in that it protects the privacy and the security and also the veracity of the data coming out of certain types of AI,

where it's possible that if you ask questions, the data coming out may not necessarily have the veracity. I'll ask Dr Costa. Have you got anything you want to add to that—the question being, if we don't get the regulatory framework right, what can be the consequences to the economy?

MATTHEW COSTA: No, nothing to add to that. I think that was a good broad overview.

The Hon. CHRIS RATH: Obviously in the '90s we saw very fast productivity growth compared to what we're seeing now. As you said, it's negative, which has huge consequences for the economy and for standards of living. Do you think that AI has the potential, for the next decade or two, to have the same benefits that the adoption of computers and the internet did in the '90s, when we saw that incredibly fast productivity growth?

PETER ACHTERSTRAAT: Absolutely. I can't put a figure on it, Mr Rath, as to whether it'll be exactly the same percentage, but it's a very exciting time. If we can do things, even stuff that's already done—the pre-filling of forms for HR et cetera, instead of someone sitting down and filling them out—it can be more accurate and it frees up all those other roles for people. You're absolutely right. In the '90s there was a boom in productivity when it was high. There are various reasons for that. At the moment it's slowed. Over the last five years, basically, productivity has gone nowhere, and in the last short period of time it's gone backwards. In a situation like that, if productivity is falling, that means there are less goods and services made per person. If there are less goods and services made per person then the standard of living is not going to go up and real wages can't go up. We need a lot of productivity levers, and artificial intelligence, I reckon, will be one of the primary ones. As our report *Adaptive NSW* shows, if we can bring in broad AI, narrow AI, predetermined and all those sorts—if we can encourage those, they will go a long way to improving productivity.

MATTHEW COSTA: I just might add to that. The scenario modelling that we conducted for our report *Adaptive NSW* was based on a methodology which sort of looked at the potential impacts of a range of technologies on productivity. We did a sort of bottom-up approach of looking at what those technologies could do in terms of productivity improvement. But from a sort of bird's-eye view perspective, one of the things we did was compare our results to that 1990s ICT boom. The figure that Mr Achterstraat mentioned about 3 per cent economic growth, that was the idea that—the scenario we looked at was that you could potentially see something like 3 per cent economic growth sustained over a 10-year period, 2 per cent productivity growth going alongside that. That is comparable to the sort of productivity boost we saw during the 1990s ICT boom, so I think the answer is yes to your question.

The Hon. CHRIS RATH: The other question is around skills and whether you think that we have the appropriate requirements in Australia and in New South Wales, in terms of skilled workforce, to actually harness the AI revolution or whether we're potentially being left behind. I suppose the flip side of all of that is the potential job losses, as we were talking about before. What do you do with the people that might lose their jobs as a result of AI, in terms of retraining or avoiding a sort of Detroit-type moment where you might have a whole bunch of people that lose their jobs and can't find alternative work? But is there anything that you wanted to talk about on skills in particular.

PETER ACHTERSTRAAT: If I can divide that question into two: Is there a skills shortage now for us to be able to harness AI? Secondly, what happens to the people whose jobs are no longer there? In relation to the first part, the skills we need to harness AI are two distinct sets of skills. There are those directly related to implementing and adopting AI—I'm being a bit general here, but programmers, people who know about data analytics and software engineers. We also need another group of people who will be able to analyse all those documentations. They may not have the IT skills et cetera, but they've got cognitive thinking and they've got judgement. They're the two broad areas of skills that we do need to completely harness this. I'm not suggesting we need a skill set to be able to develop all the AI. That would be great, but sometimes it's just as efficient to import it or to buy someone else's AI that they've already developed and modified. Do we have those skills? We can do more, and part of our report talks about the VET system and the tertiary education system doing more to encourage both those sets of skills.

In relation to the people whose roles might change, just as I was concerned with my job when I was a number adder-upper, I was able to be reskilled and it was just subliminal. I don't think anyone consciously said, "We're going to retrain these people who add up numbers." It just sort of happened and I was fortunate. I think now, with committees like this raising the attention and the awareness, and our report and others', people are very cognisant that we have to be proactive on that to start to be able to reskill these groups of people. I think the quicker we can do that, and if we can get the community on side that AI is a good thing—to do that, they need to be confident of the Government's arrangements we have around it and also have a level of confidence that they will have a place in the new world.

