COMMONWEALTH OF AUSTRALIA

Proof Committee Hansard

JOINT SELECT COMMITTEE ON ARTIFICIAL INTELLIGENCE

(Public)

FRIDAY, 16 OCTOBER 2026

CANBERRA

CONDITIONS OF DISTRIBUTION

This is an uncorrected proof of evidence taken before the committee.
It is made available under the condition that it is recognised as such.

BY AUTHORITY OF THE SENATE

[PROOF COPY]

JOINT SELECT COMMITTEE ON ARTIFICIAL INTELLIGENCE

Friday, 16 October 2026

Members in attendance: Senators Duniam and Sheldon and Ms Briskey, Mr Caldwell, Ms Chaney, Ms Julie-Ann Campbell, Mr Gregg, Mr Husic, Ms Jarrett, Mr Kennedy, Mr Laxale, Ms Mascarenhas and Ms Spender [by video link]

Terms of Reference for the Inquiry:

(1) A Joint Select Committee on Artificial Intelligence be appointed to inquire into and report on:

(a) the opportunities presented by artificial intelligence (AI) to lift Australia's resilience, productivity growth, economic competitiveness, and living standards;

(b) the opportunities for the benefits of AI to be spread across Australia's suburbs and regions;

(c) the opportunities of AI and robotics to drive research and innovation, new industries and business, and export opportunities, including in sectors of Australian comparative advantage such as agriculture, resources, health and medical research, advanced manufacturing, financial services and defence industry;

(d) the rate and extent of AI adoption across the Australian economy, and the barriers to adoption, faced by small and medium businesses and family businesses;

(e) the opportunities for Australian sovereign AI capability, including the measures required to build, retain and commercialise domestic capability, grow research and development, and attract and retain investment and talent;

(f) the importance of data sovereignty, and the opportunities presented by narrow, domain-specific AI models trained on Australian data in areas of Australian sectoral expertise;

(g) the workforce and skills implications of AI, including the adequacy of Australia's education system, job creation and displacement, and the effects on the quality and security of jobs;

(h) the adoption of AI by Commonwealth departments and agencies to improve services for Australians, and the transparency and accountability measures that apply to AI use;

(i) the adequacy of Australia's existing laws and regulatory frameworks as they apply to AI and whether there are any gaps that warrant reform;

(j) the interaction of AI with existing intellectual property and copyright laws, including the use of Australian creative, cultural and media content in the training of AI models;

(k) the risks and harms arising from AI, including fraud and scams, deepfakes and risks to the safety and wellbeing of children and vulnerable Australians;

(l) the national security and cyber security opportunities to protect Australians, as well as the foreign interference risks, associated with AI, including the security and resilience of supply chains and of systems supporting critical infrastructure;

(m) the implications of emerging AI capability for Australia's national security and strategic resilience, including the ability of regulators, the Australian AI Safety Institute and the intelligence and security community to identify and respond to emerging risks; and

(n) any other related matters;

(2) the committee present its final report by no later than **30 November 2026**.

WITNESSES

ABKIEWICZ, Mr David, Acting General Manager, Digital Markets Branch, Australian Competition and Consumer Commission 67

BASSI, Mr Justin, Executive Director, Australian Strategic Policy Institute 52

BENNETT, Professor Jill, Fellow, Australian Academy of the Humanities [by video link] 9

BOLAND, Ms Marie, Chief Executive Officer, Safe Work Australia [by video link] 61

BROWNE, Mr Bill, Director, Democracy and Accountability Program, Australia Institute 52

BUNFIELD, Ms Megan, Chief Information, Data and Digital Officer, Digital and Technology Division, Australian Competition and Consumer Commission 67

COOK, Mr Peter, Business Member, Digital Economy Council of Australia 22

CORNICK, Ms Kate, Chief Executive Officer, Tech Council of Australia [by video link] 22

DAVIS, Ms Inga, Executive Director, Australian Academy of the Humanities [by video link] 9

DENNISS, Dr Richard, Co-Chief Executive Officer, Australia Institute 52

DODDS, Emeritus Professor Susan, Policy Lead, Australian Academy of the Humanities [by video link] 9

EGAN, Ms Janet, Director, AI Policy, Institute for Progress [by video link] 1

FOSTER, Ms Nicole, Director, Global Artificial Intelligence Policy, Amazon Web Services [by video link] 14

GOODEY, Ms Amy-Rose, Chief Executive Officer, Digital Economy Council of Australia 22

HAWKINS, Ms Zoe, Co-Founder and Deputy Executive Director, Tech Policy Design Institute 36

KEAY, Dr Sue, Deputy Director, Kingston Artificial Intelligence Group [by video link] 45

LONGCROFT, Mrs Lucinda, Director, Policy and Government Affairs, Tech Council of Australia 22

MACDONALD, Dr Fiona, Policy Director, Industrial and Social, Centre for Future Work [by video link] 61

MONFORT, Mr Mark, Co-Chair, AI Working Group, Digital Economy Council of Australia 22

O'ROURKE, Mr Matthew, Head, Infrastructure Policy, Australia and New Zealand, Amazon Web Services 14

O'SULLIVAN, Mr Alec, Head, Operations, Digital Economy Council of Australia 22

QUINLAN, Emeritus Professor Michael, Private capacity [by audio link] 31

RICHMOND, Ms Jayde, General Manager, National Anti-Scam Centre, Australian Competition and Consumer Commission 67

ROSS, Mr Rick, Business Member, Digital Economy Council of Australia 22

STANIC, Mrs Rada, Chief Technologist, Australia and New Zealand, Amazon Web Services [by video link] 14

TEASDALE, Dr Hayley, Head, Science Advice and Policy, Australian Academy of Science 45

THOMAS, Professor Julian, Fellow, Australian Academy of the Humanities [by video link] 9

WEAVER, Ms Johanna, Co-Founder and Executive Director, Tech Policy Design Institute 36

WILLIAMS, Mr Timothy (Tim), Business Member, Digital Economy Council of Australia 22

WILLIAMS, Ms Sally-Ann, Fellow, Australian Academy of Science 45

WILLIAMS, Professor Penny, Professor, Queensland University of Technology [by video link] 31

WITNESSES

WINTER-LEVY, Dr Sam, Senior Fellow, ...
Technology and International Affairs Program, ...
Carnegie Endowment for International Peace [by video link].....1

WROE, Mr David, Head, AI and Security, Australian Strategic Policy Institute.....52

EGAN, Ms Janet, Director, AI Policy, Institute for Progress [by video link]

WINTER-LEVY, Dr Sam, Senior Fellow, Technology and International Affairs Program, Carnegie Endowment for International Peace [by video link]

Committee met at 08:31

CHAIR (Ms Briskey): I declare open this hearing of the Joint Select Committee on Artificial Intelligence. I begin by acknowledging the traditional custodians of the land on which we meet and pay my respects to elders past and present. I extend that respect to any First Nations peoples with us today. These are public hearings being videostreamed live via the parliament's website, and a *Hansard* transcript is being made.

I remind all witnesses that, in giving evidence to the committee, they are protected by parliamentary privilege. It is unlawful for anyone to threaten or disadvantage a witness on the account of evidence being given to the committee. Such action may be treated by the Senate as contempt. It is also a contempt to give false or misleading evidence. Witnesses also have the right to request to be heard in camera.

The Senate has resolved that an officer of a department of the Commonwealth or of a state shall not be asked to give opinions on policy and shall be given reasonable opportunity to refer questions asked of the officer to superior officers or to a minister. This resolution does not preclude questions asking for explanations of policies or factual questions about when and how policies are adopted.

If a witness objects to answering a question, they should state the ground upon which that objection is made, and the committee will determine whether it will insist on an answer, having regard to the ground on which it is claimed. If the committee determines to insist on an answer, a witness may request that the answer be given in camera. Commonwealth officers appearing today are also reminded of the Senate order specifying the process by which a claim of public interest immunity should be raised. A copy of the order is available from the secretariat.

The extract read as follows—

Public interest immunity claims

That the Senate—

(a) notes that ministers and officers have continued to refuse to provide information to Senate committees without properly raising claims of public interest immunity as required by past resolutions of the Senate;

(b) reaffirms the principles of past resolutions of the Senate by this order, to provide ministers and officers with guidance as to the proper process for raising public interest immunity claims and to consolidate those past resolutions of the Senate;

(c) orders that the following operate as an order of continuing effect:

(1) If:

(a) a Senate committee, or a senator in the course of proceedings of a committee, requests information or a document from a Commonwealth department or agency; and

(b) an officer of the department or agency to whom the request is directed believes that it may not be in the public interest to disclose the information or document to the committee, the officer shall state to the committee the ground on which the officer believes that it may not be in the public interest to disclose the information or document to the committee, and specify the harm to the public interest that could result from the disclosure of the information or document.

(2) If, after receiving the officer's statement under paragraph (1), the committee or the senator requests the officer to refer the question of the disclosure of the information or document to a responsible minister, the officer shall refer that question to the minister.

(3) If a minister, on a reference by an officer under paragraph (2), concludes that it would not be in the public interest to disclose the information or document to the committee, the minister shall provide to the committee a statement of the ground for that conclusion, specifying the harm to the public interest that could result from the disclosure of the information or document.

(4) A minister, in a statement under paragraph (3), shall indicate whether the harm to the public interest that could result from the disclosure of the information or document to the committee could result only from the publication of the information or document by the committee, or could result, equally or in part, from the disclosure of the information or document to the committee as in camera evidence.

(5) If, after considering a statement by a minister provided under paragraph (3), the committee concludes that the statement does not sufficiently justify the withholding of the information or document from the committee, the committee shall report the matter to the Senate.

(6) A decision by a committee not to report a matter to the Senate under paragraph (5) does not prevent a senator from raising the matter in the Senate in accordance with other procedures of the Senate.

(7) A statement that information or a document is not published, or is confidential, or consists of advice to, or internal deliberations of, government, in the absence of specification of the harm to the public interest that could result from the disclosure of the information or document, is not a statement that meets the requirements of paragraph (1) or (4).

(8) If a minister concludes that a statement under paragraph (3) should more appropriately be made by the head of an agency, by reason of the independence of that agency from ministerial direction or control, the minister shall inform the committee of that conclusion and the reason for that conclusion, and shall refer the matter to the head of the agency, who shall then be required to provide a statement in accordance with paragraph (3).

(d) requires the Procedure Committee to review the operation of this order and report to the Senate by 20 August 2009.

(13 May 2009 J.1941)

(Extract, Senate Standing Orders)

CHAIR: I remind committee members of their obligations under the Behaviour Code for Australian Parliamentarians to treat witnesses with dignity, courtesy, fairness and respect. Welcome. I understand that information on parliamentary privilege and the protection of witnesses giving evidence to this inquiry has been provided for you. I would now like to invite you to make a brief opening statement, and, at the conclusion of remarks, we'll invite committee members to start some questions. Ms Egan.

Ms Egan: Thank you so much for inviting me to speak to you today. I'd like to highlight that the pace of AI has been astronomical. In the last four years we've gone from chatbots that couldn't do basic maths to them solving some of the world's hardest maths problems that have stumped our mathematicians for decades. This progress doesn't show any sign of stopping. AI companies have a number of proven recipes that continue to make AI better and better. The speed of progress is actually accelerating as AI contributes to its own R&D, in the early stages of recursive self-improvement, and we've got AI systems that are now better than expert cyber defenders at finding critical vulnerabilities, yet these same capabilities can be used to exploit them at scale.

The benefits of this technology are tremendous. We should expect new cures for diseases, breakthroughs in material science and innovations that make everyone's lives better, and we can and we should be really excited about this potential. But to realise these benefits, we also have to manage the risks, and some of these are already here. This month we've seen South Korean banks targeted in cyberattacks by an individual using a combination of AI tools, and this person doesn't seem to be a sophisticated actor. Based on public reporting, it looks like a 26-year-old student who is asking the same AI models to help write his resume.

We should expect not just today's risks but more across other domains, from malicious actors creating novel bioweapons to loss of control over systems, and, potentially, workforce destruction we haven't seen before at a scale of this kind. These are genuine, complex problems that the world needs to grapple with, but these risks also won't respect national borders, and AI will affect Australians wherever it is built.

The question facing Australia today is whether it acts to capture a significant part of the AI supply chain and uses that position to shape how AI is developed and governed to make sure Australians share in its benefits, or whether Australia is willing to accept an AI future determined by others.

Australia has been dealt a hand that other countries envy; it has land, renewable energy potential and geopolitical stability and relationships to host AI training at scale. This is significant. Inference compute—the data centre used to deploy AI models—can be assembled incrementally in many places that are largely fungible and more sensitive to cost, and modelling from Sam Winter-Levy shows that Australia is not a leader based on cost alone. Its training, on the other hand, which benefits from vast, multi-block clusters wired together that host tens or hundreds of thousands of chips, is the real opportunity for Australia to have influence over AI's future.

As you've probably heard from other witnesses, Australia's current copyright regime is the blocker to substantial training compute being built in Australia, and the status quo is, unfortunately, serving no-one. We have Australians creatives whose data is being trained on overseas in the US with no compensation, no control and no ability to opt out. Australia foregoes the investment in and influence through training centres hosted on Australian soil that could provide Australia with meaningful influence in the AI future.

Thank you for engaging on these difficult issues. AI, with its implications and uncertainties, is very challenging work, but it's so important, and I'm very grateful that Australia's leaders are engaging deeply with what AI means for Australia, its citizens and the hard trade-offs involved. Where this infrastructure is located matters for two big reasons. The first is economic. Countries that host some of these data centres can capture direct economic value with construction spending, operational jobs and tax revenue, which may become quite important to public finances if AI reshapes economic activity. The second reason is leverage. Hosting compute doesn't automatically give a country control over the models trained or run on its soil, but it's easier to regulate, inspect and negotiate with an entity that has sunk capital into your jurisdiction than with one serving you remotely from another continent. For countries that rely entirely on remote providers, they face the possibility that access could be

restricted, whether because providers redirect their limited computing capacity toward other priorities or because of export control sanctions or changes in corporate policy. If countries like Australia can unlock large-scale compute sites, they can use that to bargain for frontier access with the US labs. In exchange for sites, they can seek access on par with the US so frontier models are released in the Australian market at the same time as in the US and there is equivalent access for trusted firms and critical infrastructure operators. So onshore compute does not guarantee sovereignty over AI, but it makes it much more likely that a country can influence how AI develops.

Which countries can capture these benefits depends above all on time to power. We did extensive modelling on this. Policy debates often focus on energy costs and tax incentives. Our research suggests that those are secondary. For a typical 100-megawatt data centre in the US, companies should be willing to pay at least double US power prices to operate a data centre just one year sooner, and it's similar in Australia as well. By our estimates, at the moment, Australia moves significantly slower than the United States at building data centres, which makes a typical data centre in Australia worth about 15 per cent less over its life cycle than one built in the US. The main source of delay here is power procurement—above all, the wait for grid connection.

To close, I echo Ms Egan in saying that Australia has a significant opportunity here. It's very well positioned to host compute. It already has a deep security relationship with us. It has ample land and solar energy. Relative to some of the other places, like the Gulf states, that are trying to capture a share of the build out, the security risks are low, but Australia does need to move with urgency, and that requires removing some of the obstacles that slow projects down domestically, like permitting backlogs, grid connection queues and transmission bottlenecks. Governments should not spend public funds on this industry through energy subsidies or tax cuts, but they should consider creating fast-track review processes for projects that meet certain criteria such as use of clean energy, investment in local communities or [inaudible] commitments. As long as governments can actually shorten the approval timelines, they have a lot of flexibility to shape the bargains they strike with data centre developers. To be clear: speed requires a social licence to operate, which depends on transparent deals and visible, guaranteed benefits to those communities.

I close by saying that no single country, not even the US, can build the world's AI infrastructure alone. Coordination among allies can help, and countries that contribute real value to the build-out through faster permitting, abundant land and clean power will have the most leverage to shape the future of AI. Thank you very much. I look forward to your questions.

CHAIR: Thank you very much to both of you for those. I will kick off. Ms Egan, you reflected on this in your opening remarks, and I wonder if you could expand on it a little bit further. We are at a critical juncture as a country in terms of the opportunity in front of us where we have the frontier labs seeking to expand their compute and their training, for the first time, outside the US to Australia, in particular with hyperscalers. In your view, what will that mean not just to, as you said, our ability to help shape this technology and how it's governed but also for Australians in terms of our sovereign capability? In your view, what is the opportunity here and what is the opportunity cost if we cannot seek to resolve some of the issues—in particular, as you highlight, copyright is a pretty significant one at the moment. Can you expand a little bit on some of the opening remarks and, in particular, Australians' opportunity here around our own sovereign capability for this incredible technology?

Ms Egan: Yes, I'm very happy to talk about that. Australia is in this unique position where there is interest from the leading frontier AI companies to build training infrastructure in Australia. If Australia is able to capture this, I think there's a range of different benefits. But the first thing I'll caveat by saying that the future of AI is highly uncertain. If anyone is telling you that there's a 100 per cent likelihood of X, I think they're overconfident. When I think about how to look after Australia's interests into this fast-transforming AI future, I think about how you preserve option value and optionality for Australia to protect its interests and its citizens. Maybe one of the first things to tackle there is workforce disruption. There's a range of different economic arguments about how quickly there'll be workforce disruption, which kinds of jobs will be automated and how quickly new jobs emerge, but, at the end of the day, I think it's fair to say that governments haven't traditionally had a very strong playbook on how to manage workforce disruption well and there is a risk that a lot of the service economy, which Australia is quite strong in, is disrupted and a lot of the value flows into just a handful of companies in just two countries: the US and China. I think that, if Australia builds out large-scale AI infrastructure, they will have this resource into the future that gives them some influence over and leverage over these companies that are building in terms of both future taxes and how you might distribute wealth if there is large-scale automated disruption and also understanding what's coming around the corner.

The other aspect of sovereignty that I think is enhanced by hosting large-scale training compute in Australia is the ability to just get much deeper partnerships with leading US AI companies. These companies are still very

lean. The technical talent there is still quite thin. They're growing rapidly as companies, but the key experts in AI and cyber or AI bio or deep interpretability research are still just a handful of people. These people are getting outreach from every government around the world and every AI safety institute or equivalent around the world. I worry about a situation where Australia's calls for deep engagement on what's emerging at the frontier are drowned out amongst the chorus of other nations also approaching these companies and saying: 'We need more information. We need more deep collaboration on these risks.'

I think one of the ways to clearly incentivise these companies to put Australia at the top of that list is by hosting a large part of these companies' infrastructure. At the moment, I think our AI Safety Institute is largely reliant on the goodwill of AI companies, and they should be engaging, but, at the same time, who you get access to and to the level of depth you get in technical capability can shift rapidly if they choose another country to be their frontier AI compute build-out destination.

I also think we're going to see more and more dual-use capabilities emerging, as I touched on in the opening statement. The question of whether Australian companies and the Australian government can access those capabilities, particularly ASD with regard to cyber, the Australian Cyber Security Centre with regard to hardening our critical infrastructure and working with industry to do that and ensuring we have access to the most frontier AI models—we have much greater assurance of being able to negotiate that if AI companies are reliant on our jurisdiction to develop these capabilities.

To me, these are some of the key things, like the ability to redistribute the wealth and support the Australian economy through periods of transition, to have insights into the technical frontier and where things are happening and also to be able to set rules and guardrails that actually inform us AI companies' behaviour as well. Obviously, this would be working closely in an international space too, because, at the moment, Australia could introduce the best rules in the world—the most thoughtfully written out, researched regulations—and they would probably have zero impact on AI companies' approaches to how they're carrying out their business.

CHAIR: As you reflected on in your opening statement, one of the big areas of tension that we are grappling with is how we make sure we resolve the copyright issue. We heard evidence in our session last week that there are copyright holders that vehemently argue that the current status quo is eminently workable for these big technology companies. They've got the capacity to licence, to create and to make sure that copyright holders have control over their artwork. But we also had other creators who said—I think of the comments that you made in your statement—they don't feel at all like they have control over their own IP and that they are seeing nothing in terms of their ability to seek compensation for it. Where do you think some of the tensions are from your end? How might we seek to resolve them, given that there is a real focus here on how we make sure we genuinely protect Australian creativity and creatives right across the board, whilst not potentially hindering the significant opportunity that we have for Australians to secure training here?

Ms Egan: The harsh reality for Australian creators is that, right now, companies can train on their content in the US under US law and court cases to date have all found in favour of the fair use doctrine applying. We've also got other countries that have permissive copyright laws that allow for AI training. I think Japan has been raised as one with me. I'm not a legal expert; this has been in my conversations with lawyers. At the moment there is a way for AI companies to train without paying for licences for training data, and, unfortunately, they aren't obliged to allow creatives in Australia to opt their content out.