The Hon. Dr SARAH KAINE: I want to visit the productivity issue again. With the productivity and the decline in productivity, it was my understanding—and I'm very happy for you to correct me—that one of the

issues with the reduction and the slowing of productivity was due to organisations not investing as much in the non-labour parts of their organisations, so in their technology and those other areas in which they could improve productivity. I'm wondering if you've given any thought to—we're saying that this new round of productivity is going to be based on organisations investing in the types of productivity enhancements that they haven't been investing in because labour has been comparatively cheaper. Is there anything that you see as the catalyst for that to change?

PETER ACHTERSTRAAT: In relation to why has productivity not increased as fast as we'd like—in fact, gone backwards—there are a number of factors there. As you've said, there is the lack of capital investment to a certain extent. There's also I think what they call labour hoarding, where firms who don't have enough work for their workers are not letting them go; they're keeping them on, quite rightly, but they're not as productive as they would've been if they'd kept them on. And also the lack of take-up, to a certain extent, of technology from overseas. So they're three reasons why, but there is a raft of other reasons. There are lots of ways which we can improve productivity. Artificial intelligence is one of them. But, as you said, if we can boost the capital investment, that will go a long way. If we can improve the education and the VET standards, if we can improve planning—there's a whole raft of them. Artificial intelligence is not the silver bullet, but it could be a wonderful yardstick to be able to do that.

The Hon. JACQUI MUNRO: Thank you so much for the submission and for the Adaptive NSW report as well. I think it's really helpful, in particular, that the different types of AI are broken down into their different definitions and uses. I think we probably don't do that enough, so it's good to have a clear representation of that. I am really interested in the idea of regulatory sandboxing and how we can use that effectively to make policy decisions that are, essentially, useful—like, we can carve out spaces that allow us to, essentially, experiment. You mentioned the Spain example in the submission. I'm curious about whether you think there are particular industries or topics or occurrences in New South Wales that would be particularly suitable for a sandbox environment.

PETER ACHTERSTRAAT: I will give a general answer to that. I'm not a big fan of picking winners, as to say, "We think this is the best area to do it in," et cetera. But we'd like to promote a situation where people or industries who do have an interest to be able to do that could approach public service or the Government and say, "What about this?" and then we could do the sandbox with them. If the Productivity Commission were to say, "We think industry X or industry Y or this firm or that," we tend to be trying to pick winners, which has worked in the past, but I think we've just got to be a little bit careful about that. But the sandbox idea is one which we very much support to be able to try it before you buy, as it were, and if it doesn't work—I guess the issue there, Ms Munro, is that we have to make sure that a sandbox is a genuine pilot and not just seen as phase one, because sometimes in the past people have done a pilot, invested all this money, and they probably should have stopped it. But then they decide, "Well, we've done all this. That's phase one. We're going to keep going."

The Hon. JACQUI MUNRO: Sunk-cost fallacy, yes.

PETER ACHTERSTRAAT: Sometimes we have to make hard decisions in relation to pilots and sandboxes to stop it.

The Hon. JACQUI MUNRO: Thank you. Dr Costa, did you have anything?

MATTHEW COSTA: I would just add I'm not able to nominate particular areas; however, I suppose we're in favour of the idea, in general, and particular regulators in the New South Wales Government will be considering and thinking about those opportunities. The question is well directed to them, I think.

The Hon. JACQUI MUNRO: Who are they?

MATTHEW COSTA: It's going to depend on the area in question. Yes, the whole—

The Hon. JACQUI MUNRO: Are there any regulators that are currently talking about that possibility?

PETER ACHTERSTRAAT: Forgive me, Doctor. It may well be that industry New South Wales people might approach them and say, "We've got this great idea. We want to experiment." I would encourage those sorts of departments to listen to those people, and we'd gladly assist if people do approach them. But I guess at this stage, Ms Munro, it's more demand driven—they come to us—rather than we say to you, "We want you to do it in that area."

The Hon. JACQUI MUNRO: So it's about ensuring that departments are aware of the possibility of doing something like this so that it's not just, "No, that's not possible," it's, "Okay. Let's start this conversation"?

PETER ACHTERSTRAAT: Correct. You've got me thinking. Maybe on the proactive side, if departments can see something working in other jurisdictions, then that might be an opportunity to talk to the

industry association for that group, to say, "Look, this works well in Ontario, or wherever. What about we try it here?" So that's actually a good thought.

The Hon. JACQUI MUNRO: Just one more question about getting buy-in from the public and how we help people feel comfortable with new technologies. I guess it's kind of general, but in this specific way are there examples that you are aware of from other countries where governments have done a good job of helping people understand AI, its risks and opportunities? Where are they? How did they do it?

PETER ACHTERSTRAAT: Can I answer that in the general rather than the specific? In the general, I've found with productivity improvements and asking whether it be new technology or different regulation or whatever—from my experience, back in the old days, you could sell it to the community. Even if 90 per cent of the people in the short term were winners and 10 per cent were losers, you were able to explain that in the long run everyone would be better off, even though in the short term, only 90 per cent would be. More recently, I have got the feeling—there's no evidence for this—with social media and things like that, that we tend to need now about 99 per cent of people to be short-term winners, which makes it more challenging for us to be able to sell new ideas. So what we've got to do is focus on the long-term benefit.

What I've noticed in my experience is that if we give people a solution without telling them why, they're going to resist it. If we say, "You've got to do this artificial intelligence," what we have to do is present the case. The case at the moment is that the amount of goods and services produced per person each year is not increasing, and if you want real wage rises, we need to do something to improve that, and artificial intelligence may be one of those ways. I think we've got to keep sending the message about why we need productivity levers, and then we've sown the seed, then we can be a bit more specific and talk about particular types of productivity levers, whether it's artificial intelligence or whatever. But you're all absolutely right. If people start thinking artificial intelligence means losing my job, it's like people who think productivity means cost cutting. Productivity doesn't mean cost cutting. It might mean increasing the cost, but if you can increase the output you double that, so it's a matter of the narrative.

The Hon. JACQUI MUNRO: Yes. Just a little note that I know that the Taiwanese Government have been very proactive in showing citizens, for example, what a fake video might look like or how a video might replicate somebody speaking about a particular issue and ensuring that people are protected in some way through knowledge of the possibilities of AI.

PETER ACHTERSTRAAT: I can't comment on other countries on that, but the point—I think one of the five principles of the assurance map is that privacy and security is important, not just for itself but for community confidence. When the Excel spreadsheets came in, if I didn't trust that it was going to add up properly, I'd probably still do it all myself. We need to have a governance arrangement around it so that people will have confidence that the data coming out is accurate and that, if there is any fictitious stuff, it's weeded out either by regulation or by the market itself.

The Hon. CAMERON MURPHY: Thank you very much for your submission. I think it's really interesting and enlightening. I just wanted to put to you an issue that came up through some of the evidence from earlier witnesses. In your submission, it's really great that you identify the four groups of emerging technologies. One of those stands out to me, which is reinforced AI. And we just heard earlier from some of the creative industry groups that they're saying that most of the generative AI is based on simply uploading or stealing all of their creative content. Do you see that there are any risks associated with adopting that sort of technology wholesale, without dealing with issues about compensating copyright owners and putting in place regulation or protection for moral rights over the use of that material?

PETER ACHTERSTRAAT: I'll answer generally, Mr Murphy, and thanks for that question. I'll get Dr Costa to go on further. We do need to be in a situation where we encourage innovation, encourage R&D. Just like in relation to developing pharmaceuticals, where someone spends a lot of time developing a pharmaceutical, they should be entitled for some time to be able to use that themselves, rather than other people use it. But there will come a time when it might be in the general population. Similarly with other copyright issues and things like that, where reinforced AI—I don't have the complete answer. The concept I'll hand over to Dr Costa, but we need to make sure that there's not a disincentive—for people not to produce content because they think, "It's going to be taken by someone else, disseminated, and I won't get a return on it." I don't know what the regulatory answer is, but I think you've hit the nail on the head. I think we've got to be most careful that we can encourage that innovation and content but, at the same time, disperse it in a regulated manner. Dr Costa, have you got a thought on reinforced AI?