I think the issue here is that AI companies seem to have the very strong view that, if they pay for AI licences for training in Australia, that could undermine fair use doctrine in the US. I'm not making a moral statement about companies' behaviour in this space. I'm making a pragmatic argument that, at the moment, they have options to train in jurisdictions that don't require them to pay licences, and they have the legal advice that, if they pay licences for training data, they could create the market for that data which could actually—because, under fair use doctrine in the US, it needs to be transformative use but there also needs to not be an existing market of paying licences for data for that purpose. I understand that they have legal advice that says, 'Don't take the risk.' It's less of a: 'Here are lots of pros for Australia, and here's a con. Maybe the pros will come out on top.' My understanding is that AI companies have taken a strong line: 'This is a red line for us.'

From a pragmatic perspective, looking forward, Australia faces this choice on what we can do to unlock this investment, and also have future influence on a technology that is already impacting creatives but will impact far more than creatives. But it's also: how do we sustain and support the robust and rich creative industry of Australia into the future? Australia has a world-leading creative industry, and we should be incredibly proud of that. I worry as to copyright licensing; I can't see this being effective in the long term if AI companies continue to train in other jurisdictions. The thing is, though, that companies are licence sensitive but not price sensitive. I've seen proposals from an organisation called Good Ancestors that outlines how you could just charge a fee if it goes into

government before being passed out to creatives. It's not called copyright. We could get an ongoing stream of revenue on an annual basis that supports a creative industry in the longer term, potentially giving creatives more money and more access to opportunities to create great works than we have under existing government support.

I also think we want to have something like this going forward, because one of the risks here is that the creative industry has had the protections of copyright to date. But I think of software engineers whose legacy of contributions to the open-source software repositories of the world, through free time and effort and their professional contributions—all that data has been sucked up into these models and has made these models excellent at software engineering. Now the reality is that we're seeing a softening of the market for employment for junior software engineers, but there is no channel to provide support to these sectors as they are impacted by increasing automation of aspects of the economy. I'm optimistic there'll be new jobs and new ways of contributing, but I worry that there's a lot of knowledge that has been incorporated into these AI models that don't have a way to reimburse the people who've created that. I'd like to see a solution that supports both those in the creative industry—not just of yesterday but future creatives as well—and other industries that may be impacted by AI disruption as they find new ways to contribute to society and new economically valuable contributions to make.

CHAIR: Thank you. Deputy Chair.

Senator DUNIAM: Ms Egan, I'm picking up where you left off with the questions from Chair, with regard to creatives and copyright. That is that intractable issue that we are grappling with at the moment, and I think the committee has become a bit of a lightning rod for both sides of the debate here. I personally am not an all-or-nothing kind of guy on this issue. I think there must be a way forward, very much to the point that you've made there. What I'm interested in, though, is the nuts and bolts of how we might resolve this issue. There are entities and individuals who are advocating for no change because the system works for them and they believe that the companies and the AI labs will be able to enter into arrangements, even if it is a protracted process, in order to comply with Australian copyright law. And you've talked about what that might mean for their capacity to operate in overseas jurisdictions. You talked about a fund. I guess I'm trying to find any pointer to what we might look at doing here to deal with the entities and individuals when it comes to the rights holder cohort who would be impacted by AI. Do you have any views or models or examples in other jurisdictions that we in Australia could look to that will create least harm for anyone in the rights holder cohort?

Ms Egan: I think that I wouldn't necessarily recommend models in other jurisdictions, like Japan's text and data mining exemption, because Australia has a good hand to play, so it can actually ask for more than just giving out a carte blanche for AI training. I actually am pretty enthusiastic about options that increase optionality to both parties. You could keep an option for existing copyright law to still apply and then introduce a separate channel that companies could choose to go through that allows an opt-out for creatives, which would give them more control than we currently see under the status quo.

There's still an option, should companies be willing to reach a copyright licence agreement and should creatives say: 'No, I'm retracting my data. I'm opting out of this separate channel. If you want my data, you have to engage with me through the traditional copyright process.' You could attach terms and conditions on this streamlined process where, if you want to have a way through that isn't copyright, you contribute to a fund and you meet data centre expectations, which I think the government's already working on. This just increases the options available to parties and allows them to find agreements that work for them. So every creative in Australia can still say, 'I would like to opt out. I'm setting this up so that you aren't going to use my data absent a copyright licence,' but you could also then have those who wish to benefit from this fund allowing their works to be trained on.

Senator DUNIAM: You made the point before about creatives' works being used for training overseas. Do you have examples of that that you can point to? It's something I've referred to in a nebulous sense, not being a holder of examples myself. Do you have specific examples you can highlight to the committee?

Ms Egan: I don't have examples of an Australian work that has been used. But, for example, models trained in the US have vast amounts of data from the internet and some of that content, while it's fair use in the US, would not be covered under that if it were trained in Australia. Another example is that, if you have legally procured books or a vast library of texts, in the US you can train on that under fair use. In the Australian context, you wouldn't be able to do that. I'd be fairly confident that there'd be at least one Australian work in that dataset.

Senator DUNIAM: Yes. Dr Winter-Levy, you talked about the speed with which our jurisdiction needs to move in order to take advantage of the opportunity. I want to go to your point around approval timelines. Obviously, under our federated system, there's only so much a Commonwealth government can do when it comes to planning and approvals. Do you have any guidance on what Australia might do to deal with the realities of

planning and environmental law at a state and territory level and what we might recommend to the Commonwealth through this committee?

Dr Winter-Levy: In the United States, one of the reasons the build-out has been able to unfold so quickly is that different states have chosen to strike different types of bargains with the companies. Some states, like California, have chosen to move much more slowly with much more restrictive environments. Other states move much more rapidly, like Texas, Indiana and so forth. You don't need to have a centralised system to enable fast speeds of data centre build-outs.

I would say, from the company's perspective, one of the key things they're looking for right now is clarity and predictability over the timeline. Anything that you can do to credibly signal, 'If you comply with certain requirements around whatever it might be—clean energy commitments, investments in local communities—we can credibly say that you will get this data centre up and running on a certain timeline,' is very valuable to the companies. It's about issuing guidance for local communities and for states, providing predictability and providing standardised processes that they can move through at speed so that they're not going to have to comply with one set of processes and then have another review process that's stuck on at the end of that. Anything that you can do to provide standardised, streamlined, predictable processes is extremely valuable for the companies.

Then, in turn, if you can offer that, you can extract quite big concessions on the other end, whatever it might be. The thing that companies are prioritising right now is speed, and anything you can credibly signal about speed and predictability is the most important thing for the companies.

Senator DUNIAM: Thank you.

CHAIR: Ms Chaney.

Ms CHANEY: Thank you for your evidence. It's really useful. I'm interested in how enduring a couple of things are likely to be. As you say, so much is uncertain and we're making decisions in that uncertain world. I'm interested in your views on how enduring the advantage of frontier models will be over other models and how enduring the advantage of centralisation in training compute will be. Obviously, if we're making a big commitment to training, we need to know that that will continue to be a useful thing to have.

Ms Egan: Frontier over open source or fast follower models—one of the things that's currently driving China's AI progress is its ability to distil models from the US frontier. By distil I mean you do a lot of queries of a more sophisticated model and then use the outputs of those queries, like the conversation and any reasoning traces you can extract, to actually train a model at a much-reduced cost. There's been good reporting on China's AI ecosystem rejigging itself to really focus in on distillation, but we're also starting to see movements by the US frontier companies, like Anthropic, to actually curb and limit the effectiveness of distillation.

There's an open question as to how robust this is. This is an area of uncertainty. But I think it's not just about having the most capable model but also about being able to deploy it at scale. We see that Chinese AI models are actually quite far behind US AI models, for example, when it comes to long-term tasks—agentic behaviour that goes over longer time horizons. These same behaviours are those that are going to continue to benefit AI research and development through recursive self-improvement as AI continues to improve AI. There are some people who think that the US lead and compute leadership—having more compute, more agents, being able to be deployed at the same time on more complex things—will lead this flywheel of advancement and we will see this emerging gap.

It hinges on things like whether we curb distillation—and there are efforts moving along in that space—and then whether we will see this flywheel of improvement from the frontier companies. I think that, yes, there's some uncertainty in these things, but I wouldn't want to—the reason Chinese models are so close to US models is that they are distilling, and the US administration has said very strongly that it is willing and ready to crack down on that. So I think the option of distilling a model from a US frontier is not available to close allies and partners of the US and may soon not be available to China. So, long story short, I think frontier models are incredibly capable and incredibly valuable from a strategic perspective in a way that the open source models that follow will not match and will probably fall further behind.

Ms CHANEY: What about the enduring benefit of centralisation for training models?

Ms Egan: There has been some reporting about companies using more reinforcement learning and—sorry, it's a bit technical—high-bandwidth memory environments, like inference-esque compute clusters, to help with training. AI companies are still continuing to expand what we call pretraining, which is when you use a massive amount of compute as a single machine. I've recently been talking to people at the AI labs who think that that is where you'll get real intelligence dividends in the longer term. Reinforcement learning is helping you elicit more capabilities from a set level of intelligence, but it's the centralised compute that they expect to have massive

capability jumps going forward. You can also use it as a proxy—the companies who are closest to the technology and deepest into the technological developments are paying a big premium to build out and construct massive centralised compute—because it would be much cheaper if you could just spread it around the world and not worry about that.

Ms CHANEY: We've had some other evidence from people saying that they think most of the productivity benefit from AI doesn't actually require frontier models and that we'll see a fit-for-purpose approach where you might have a range of AI models at your fingertips, and you use different models at different price points for different purposes. I'm interested in your views on how much that frontier advantage is real in business and in government for different functions. Is it the case that we only need frontier models for five per cent but that it's a really important five per cent? What are the different scenarios on that question? I'm sorry I'm not directing the question to either of you!

Dr Winter-Levy: I can offer some thoughts if it's helpful. I think this is another area where there is huge uncertainty. There will almost certainly be a wide range of tasks where you don't need the best of the best and the most exquisite frontier capability. There will almost certainly be a broad range of tasks where a good enough model, whether it's an open source model or a sub-frontier model, is good enough. There will likely also be a set of tasks—particularly potentially in the national security domain, in anything that's zero sum, where small improvements in intelligence could have quite big returns potentially; in scientific R&D; and in financial trading, where being a bit quicker has really big benefits to being a bit slower—where a frontier model is likely to capture a lot of the value, particularly the scenario you laid out where you mix and match across different models. There's likely to be an orchestration layer above that where one model is figuring out which tasks to delegate to which other types of models. That's likely to require frontier model.

With the caveat there is a lot of uncertainty here, I think premising a national strategy on the idea that you don't really need access to the frontier and you can get by with good-enough fast-following second-best alternatives is a risky bet. It might prove true, but it might leave you very exposed if there really are a set of capabilities that you just need the best systems for. To one of Ms Egan's earlier points, I think access to those frontier models is not guaranteed, partly because compute is just in such short supply that the frontier labs are rationing who has access to the best models, but also because the US government and the labs themselves are worried about the security risks associated with access to the best systems and so are increasingly restricting who can access those. For many countries around the world, the default assumption is that they will just get access to the best frontier systems because US companies want to reach a large market. That may be true, but I think there are some trends that push towards a much more restricted set of actors who actually get access to frontier models or a much more staggered series of rollouts where trusted US firms—maybe trusted US allies, but not necessarily—get access to the best systems and then later the rest of the world does, by which time maybe the US has already secured a lot of the economic gains or a lot of the strategic gains that might come from these frontier capabilities.

Ms CHANEY: It's a risky strategy to say, 'We can make do with second best'?

Dr Winter-Levy: I think so. I think that is a bet on one particular technological future playing out, which may be correct but may not be correct. I think that you certainly need back-up plans in case the frontier proves to be—another way of putting it is that, no matter how many good-enough second-best open source Chinese models you have right now, you cannot solve the types of math problems that OpenAI is solving with its best systems. It doesn't matter. There are a certain class of problems that the frontier unlocks that, right now, other systems do not unlock.

Ms Egan: I'd add to that point quickly. If we are facing a world that is compute constrained, and all signs seem to point towards that, because the builders of the chips are notoriously conservative. We're not taking AI demand signals seriously, and it takes three to five years to build out new clean-room space to make more of the advanced chips, so we're in this lag between having enough compute capacity. If compute is limiting bottleneck, you probably want that unit of compute to be going towards the greatest capability, and what we're seeing is that more capable models are actually using fewer tokens or less amount of compute to achieve the same result as less capable models. So less capable models are starting to need more compute to achieve similar results to more capable models.

CHAIR: Ms Chaney, I'm conscious of time. It may potentially be that you ask a question, and the answer is provided on notice or as a short comment. Senator Sheldon has one final question before we finish the session.

Ms CHANEY: Very quickly, I'm interested in your views on the role of Australia's AI Safety Institute, its level of resourcing and how it could be contributing globally with greater resourcing.

Ms Egan: It does not have the resources to currently be taken seriously. I think it also has an incredibly broad remit, which is providing AI advice to the entire government. But the $7.5 million figure is nothing compared to what other countries are throwing at this. The conversations in other countries are, 'Well, we probably need a lot more investment income capacity.' It gives you the situational awareness to understand what is happening. It gives you the ability to work on these unsolved critical questions around alignment and control. I don't think Australia's AI Safety Institute is currently resourced enough to be taken seriously on the international stage either.

CHAIR: Senator Sheldon, one final question, very short and sharp.

Senator SHELDON: I'll be as short and sharp as I can be. Thanks again for joining us this morning. There's obviously a lot of pressure across the community, across the world, about how we respond to both the opportunities and the challenges with AI and its rollout. What do you see as the timeframe for Australia to be on the frontline to take advantage of AI? What sort of period? Also, what do you see as where the government should be concentrating a lot of its energies to deal with sovereign capability? Where should be our priorities, in your mind, in light of the evidence and conversations you've had in the work that you do?

Dr Winter-Levy: On timeframe, I think there are different ways of thinking about this. In terms of compute build-out, right now is a real moment of opportunity. We're seeing this data centre backlash growing in the United States. We're seeing other countries angling to seize some of this build-out—the gulf states in particular. But there are big risks to building in the gulf. Europe has challenges unlocking large-scale compute build-outs, right now there really is this window for Australia to seize this first-mover advantage in unlocking compute build-outs while the labs are so desperate for whatever source of compute they can find. On compute, I think there's a narrow window that may be closing.

On the broader set of risks, the obvious risk now is cyber risk—securing critical infrastructure from a range of cyber threats right now. People in the labs and Washington are worried about biological risks being the next thing that's coming, the thing that everyone can see coming down the pipeline in the same way that everyone could see the cyber risk coming down before Meta. I think biological risks from the models is the crocodile closest to the boat after cyber. Those are the two things I think that apply right now.

In terms of the timeframe, there's huge uncertainty. If you talk to the labs, they talk about the next year, the next two years, closing the loop on these recursive self-improvement processes. Then, it's not clear what happens after that. I suspect policy will continue after that, no matter what happens. I think it depends on the type of risk, the type of scenario, you're worried about in terms of really playing this out into the future.

Ms Egan: To add quickly to that, I think the timeline is moving quite quickly. AI companies are making decisions about where they build out large, 10-gigawatt scale data centre campuses, scaled to that amount. They're going to make a decision and follow through on that. Interestingly, the data centre moratoriums in the US have actually not slowed build out that much. There's some great analysis by SemiAnalysis on this that show there are a lot still coming online and still more room to move. The department of energy in the US is building out an eight-gigawatt cluster in Portsmouth, Ohio, which OpenAI will be using. The buildout continues, and so the question is: does Australia want to have this influence over AI? If so, the earlier that you start building, the more influence that you will get. As AI compute is built out elsewhere, after companies make big decisions on where to situate training compute, the hand is much less strong when asking for concessions and ongoing support for the creative industry because they've already have made significant investments in other jurisdictions.

CHAIR: I will call time on this session. Thank you so very much, Dr Winter-Levy and Ms Egan, for making yourselves available to give evidence to our committee today. Please coordinate with the secretariat on any further information that you would like to provide to the committee or on anything that you've taken on notice and provide that within the next two weeks.

BENNETT, Professor Jill, Fellow, Australian Academy of the Humanities [by video link]

DAVIS, Ms Inga, Executive Director, Australian Academy of the Humanities [by video link]

DODDS, Emeritus Professor Susan, Policy Lead, Australian Academy of the Humanities [by video link]

THOMAS, Professor Julian, Fellow, Australian Academy of the Humanities [by video link]

[09:15]

CHAIR: Welcome. I understand that information on parliamentary privilege and the protection of witnesses giving evidence to this inquiry has been provided to you. Could you please state the capacity in which you are appearing today?

Prof. Thomas: I'm the director of the ARC Centre of Excellence for Automated Decision-Making and Society.

Prof. Dodds: I'm a professor emeritus at La Trobe University.

Prof. Bennett: I'm a professor at UNSW.

CHAIR: It's fantastic having you appearing before us today. We really appreciate you making yourselves available. I'd like to now invite you to make a brief opening statement, at the conclusion of which we'll get underway with questions.

Ms Davis: Thank you so much, Chair and the committee, for the opportunity to appear today. I'd like to begin by acknowledging the Ngunnawal and Ngambri people, the traditional custodians of the land on which the Academy's HQ is located, and I pay my respects to elders past, present and emerging.

The Australian Academy of the Humanities is one of Australia's five learned academies. We are an independent and non-partisan national body representing more than 740 of Australia's leading scholars and practitioners across the humanities and creative arts. Our fellows bring humanities' expertise to bear on the defining challenges of our time. We advise governments, support emerging scholars and contribute to public discussion, bringing context, insight, imagination and long-term perspectives to national challenges.

To explain what the humanities are, we encompass a range of disciplines, including history, philosophy, languages, literature, archaeology, cultural studies, Indigenous studies and the arts, to name a few. These disciplines really help us understand how people think, create, organise societies and imagine our futures.

I would like to mention that our academy includes 37 Indigenous fellows, whose leadership in Indigenous studies and Indigenous knowledges offers important insights to Australia's approach to AI. Our strategic plan is called *Australia's humanities moment*, and we chose that title deliberately. We're living through a period of profound change, and that is being largely driven by AI. It's not simply a technological transformation. It raises profound questions about knowledges, cultures, creativity and what it means to be human, so we argue that humanities have never been more relevant than they are now and to our future. Technological solutions must work within social and cultural contexts and depend on people adapting to change.

Unless we understand these human dimensions, even well-intentioned reforms can fail. We argue that Australia's sovereign capability must be understood more broadly than industrial, technological and defence capabilities in those terms. It must include sovereign intellectual capability. That's our ability to understand our own history, cultures, opportunities and challenges and to make decisions about our future on our own terms and in our own interests.

In our 2026 pre-budget submission, we called for a sovereign AI research fund. Our submission to this inquiry highlights three priorities: firstly, a national AI research plan, which would include humanities led research and development; secondly, protecting Australia's knowledge and data sovereignty, including Indigenous led approaches to AI and considerations for how best to manage national research collections infrastructure; and thirdly, we'd like to see a national capability approach to AI, which recognises that Australia's future needs more than technical skills. Our central message is that Australia cannot realise its ambitions for AI through investment in technology and infrastructure alone. We must also invest in knowledge, people and institutions needed to shape AI in the public good.

To finish up, I'd like to introduce my three distinguished fellows of the academy who have joined me today. Professor Sue Dodds, our policy lead, is a philosopher with expertise in ethics and public policy. She has particular interest in the regulatory safeguards and legal aspects of AI and in meeting the need for a diversity of skills and aptitudes to be deployed by an AI enabled workforce. Sue is also deputy chair of the ARC board. She is not appearing in that capacity. Professor Julian Thomas's experience can inform the development of a national AI research plan to protect Australia's knowledge and data sovereignty. That can ensure that AI is developed

inclusively and that there are equitable benefits drawn from AI for all Australians. Last, but not least, Professor Jill Bennett is director of the Big Anxiety Research Centre. Her work, and the work of the team that she leads, uses creative arts and humanities research to develop technologies that respond to human needs. Jill can demonstrate the potential to harness the expertise of the humanities to build better AI tools to address pressing social issues.

We welcome any questions you might have, particularly about Jill and Julian's research, but also any other questions drawn from our submission.

CHAIR: I'll kick off. In a broad sense, I'd like your reflections, and those of anyone appearing with you today. We are hearing evidence that the position that Australia is in right now is a relatively unique one globally, given the interest of those at the very frontier of AI in seeking to come here to train their frontier. As you say, there's a lot of focus on what that means for our sovereign capability when it comes to this significant infrastructure and on the ways in which Australians can gain value right up the stack because of this investment.

Do you have some reflections on what else we should be seeking as part of our opportunity for leverage over those at the very frontier of this technology—by having them here, having them bring compute and having them invest here in the compute, capacity and training? We heard, including from our previous witnesses, the extent to which we can help shape global governance, safety and regulations in that space. To your point around the protection of humanity and how we keep humanity at the forefront given the rise of the artificial, what are the things we should be looking out for in Australia's interests, from your perspective, when leveraging this interest at the frontier?