MATTHEW COSTA: I think it's a valid point, and I suppose a general point is that every new technology seems to pose challenges to existing regulatory frameworks. I suppose there are two issues we call out in relation to that. One is having outcomes-focused regulation—really asking what is the spirit of the regulatory

outcome you're trying to achieve. In this case, it's fairness and protection of intellectual property and incentives for development of creative content and so on. Those are things to be considered in relation to the new technology. Another point we make is fairness. If new technologies are not seen as producing fair outcomes, that can undermine community support for those technologies, so it's a really important thing to deal with, the issue of fairness, for increasing confidence and promoting broad support for the use of these technologies.

The Hon. CAMERON MURPHY: How do you deal with something like moral rights over a copyrighted work, where it's not simply a matter of money but an LLM is just hoovering up written material, copyrighted material, and then is using that to generate an output but an artist says, "I don't want to be associated with that. I don't want my work to appear as the output, because I don't like the group that is going to utilise that work", for example? How do we manage things like that, in the adoption of this technology?

PETER ACHTERSTRAAT: Mr Murphy, I don't know the answer to that one. I'm terribly sorry. It's a wicked problem. It wasn't directly part of our research. We might take that on notice. No, we won't take it on notice. It's just too big an issue.

MATTHEW COSTA: I agree with that, but these are general-purpose technologies, which means that they permeate and can be applied across every sector and area of the economy, which really means that, to respond to them appropriately, we're going to need to marshal expertise from all over the place. I think you're talking about issues to do with intellectual property law, and intellectual property experts are the kinds of people that you want to be working through these challenges with. They're certainly—

The Hon. CAMERON MURPHY: Yes. But have you identified in your productivity work that that's a significant risk—that you may end up with disputes over the use of this material or you may end up later on with litigation that results in a cost to people that use this technology without paying for it and so on? Is that something you took into account in your modelling?

PETER ACHTERSTRAAT: Mr Murphy, we did not cover it directly in our report. But, just listening to you now, I can see it is a huge issue. It's not one we covered, but it's clearly one which needs to be addressed before we go too much further down the path of AI because it will—if it's going to stifle innovation, if people are not going to write songs or whatever, then we're going to go backwards. So I think—

The Hon. CAMERON MURPHY: Yes. I worry we might go five years ahead and then the lawsuits will start to catch up with that innovation, perhaps.

PETER ACHTERSTRAAT: Absolutely. If people are already putting things out there and they answer questions and are dragging information from elsewhere and then they find out they had no right to get that information, you're right; there could be a Pandora's box of—could be a sleeping log there of lots of litigation. It's an issue which must be looked at. It's not much of an answer, I'm afraid.

The Hon. JACQUI MUNRO: I just have one question around the interaction between this work and something like the Intergenerational Report and how they feed into each other, how we are, potentially, affecting intergenerational equity by not engaging in this regulatory space with AI. I guess it's similar to what Chris was asking before.

PETER ACHTERSTRAAT: Thank you, Ms Munro. If I can maybe paraphrase your question, if we don't seize an opportunity now to use this, our children and grandchildren won't have the same standard of living that we have. And, if we don't pull some productivity levers over the next few years, that may well be the case. This is a lever that's standing out there for us to pull. But, as all members of the Committee have made clear, we've got to make sure that we pull that lever the correct way and we don't cause problems by doing so. But you're absolutely right. If we just stand there looking at the lever and don't do anything, then the productivity will go backwards, and then our children and our grandchildren will have a lower standard of living than we have.

The Hon. JACQUI MUNRO: Are there any other productivity levers that you see at the moment, that are so clear and obvious?

PETER ACHTERSTRAAT: It depends on the eye of the beholder, what clear and obvious is. Some people will say if we can enhance the training of our population, that will help. Some will say if we can change some of the planning regulations, that will help. Others say if we can do other things in relation to competition, that will help. But each of them, I think—the Committee made the point earlier that we do have to bring the community on side with us. When we say something is obvious—this is clearly obvious to me—it's a matter of those who think that to be able to sell that and market that to not just the population but also to businesses that may not have been as quick adopting technologies as some of our peers in other countries have been.

The Hon. JACQUI MUNRO: So it goes to competition as well.