Ms Davis: To leverage maximum impact from this frontier coming to Australia?

CHAIR: Yes.

Ms Davis: There are tremendous implications for our sovereign capability and for employability for humanities graduates. I think it's probably worth noting that the frontier AI companies are founded by and are indeed employing humanities trained experts. Anthropic was co-founded and is run by co-directors who have English literature degrees and creative industries degrees. We're also seeing some of these frontier providers employ a range of humanities trained professionals, including philosophers. We're seeing that expertise infiltrate many levels of the AI knowledge ecosystem. You are hearing from experts with this kind of training—Kate Conroy, for example, who runs the AI Safety Institute, has a PhD in philosophy, and Sue Dodds, presenting today, has a PhD in philosophy. There are opportunities for us to be developing our sovereign cultural and social capabilities and to be valuing that training in Australia's interests and for opportunities to be spilling over into other sectors. Sue, would you like to add a further comment?

Prof. Dodds: The opportunity is really for Australia to recognise what it's got at stake and look for opportunities to create guardrails—regulatory, legislative or principles based—which ensure we don't lose what we most value. A clear example are things like First Nations data sovereignty. We've got a unique set of enduring cultures within Australia that, on the one hand, may benefit from improved access and capability in AI but, on the other hand, have been working hard to ensure that the care principles around Indigenous data sovereignty are protected as we go forward. That helps to recognise what's at stake.

Similarly, one of the examples within our submission—the language data capacity—is that we've got more community languages in Australia than many other successful multicultural societies. That means that we've got responsibility with regard to Pacific Island languages to a range of Australian Englishes, as well as to community languages, which we haven't yet been able to fully explore. There's both a risk and a benefit here. There's a potential for us to get much better input into our knowledge of ourselves, but there's also a risk that, if models are developed that don't have that diversity and depth of knowledge within them, that we'll get outputs which are not successful. It's looking for the opportunities there.

Similarly, with questions about the purposes of government and the services they provide, ensuring that there are appropriate mechanisms for when you put the human in the loop or how we have explainable AI—if it's the case, in a complex liberal democracy that, in principle, people who are affected by policymaking need to be able to understand why the outcome has gone the way it has——then we need to have mechanisms which allow for some degree of transparency within the systems. There are aspects of our unique legislative, cultural, social and environmental background that we can seek to develop appropriate measures and guardrails to address as we go forward as well as recognise that the unique characteristics of Australian character risk being homogenised in a system that doesn't have sufficient granularity within it. Julian might have further comments to make on this point.

Prof. Thomas: I come to this thinking about some of these questions around how we share the benefits of AI going forward across the whole of the Australian population. The context for it is that we still have a relatively undeveloped digital infrastructure in this country—not just for computation but also for communication and for the delivery of essential digital services. When we're thinking about very substantially augmenting that infrastructure for the purposes of either inference or some training for AI, we also need to be thinking hard about what that means for Australians in remote and regional Australia and whether there is an opportunity to expand and extend the opportunity to participate in an AI driven digital economy for everyone.

There's another element there as well, which is to what degree we might want to say that some of that extraordinary new computational infrastructure might be available for public-interest purposes—for addressing, for example, the kinds of digital disparities and inequalities we have, the difficulties we encounter delivering good digital services across the whole country, and the capacity of our public and not for profit sectors, which are so important in delivering digital services that are increasingly AI driven. How are we going to manage that challenge? If we are looking at what would be a significant step-change in digital infrastructure for the whole country, I think we would be hoping that we could take a national, holistic view of what would be required, and undertake the necessary consultations accordingly so that everybody is at the table.

Prof. Bennett: It's very important to do the bottom-up work. For example, I lead an area that we call psychosocial AI or psychosocial technology. The method for this is to work from communities to work out where the needs are and how we can service them. For example, we've developed AI companionship systems for older people and people with dementia, among others. This is an area of need which is not effectively serviced top-down, particularly if you're looking at massive problems and service challenges across regional Australia, where we think there's a real opportunity. We've been looking at loneliness and isolation, and the shortfall in services across these thin markets. Often we find people who are, maybe, diagnosed with dementia or mild cognitive impairment and they don't see anyone for two weeks other than a home care worker for an hour. That, in itself, increases risk of dementia by 50 per cent. We're looking at how we can use technology—in this case, an AI companionship system—to do the cognitive and emotional support, the everyday reassurance and companionship that is absolutely vital for remaining healthy at home. There's a real opportunity here for making change at scale, but you can't do that with a top-down, off-the-shelf model.

We spend half our time counteracting the normative biases in LLMs in terms of their US idioms—they obviously don't speak to people in the bush with dementia, right? It has to be a really attentive bottom-up system working with communities. This is why we call it psychosocial—it's all of the social communications disciplines as well as psychology and anthropology coming together to develop the tools. There is actual development in there. We create an AI module to do that interface. Globally, this is the case. We also work in China, where there's no problem developing the technology and rolling it out but there is a problem in the human-centred implementation. That's the challenge. That's something that with our discipline—I don't like using those terms anymore, because everything is interdisciplinary now. These are all new mixes of disciplines and practices, but, as we've been saying, we need these humanities skills in the development process for these tools that are ready to be implemented.

CHAIR: Following on from that, given all of what you've just provided in terms of the importance of not just our AI sovereignty but our own data sovereignty, our own culture and what it means for Australia—in particular, I take the point around First Nations, which often doesn't seem to be a key focus in some of these conversations—what's your view of what the potential opportunity cost is of our ability as a nation to be able to achieve those things that you've just been talking about, and the importance of those things if we don't have training here? If we miss out on training and these big frontier models continue training elsewhere or it continues to happen without us and we just continue as we are, what do you think the opportunity cost is for us to truly capture the key points you've been making?

Ms Davis: I might hand that one over to Julian, but before I do—some of the examples that you were hearing today from Jill and Julian in particular aren't benefits that are going to automatically materialise because we have attracted frontier technology training to Australia. The regulatory and policy settings need to be right. There needs to be some kind of incentive in place and a programmatic approach led by government to say, 'We really value this, and we are going to look at funding it through a programmatic approach.' It's got to be stimulated. It's got to be called out. That's why we've made the call for an AI research sovereignty fund, which we're happy to talk about. Julian, did you want to talk about the missed opportunity? Chair, your question was about opportunity cost and whether or not we should attract the frontier tech to Australia.

Prof. Thomas: Thanks very much, Chair, for the question, which is a terrific one. I'm sorry to say this, but I think it's a little unknowable actually. We know there would be a cost. We know that there will be a cost if that

does not happen in Australia. We can well understand why it is considered so desirable that we do have training capabilities in Australia not just for frontier models by the way but for the development of our own AI systems across the board, from the little things to the big things—all of those things, we think, are going to be very important. But there's absolutely no doubt that, if frontier model training were occurring in Australia, there would be—and this has been pointed out in a number of the submissions you've received, and we do not argue with these—significant positive spillovers in many ways.

But I suppose we do think that we might be losing a little bit as well if that were all that we focused on. There's a risk that way too. That is why we are talking about why we think a coordinated national plan or research around AI, some ongoing attention to the possibilities and potentials of AI driven national research infrastructure that we are starting to see—and we welcome the exposure draft of the expert group on NRI, which has recently been released by the Department of Education. That's an important document. Susan already mentioned the very significant work that's going on in that area through LDaCA, the language data commons. That's just one example. There are many others.

I'm sorry if that doesn't answer your question directly, but I think that's the way we would approach it. We can see the benefits, and they are likely to be very considerable. They're a bit unknowable because, as I think many of your interlocutors in these hearings have been saying, there is just a great deal that we do not yet know about how the technologies we are talking about are likely to develop, where they will be located and what sort of benefits might flow from them. But our sense is that in this situation it's also equally important not to be too distracted by what looks like a single prize or a single trophy. We have a lot of hard work to do across government in our universities and in the academies, of which we are one.

We have a lot of hard work to do to coordinate, augment, extend and develop our national capability more broadly, and we would really like to see some attention on that.

CHAIR: Sure. I appreciate that.

Prof. Bennett: I think we should also say that you have to do the social and human dimension of the implementation science rather than only focus on the infrastructure, and that has to be done alongside and upstream, not as the end point. If we wait until the implementation happens, that's when you get robodebt and a repetition of these unintended consequences.

CHAIR: Thank you. Ms Chaney?

Ms CHANEY: Thank you very much. I have just a couple of questions. Firstly, I want to look at election integrity. The Australian Ad Observatory has exposed hidden political influence and transparency gaps. With generative AI making political messaging more personalised and hard to observe, I'm interested in your thoughts on what we need to do before the next election to safeguard democracy given those changing risks.

Ms Davis: That's definitely a question for Julian, because that research is [inaudible].

Prof. Thomas: Thank you very much for the question. I think it's a good example of the kind of very positive humanities-led R&D that really helps us address some of the aspirations of the national AI plan in relation to security assurance, trust and protecting our institutions. The Australian Ad Observatory is a project of our ARC research centre. Also, I'm really pleased to say that it's evolving into a really significant element in our national research infrastructure. It's part of the Australian Research Data Commons humanities and social science and Indigenous data commons program. It has been funded within what we call the Australian Internet Observatory, which aims to extend the capabilities and the experimental methods developed with the Ad Observatory into larger areas.

This area of political advertising is absolutely critical, so there is a lot of work underway, I would say, on this. What we have been doing is developing the tools to enable us to understand what kind of advertising Australians are exposed to on digital platforms, especially during election periods, where this is a really critical issue. The problem we're dealing with, as you well know, is that, to date, these platforms have been a black box. The problem is that, in the broadcasting environment, it was all transparent and public. We could see who was advertising what and when, and in a regulated environment around elections it was possible to monitor that. That then ceased to be the case with the emergence of the digital platforms that we now have, so there is a huge problem about accountability, transparency and the integrity of the process.

What we've done is develop the tools. We're working with the users of those platforms rather than the platforms themselves. It's a kind of citizen science project where people donate to a research project the data they encounter in their interactions with digital platforms. We put all of that together and, using AI, we can understand what the patterns are, what people are seeing, what we know about this advertising and what kind of thing is happening there.

Ms CHANEY: Are there any actionable findings out of that so far in terms of transparency or other requirements? Given that time is of the essence and things are changing rapidly—I saw an AI video of Jacqui Lambie yesterday saying things that she did not say—what can we do in a short timeframe?