The Hon. Dr SARAH KAINE: In your submission you talk about the impacts on social inclusion and associated issues. The one that I'm most interested in is your comments on ensuring appropriate protections exist for workers. One of the things that I think you spoke about earlier was the change in jobs and that new jobs are created and that not all jobs are created equal—even if they're created equal in number, not necessarily in quality. You do talk about—and I think you're quoting your Adaptive NSW document—needing to strike a careful balance between ensuring adequate worker protections and allowing essentially what you've been talking about: not stifling innovation. I'm pleased to see that included but it is a very normative—it's quite a judgement, isn't it? What is the careful balance? How do you strike that balance? What's the balance? If you asked probably a couple of people on this Committee, our idea of what's balanced might be quite different from each other. I wondered if you had any further comments to make on that section of your submission?

PETER ACHTERSTRAAT: Thanks very much, Dr Kaine. When the expression "worker protection" is used by people with new AI, I like to use the expression "worker opportunity". If someone is in an industry, say, the old typewriter mechanics—people who used to repair typewriters; we can't just keep them doing that. We've got to retrain them. Worker protection I see as enhancing them with a skill set to be able to either work in an enhanced version in that industry or even move to other industries. So I think the key and the flavour—what we've tried to bring out through the whole of our Adaptive NSW report—is to start preparing people whose jobs may change so significantly that they will need to be retrained. So I guess it is worker protection, but protection in that they'll have a better job, hopefully.

The Hon. Dr SARAH KAINE: Could I ask for a bit more on that? I think that's a wonderful aim but is the thought—and the basis of what you're talking about—that the market will provide that because it needs those workers? I'm a little worried about that because what we've seen so far in some of these disruptive areas is actually a movement from more secure jobs to less secure jobs. Are your assumptions that the market would be providing that move towards those protective jobs, or that there would need to be some intervention in that market?

PETER ACHTERSTRAAT: Thank you very much, Dr Kaine. In fact, that's a very good point. There will undoubtedly possibly be market failure where the market won't necessarily say to these particular people "We're going to retrain you", et cetera, so there may well be a role for government then to be able, either through the VET system or through a training system, to support those people. It's certainly not exclusively the market to do it; I think the government will have a role, or governments all over the world will have a role, to create some sort of a safety net. That benefits everybody: It benefits the individual worker who has to change their role, and also benefits society by enhancing productivity so that these people, where possible, will be able to have a job or a role which is more productive. They may not be. They may be people who have to be supported by governments in another way.

The CHAIR: Mr Achterstraat, I've got a question at the end of a long week and on a Friday afternoon: Is there opportunity for AI—for us—to move to a four-day working week?

The Hon. Dr SARAH KAINE: Great question.

The CHAIR: Yes, it is a good question. We were provided with a report from the CSIRO that said with AI, one of the great productivity gains we could get in terms of work-life balance—and in fact, the Swinburne University study found that company managers who took part in a four-day-per-week trial gave an overall success score of 9.25 out of 10. Is AI the opportunity for us to get productivity up by resetting our work-life balance with potentially a four-day working week?

PETER ACHTERSTRAAT: I think we'd have to make sure the horse is before the cart. Before we would jump to a four-day week, we'd have to decide whether AI has been able to generate the productivity to be able to do that. What I can say is that there is documentation which suggests that the average worker in the USA produces in four days what the average worker in Australia produces in five days. There could be various reasons for that: It could be better training; it could be more capital enhancement, as Dr Kaine said; it could be AI; it could be all sorts of reasons; or it could be different regulation. I haven't done any analysis to say whether we should switch to a four-day week. Going from six days to four, is that what we're saying? Or seven days to four?

The CHAIR: Seven-day working week?

PETER ACHTERSTRAAT: Forgive me; it's Friday afternoon.

The CHAIR: From five to four, I think we're going.

PETER ACHTERSTRAAT: To answer your question, we haven't done any analysis on that but if we did want to reduce the working week, we would have to have tools such as the AI. I think it was John Maynard Keynes that said once that if productivity enhancements in technology keep going, we'd end up with a 15-hour

week. I think that was what John Maynard Keynes said. It didn't quite happen yet but I think if we don't have AI, it's less likely to happen.

The CHAIR: Thank you very much for your fantastic work generally—we really appreciate it—and also for the wonderful submission. It's going to inform our deliberations greatly. Thank you for taking the time today to come and give evidence to us. It has greatly enhanced the work of the Committee.

(The witnesses withdrew.)

The Committee adjourned at 16:45.