Prof. Thomas: What I think may be most useful for you is if I send you our most recent reports out of our work on this, because we believe there are things that we're finding and that it is necessary that they be addressed. What we are doing is making the research infrastructure more robust so that we can, for example, start to see what is going on on mobile devices, which were always extremely difficult to access. It is a bit of a slow process to get this right. To achieve observability of this area is no small research challenge, and I think I can say that Australian researchers are actually at the cutting edge of the field in relation to this.

Ms CHANEY: Are you able to say what some of those things that you're starting to see that could be done are? We really want to focus on actions that we can implement before the 2028 election.

Prof. Thomas: I wouldn't want to try to say that now, if that's alright. Let me pass that information on to you. I'll take it on notice, if that's okay.

Ms CHANEY: That can be a question on notice. That would be great. My other question is about liability in relation to agentic AI. Maybe that's a question for you, Sue. I'm interested in your thoughts on how we allocate liability for this really novel legal issue.

Prof. Dodds: I think it is a really important issue. I don't have a ready answer—I'll start with that—but I do think that, when thinking about where the guardrails should be, taking a more holistic view than we might have in the past—it's not just a simple piece of tech that we're looking at. If we were to have a frontier data centre here that was working on this front of work, we would be looking at tech that's had a lot of Australian data built into it. Thinking about where accountability more generally sits—liability being obviously just the bit where you pay out—legally, accountability and responsibility may sit between the previous owners of the data, those who have control over the information and those who are using the outcomes of the data. We have the capacity to look at this in terms of various slices, but it really needs to be thought through. It's often shifting a bump under the rug. That's what liability is often about—seeing where we as a society think we want that responsibility to sit. Given that a very significant amount of our data is related to government purposes, whether it's Commonwealth data, and that we've got the question of offshore owners of the technology, how do we work with the intersection there?

But, also, if you're thinking about the example you were just talking about, around elections, how do we make sure that we've got appropriately skilled and thoughtful guardians overseeing the way in which the implementation is occurring so they can anticipate and be responsible for the ways in which some of the negative consequences may occur? This is where I think that bigger research plan and the idea of AI safety come through, as does making sure that they're appropriately resourced so that they're not just reactive. You might think that, with regard to liability, part of it is 'what could we have avoided?' not just 'how do we respond when something's gone wrong?' Really building that in from the beginning is going to be an important part. That requires people who are not simply expert in the tech but alert to the broader social, legal, cultural and environmental impacts of the application of AI. That's not a clear answer, but I do think we need to be starting, right now, to be looking not at liability limitation but at understanding the broader impacts and then looking at where liability might sit. I hope that's helpful.

Ms CHANEY: Thanks. There are no more questions from me, Chair.

CHAIR: Excellent. I don't think we've got any further questions from our committee members with us today, so we'll leave that there. Thank you so much, again, for making yourselves available and for the really thoughtful contribution to the discussion. I think it's really important that we keep front of mind the humanity element in this discussion. It is crucial to who we are as humans. As a person from the humanities—I'm a psychologist by training—I really appreciate and value the discussion that you've had with us this morning.

FOSTER, Ms Nicole, Director, Global Artificial Intelligence Policy, Amazon Web Services [by video link]

O'ROURKE, Mr Matthew, Head, Infrastructure Policy, Australia and New Zealand, Amazon Web Services

STANIC, Mrs Rada, Chief Technologist, Australia and New Zealand, Amazon Web Services [by video link]

[09:51]

CHAIR: Welcome. I understand that information on parliamentary privilege and the protection of witnesses has been provided to you. Do any of you have any comments to make on the capacity in which you appear?

Mr O'Rourke: I am head of infrastructure public policy for Australia at Amazon Web Services.

Ms Foster: I'm based in Toronto and I lead AWS global public policy for AI.

Mrs Stanic: I'm based in Sydney.

CHAIR: I now invite you to make a short opening statement. Then we'll proceed to committee members for questions.

Mr O'Rourke: I would like to start by acknowledging the traditional custodians of the land on which we each meet today and pay my respects to elders past and present. Thank you for the opportunity to appear before the committee. The questions this committee is examining are ones we engage with every day through the customers we serve, the infrastructure we operate and the communities we work in across Australia. We acknowledge the letter sent by the committee to the Prime Minister earlier this week. We are continuing to digest and work through the detail of what is a substantive piece of work. However, we certainly echo the committee's recognition of the importance of these matters.

At AWS, we have been building and operating in Australia since the launch of our Sydney region in 2012. Since that time, we've invested $18.9 billion building our Sydney and Melbourne regions and have announced plans to invest a further $20 billion through to 2029. That investment translates to data centres, jobs and computing power that hundreds of thousands of Australian customers, from startups to government agencies, rely on every day.

Our research, by Strand Partners, shows that 61 per cent of Australian businesses have now adopted AI, up from 50 per cent last year, and 79 per cent of those report meaningful productivity gains. That's encouraging, but it also means that there are some businesses and workers who haven't yet been able to capture some of those benefits. In Australia today, researchers are using AI on AWS to reduce epileptic seizure frequency, get patients to a neurologist faster and help stroke survivors find their words again. These aren't hypotheticals. They're built by Australians in Australia, and we want more stories like these, not fewer.

We recognise that the pace of the technology raises legitimate concerns about safety, transparency and fairness. We take those seriously and have built our approach to responsible AI around them. We also recognise that no single entity can solve these challenges alone, which is why we work across sectors with governments, society and other companies and welcome the work of this committee. We're happy to take the committee's questions.

CHAIR: Thank you, Mr O'Rourke. I'm keen to understand the opportunity for Australia in terms of the investment in compute, in particular, and our sovereign capability. How are we placing ourselves as a country and what leverage are we able to use with the interest in building that out? I'd be keen to get your perspective on some of those things. What proportion of future compute capacity, from your end, do you expect will be used on AI and what proportion of your Australian compute capacity will be used to provide tokens overseas?

Mr O'Rourke: I might start by just talking about the opportunity aspect there, because it is extraordinary no matter how it's measured. I think the Productivity Commission has identified that AI can contribute $116 billion to Australia's GDP over the next decade. The Treasury has described AI as the first 'substantive, credible global productivity growth accelerant' the world has seen in the last two decades. So there's a meaningful opportunity there for Australia to catch up on some of the sluggish productivity growth over the last decade.

But I think the numbers that we are most excited about are the human ones, which are the ones I touched on in my opening statement. Today in Australia there is a research group which has been able to reduce the frequency of epileptic seizures by 30 per cent and increase the speed with which patients can get in front of a neurologist by 70 per cent. In Queensland there's a research centre which has been helping stroke survivors find the words that are on the tip of their tongues in real time using AI that's on AWS. And there's a new entrepreneur which has built, on AWS using AI, a bushfire detection platform which has been able to identify 30,000 bushfires since 2021 and, more importantly, notify the first responders about five times faster than human reporting. These are all

Australian researchers and entrepreneurs who are solving Australian problems on infrastructure that's here in Australia. And they're the types of activities that we get most excited about and would like to see more of.

Our research from Strand Partners shows that AI native Australian startups are averaging 156 per cent annual revenue growth, which is 92 per cent more than Australian startups generally. Within that cohort, we would love to see the next Canva, SafetyCulture or Nera, which were all really successful startups that started their journey on AWS and continued.

Turning to the proportion of compute, everything we build in Australia is overwhelmingly for use by Australian customers or for Australian end users. Those examples that I spoke to are Australian businesses, not-for-profits or researchers, which are building on AWS to serve their own business needs or improve their own operations in order to provide services to their end customers. There are also businesses which might not be Australian but are serving Australian customers. A classic example there would be Netflix, which is using AWS to move its streaming content closer to Australian end users.

There will be some edge cases. For instance, we launched our New Zealand region in Auckland, I think in September last year, so there will have been, before that, some New Zealand based customers who were building on our infrastructure here, but they're probably in the process of migrating their workloads back to Auckland. Likewise, there will be Australian customers here who build on our infrastructure here, but whose customers might be in other countries; they're exporting their services. So they can easily still deploy on our infrastructure in other countries if their customers are in, for instance, Germany, the US or parts of Asia. So, overwhelmingly, the compute capacity we're building here is for our Australian customers or for those who are serving Australian end users.

CHAIR: To that infrastructure build—and I understand Amazon has made significant investments in infrastructure—can you talk to the committee about the energy element of this? Obviously, a pretty significant part of the AI infrastructure build-out will be the consumption of significant additional energy and the contribution to the grid. From your end, what proportion of the future electricity demand could realistically be matched by new renewable generation?

Mr O'Rourke: We're focused on matching all of our energy consumption over time. For the last six years, before renewable energy matching was contemplated as a regulatory requirement, we have been investing in renewable energy in Australia. As a result, we were the largest corporate purchaser of renewable energy in Australia last year, up from third place. We have invested about $2.8 billion over that period in renewable energy generation and storage assets. Our portfolio is about 20 projects, equivalent to almost a gigawatt of capacity. Assuming a very conservative capacity factor of about 25 per cent, that is equivalent to about one per cent of the national electricity market, so it's quite a significant scale of investment that we've been working towards. Our motivation has been driven by our goal to achieve net zero carbon by 2040. That's been one of the driving factors behind our appetite, and we continue to double down on that, recognising that it's now becoming an expectation and a requirement at both state and federal levels. We're certainly continuing down that path. We recently announced a couple of standalone battery projects, which is certainly a first for Amazon globally. This was in Bairnsdale in Victoria. It's a new standalone 50-megawatt battery. Other battery projects that we have done have been hybrids integrated into the generation project, but the standalone battery is really to support the firming of the grid. That's been a another opportunity where we've been able to contribute something quite distinct.

CHAIR: Excellent. That's enough from me. Deputy Chair.

Senator DUNIAM: Thank you very much for making yourselves available today, both remotely and in person. We're trying to be as constructive as possible. There are some who've confused our constructive approach with deference to certain witnesses, but I think it's just the right thing to do to try and find a way through this policy issue for Australia. I just wanted to round out on this issue of energy. Mr O'Rourke, you mentioned—I can't remember your exact language you said—matching your energy use with the contribution, and then you said 'over time'. Is it the case that you intend to expand your compute power and/or capacity and therefore power usage with generation simultaneously, or is it a catch-up kind of program? How does that work in terms of what you mentioned before?

Mr O'Rourke: To date, because we have been focused on achieving our 2040 target, we have been focused on matching our forecast over that period. Essentially, we've been chasing a moving target. The introduction of the regulatory schemes now is obviously changing the timeframes for some of that. We're really just bringing forward some of the projects and investments that we would otherwise have been making. In our case, we're not looking to match one-to-one on a facility level, because, as a cloud provider, we really take the perspective that it's a portfolio. We're looking across all of our electricity consumption, and we're just trying to match that at scale.

The advantage there is that we are able to do large-scale offtake agreements with renewable generators because we're looking across our portfolio and across our entire growth rather than at individual facilities.

Senator DUNIAM: Sure. Thank you. In the wake of the Services Australia unauthorised access issue that we confronted a week or so back at the Sydney hearings of this committee, obviously there's a degree of risk around government data and how it's managed and being protected. There's a risk, of course, when you feed that data into an AI model, that that data could be taken offshore by that model operator and used for training purposes. This is something I look forward to asking other frontier labs and model operators when we have them back. What's Amazon's approach to this? People are concerned about their data, whether that is personal information or data that's related to our nation's security et cetera. What's Amazon's approach to this, and how do you deal with any risk that presents?

Mr O'Rourke: I might ask my colleague Mrs Stanic, our chief technologist, to answer this. It really relates to our B2B model as well, which I think is a distinction from some of the other witnesses.

Mrs Stanic: Thank you for that question. Rightfully so, data and unauthorised access of any sort is front of mind in the way that we design and make available our infrastructure to our customers. And I'm pleased to say that in the way our global cloud footprint is designed, including our regions in Sydney and Melbourne running here, we provide full flexibility and ability for our customers to store all of their data on Australian shores. That means they're fully in control at all times of where their data resides, and they can choose for all and any data to stay in our data centres and what we call our regions in Sydney and in Melbourne. This is very, very important to our regulated industry customers—banking, health care and government, of course. But there also are businesses who may choose to host some of their data overseas. What I'm saying is, there is flexibility to fully remain in control and to contain that data in Australia. This applies beyond what we've traditionally provided with cloud. This also applies to the way that AI capabilities are built and used on AWS cloud. We provide mechanisms and ability to keep any inference to any of the models that we make available fully in Australia. No prompts and no data ever leave Australian shores, when we're talking about AI specifically. Those measures and controls exist and provide full control to any end customer of Amazon Web Services.

Ms Foster: Just add a little bit more to that, I appreciate the question because it's a really important one that we do get asked a lot. First of all, no AWS customer data is ever used to train models, either an Amazon AWS model or any of the other third-party models that we make available to our customers. Australian data that may be used in your own application is never used in training the model, going back to either AWS or the model provider. A good example of that could also be that there are questions even around the prompts that might be used. When you deploy a third-party model in Bedrock, which is our broad service offering to access generative AI, as you may be using prompts to query that model, even those prompts are kept private to your own workload. Even those prompts do not go back to the model provider or back to AWS. Those are held within the customer's environment. It is completely a secure environment for leveraging generative AI.

Senator DUNIAM: Alright. Thank you very much for that. That is helpful. As I say, it's one question I failed to ask of the others at this end of the spectrum previously, but I'll revisit that. I have two more questions. One relates to the issue of copyright here—whether there are any perspectives or views on a way forward to resolve the copyright issue in Australia. Then I'll talk about investment windows. But any perspectives on copyright, I would welcome.

Ms Foster: Amazon is a rightsholder as well as being an AI deployer and developer. We very much are respectful of copyright holders' concerns on this matter, and we are very keen to continue to work with the Australian government on finding pathways forward and solutions that would enable further AI capacity to be built in Australia, including AI training capacity.

Senator DUNIAM: So the ball's in our court, or the Australian government's court, to come up with a solution. I appreciate that. The final question I have is on time. You're a multinational company, and you're looking at a range of jurisdictions. There are mixed views in this country about which way we go, and indeed whether we progress at all. For a company like Amazon or AWS, what sort of timeframe are you looking at for investment decisions to be made around a presence here in Australia?

Ms Foster: They seem to be rolling investment decisions, so I don't think that there is, at this point, a risk of missing the boat right now. I think that is potentially a concern. As growing investment decisions are made, we may reach the point where we feel like we have enough capacity. But at this point there seem to be continually growing appetites for capacity. In addition to that, in terms of that growing need for capacity, we are very much reading our customers' signals about where they want that capacity. That really is the nudge to us about where we might build specific AI training capacities—where our customers are looking for it. I don't know if Matt might want to add something further on.

Mr O'Rourke: The $20 billion investment that we announced almost 12 months ago is investment that we're currently ploughing into the economy and building up. That is really reflecting the growing demand amongst our customers for AI but also cloud compute capacity. One really drives the other; customers want their workloads near their AI so they can experiment, innovate with it and integrate it into their use of AI models. And vice versa, if they're wanting to innovate with AI, they might be bringing more of their data into the cloud as well. As Australia's digital economy grows, we are really needing to continue to expand ahead of that to make sure that there's sufficient capacity for all of those Australian businesses and other customers to grow into.

Senator DUNIAM: Thank you very much.

CHAIR: Ms Chaney.

Ms CHANEY: I'm interested in your views on allocating safety accountability across the tech stack. You host your own models and third-party models through Bedrock, and those Bedrock guardrails are optional. I'm interested in what know-your-customer checks or misuse monitoring AWS applies or could apply to prevent its platforms from being used for offensive cyber operations or biological misuse or other harmful purposes. When a customer builds something harmful on AWS, where does Amazon think responsibility sits?

Ms Foster: In lots of different places, actually, so it's not going to be a straightforward answer. Responsible AI has to be built in at the beginning all the way through to the end. Keeping bad content out of your training data is one of the first things that you would do. In terms of the responsibility, that's going to be the responsibility of the developer. Sometimes that is Amazon, with our first-party models, and sometimes that's the third-party model provider. Then, as you go through, there are a number of ways that you train the model and also fine tune and optimise the model to ensure that it's behaving in the way that you want it to behave.

As a model provider and as a service provider, we do a number of different things to provide different types of responsible AI documentation and provide guidance to customers about in which ways the services or models have been designed to be used so that they don't become inadvertently used in the wrong way so that you get slightly less optimal performance for how they're deployed. We impose all kinds of different testing and evaluations. Some of those are done in-house. Sometimes we use external providers to assist with red teaming, particularly in areas where Amazon might not have domain expertise.

Ms CHANEY: Just to clarify, is this in your model development?

Ms Foster: Yes, in our model development. We published the way in which we approach this in our frontier model safety framework. We are also very clear about what we're doing and how we're doing it.

Ms CHANEY: I'm just wondering about that at that hosting level. What accountability do you see at that hosting level, particularly?

Ms Foster: We see a couple of different things. The really in-depth red teaming and evaluation we would view as the responsibility of the model developer that we're hosting. But we do a security and safety evaluation before we host those models. In addition to that, we also are providing a secure environment for those models, so that also provides a level of customer security. We also provide customers a model evaluation tool. Different models will perform differently depending on the data and the particular scenario that you're using. Some may be more cost efficient. Some may just be more effective. We have a model evaluation tool that customers can use to evaluate whether or not that model is performing the way that they want it to perform, and to evaluate the pros and cons of that performance. Then on top of that, we have the guardrails that customers can tune up or tune down depending on their use case as well.

We were the first cloud provider to certify several of our services against the ISO 42,001 standard. That is a risk management standard. We also provide a lot of responsible AI documentation and we also do a lot of additional transparency reporting through the framework, as well as the Hiroshima AI process.

Ms CHANEY: I'd love to hear a little bit more about the security and safety evaluation you undertake before hosting. What does that look like?

Ms Foster: It's not as in depth, for example, as what the model provider will do in terms of their own deeper safety evaluations and thresholds. But I would also say this is an area of work that is evolving. For example, on some of the more capable models that have been released by some of the labs, we've done much more in-depth work around some of those models just because they represent higher capability.

Ms CHANEY: On voluntary versus mandatory, obviously, a lot of the things that you're doing at the moment are voluntary. I'm interested in whether Amazon would support making some core elements of safety and transparency reporting and other measures mandatory in Australia, such as pre-deployment testing of frontier models by the AI Safety Institute, mandatory incident reporting. Would you support that being mandatory?

Ms Foster: Yes. Keep in mind, in addition to being an organisation that makes models available through Bedrock, we're also the operators of a lot of very important digital infrastructure, so ensuring resilience of systems and digital infrastructure is also a very important priority for our organisation. Security is probably the most important priority that we have as we think about how this technology is evolving, and ensuring that we have the ability to maintain security of systems. There are a few things that I would suggest as we think through mandatory. One would be that we start with the standards first and that we have some agreed-upon metrics, methodologies and risk thresholds that we're evaluating against. The other is proportionality—making sure that what we're evaluating for is actually aligned to the threshold of the model that we're evaluating—not all models have the same capability—because we also don't want to create barriers to entry for new model providers.

The other would be trying to find opportunities for international cooperation, and that can take a bunch of forms. Australia is already sort of part of this international network of safety institutes, and it's a really important opportunity for a couple of reasons. One is that the safety of the overall digital infrastructure is not something that can be isolated in one country. We need to be able to cooperate in order to evaluate and mitigate any of those risks, so that international cooperation is hugely valuable for understanding both what the risks are and the best methodologies for evaluating and mitigating those risks. So that piece is really important.

Also, for model providers, having either mutual recognition or shared evaluation is going to be super important in terms of making sure that we're focusing on managing risk, as opposed to checking compliance boxes. So I think we want to make sure that we are putting the efforts on the right things. Some ability to have an international or at least multilateral approach to those things is going to be really important to enhancing security.

Ms CHANEY: How equipped do you think the Australian AI Safety Institute is to be a credible voice in those international discussions with its resourcing of 25 people and $8 million?

Ms Foster: I think the important thing is quality over quantity. I think that the way that Australia is going to be able to punch above its weight is to ensure that it's the right 25 people or that you are really attracting the right kind of individuals that have the deep expertise that you really need to do this work. A lot of work can actually be enhanced by using AI to do evaluations of AI, so you want to make sure that you're hiring the strongest people. The other way to maximise those resources is leveraging that international network and relying on that collaboration with other international safety institutes. So I think there are some opportunities to maximise—

Ms CHANEY: So you're saying it might be adequately resourced at the moment?

Ms Foster: It could be. Quality over quantity, I would say.

Ms CHANEY: You're the first witness who's said that they think it's adequately resourced at the moment.

Ms Foster: I don't know that I can give a number that would be the right number or the right way to pitch it. But also I think that oftentimes organisations need to grow organically as well, so I think starting with the right talent at the beginning is a good start.

Ms CHANEY: I've got one more one more, but I'll stop there, Chair.

CHAIR: No worries. Thanks, Ms Chaney. Senator Sheldon.

Senator SHELDON: Thanks again for joining us, Ms Foster. I just wasn't quite clear on one of the answers you gave to the excellent questions from Ms Chaney. I just wanted to ask about the mandatory reporting to government. I'm just not clear. Are you supporting the view that it should be mandatory reporting to governments, or are you not?

Ms Foster: I think that there should be reporting, whether it's mandatory or voluntary. That transparency, I think, is the right approach. There are a couple of things that I would add to that. One is the importance of having international or multilateral standards around how that is done. The other thing is that the reporting needs to be done in a responsible way, similarly to how it's thought about in cybersecurity reporting. We want to make sure that for any potential vulnerabilities, before that information is made public in any way, mitigations have been put in place so that you don't actually enhance risk or provide a road map to adverse actors. So there are two important things we need to align, I think, around the standards of how we report and what constitutes an incident that should be reported and also just making sure that we are careful about how that information is made available or not and we ensure that we're able to mitigate risks before that information is publicly available.

Senator SHELDON: Thanks. Mr O'Rourke, thank you for your opening statement. I think you were here just at the tail end. Maybe you saw some of the evidence from the Australian Academy of Humanities. We've had a number of witnesses through this inquiry talk about not only the critical questions of technical responsibilities with AI and the technology that is evolving from that and is here now but also the humanities aspect of what we need to be preparing at the same time—the human interaction. Part of our inquiry includes questions about the

effects on the quality and security of jobs as well. It covers more broadly the effects on living standards and the community more generally.

You'd probably be aware. I'm sure, in all your preparation for today, you would have expected a number of questions about social licence as well. I take heart from a comment you made in your opening statement: 'Technology raises legitimate concerns about safety, transparency and fairness. We take those seriously and have built our approach to responsible AI around them.' What I've got a concern about is that there's been a court case—I've made some documents available to the witnesses and the committee—in Canada, where there's been the use of AI to surveil workers. In that use of AI, the courts found damaging evidence and findings against Amazon. What's your knowledge of the Atlas system that's been used by Amazon?

Mr O'Rourke: I was just given hundred-page documents as I arrived, and I'm not familiar with the issues or the documents themselves. We're here representing AWS, and we just don't have any familiarity with those issues. I'm sorry.

Senator SHELDON: I'll go more broadly if you're not aware of the specific matters that I've put to you. More broadly, are you in a position to talk about the surveillance systems and the social licence questions that Amazon are facing?

Mr O'Rourke: From AWS's perspective, I can only speak to our infrastructure issues.

Senator SHELDON: We're in an inquiry we're carrying out into the use of AI. I thought it was a particularly poignant point you made about how technology raises legitimate concerns about safety, transparency and fairness and the responsible approach. I'm not directing this at you, Mr O'Rourke, personally. I'm directing it at the company, and Ms Foster may be able to add something as well. I've got a deep concern about a company coming in front of an inquiry in this country and not having people who can give evidence about the humanities question and the impact of the technology that's rolling out, specifically where they avoid both in their submission the impacts of how they're using AI within their own company. I've got to say, I'm frightened about the company using it inappropriately with its own workforce and then allowing that or disregarding how it might be misused by other actors. Do you think I should have those concerns?

Mr O'Rourke: I think, with any concerns around how AI might be used by other actors, Ms Foster can talk to our safety and transparency frameworks around the services that we make available within AWS.

Senator SHELDON: Ms Foster?

Ms Foster: I'm not familiar with the specific project or issue that you've raised. I am less familiar with the operations on the Amazon side of the company, and I'm more familiar with how AWS designs its services and makes those services available and understanding the policy implications of regulation.

Senator SHELDON: Going to the fairness question then—maybe I'll do this broadly. I'm asking the question in a broad context, but I'll just add for the record that it is about the Atlas model that Amazon uses. The company uses a labour risk framework to clarify how Amazon evaluates emerging workforce issues, labour union targets and internal or external labour threats. That's within Amazon's own documents. The Atlas model ingests hundreds of data points. This goes to the critical question about the humanities. It's not just the technical question but the humanities question on the use of AI and the social licence question across the world, let alone here. The Atlas model ingests hundreds of data points: employee trainings completed, health and safety incidents, simple first aid cases to serious injuries, feedback from workers on management, worker tenure and turnover, the proximity of a warehouse to another workplace experiencing union activity, local minimum wage rates, other economic conditions and even the temperature inside and outside the warehouse. It also, actually, looks at private conversations that take place on Reddit and elsewhere by the workers, amongst themselves, about Amazon and, of course, any other matters. I'll ask this now in general terms. I guess this is the humanities question. Would AI being used as a way of surveillance on a company's own workforce and on their own personal interactions and for general blanket use across the entire workforce be appropriate for any company?

Mr O'Rourke: I'm not familiar with any of the particular issues that you've just raised, but, in terms of AWS, in our data centres, we will make use of CCTV for security and safety reasons. For example, one of the suggestions that actually came up from staff—there's always constant idea generation and innovation about how work can be improved and made more efficient and safe—relates to the use of CCTV for identifying person down instances. This would be where a worker might be potentially working by themselves in a data hall. If they fall, that can be detected and an alarm can be sounded if help or assistance is needed. In our data centre space, it's an example of how AI is used to detect and respond to safety issues to ensure that the workforce is working safely, productively and happily.

Senator SHELDON: I certainly respect the opportunities that AI gives. There's opportunity and risk. Risk goes to the humanities question and the social licence question. Would you be surprised, Mr O'Rourke, that, actually, in your Melbourne warehouse, within these documents from the court—as you mentioned, it's 153 pages, and I'm not expecting you to have read through all of them. If you take this on face value, what I'm putting to you is that not only did it apply a global surveillance in spying on its workforce but it also included sites in Melbourne in that global spying. It's in the documents that have been available. Would you be concerned if that were the case?

Mr O'Rourke: Sorry. As I said, I've got no familiarity with the issue. I can only talk to AWS's issues.

Senator SHELDON: The surveillance cameras you've just been talking about at the Melbourne site have specifically been used to turn around and surveil workers on the basis of their union activity if they had problems or concerns about how the management was operating. Do you think that's an attack on human rights?

Mr O'Rourke: I've got no familiarity with those issues.

Senator SHELDON: If that were happening, would that be an abuse of human rights?

Mr O'Rourke: I can't comment with no context or full background. I'm just not familiar with the issues. I'm sorry. We're here, representing AWS, to talk about the use of AI models and services that we make available, and infrastructure that we build to make them available to Australian customers.

Senator SHELDON: Ms Foster, are you in a position to make any comments about the critical questions of not only the technology and the technical aspects—they are important, and I respect the answers you've given—but also the humanities question? That is really critical to the social licence question across this country, not only for Amazon but also, importantly, for how we evolve and how we develop and approach the use of AI.

Ms Foster: I can maybe give some context around how we think about responsible AI and responsible use, more broadly, of our services. I don't know that that will satisfy the exact question you've asked, but it may give you some context in terms of how we think about our services. We use responsible AI dimensions that are more around technical language: safety, veracity, controllability, fairness, explainability, robustness, privacy and security. So we evaluate our development of AI services using those dimensions, which can be technically defined.

There are another two policies that I can point you to that might be helpful—well, three things, actually. We have specific service terms as to permitted use and non-permitted use of our services and an acceptable-use policy which details that for all AWS cloud services. Then, to deal with some of the specific questions that come up around AI, and, especially with generative AI, we have a responsible AI policy that guides the circumstances under which we make our services available or not available. So I can point you to those policies, around what we view as acceptable use of our services.

Senator SHELDON: I'd be very happy for you to make those policies available. I'd draw this to your attention as well, and ask you, if you're able to, to give me a response on this. The policies that have been applied at the Melbourne site, and at a number of sites, have involved serious injuries to workers. Going to the question of the use of AI and the issues of safety, privacy and security, my deep concern is that, under these court process documents from Canada, Amazon is doing the direct opposite of what you've just outlined as policies, in how they approach this. There is not privacy. Workers have not been told that they are being surveilled. As to how they've been treated, it's like a digital prison without them knowing that the guard—AI—is oversighting them. In numerous cases, people have been terminated—certainly in the American examples, the Canadian examples and now in the Melbourne examples. It's a bit like they're using AI to hunt down the workforce that are pursuing their human rights. It's a bit like the workforce being hunted down by the Terminator. I'm deeply concerned about the humanities aspect of how AI is also used.

It's really important, and critical for social licence, that companies such as Amazon take responsibility for their direct workforce, because, if they can't do that for their direct workforce, then heaven forbid what they're allowing or encouraging to happen elsewhere. Amazon is turning its warehouses into digital prisons, in the way they operate. People are under surveillance—more like criminals. They're mums and dads with kids. This is deeply concerning to me. I think that the point about the humanities question and the social licence question really needs to be seriously answered by Amazon.

Thank you, Chair. I don't think there's much more I can add, but I'd certainly be happy to hear any comments from Mr O'Rourke or Ms Foster.

CHAIR: Mr O'Rourke, any final—

Mr O'Rourke: Nothing else.

CHAIR: Anything final from you, Ms Foster?

Ms Foster: Maybe I'd just to highlight that in many jurisdictions use cases in the area of employment are often an area where AI is regulated. From a policy standpoint, that's a trend that we are seeing around the world and may be of interest to the committee.

CHAIR: Thank you, Ms Foster. We are at time. Thank you very much for making yourselves available to answer the questions from committee members. If there has been any additional information requested on notice or additional information that you're able to provide the committee, please do so through the secretariat within the next two weeks.

COOK, Mr Peter, Business Member, Digital Economy Council of Australia

CORNICK, Ms Kate, Chief Executive Officer, Tech Council of Australia [by video link]

GOODEY, Ms Amy-Rose, Chief Executive Officer, Digital Economy Council of Australia

LONGCROFT, Mrs Lucinda, Director, Policy and Government Affairs, Tech Council of Australia

MONFORT, Mr Mark, Co-Chair, AI Working Group, Digital Economy Council of Australia

O'SULLIVAN, Mr Alec, Head, Operations, Digital Economy Council of Australia

ROSS, Mr Rick, Business Member, Digital Economy Council of Australia

WILLIAMS, Mr Timothy (Tim), Business Member, Digital Economy Council of Australia

[10:48]

CHAIR: Welcome. I understand that information on parliamentary privilege and the protection of witnesses has been provided to you all. I now welcome you to make a very short opening statement, at the conclusion of which I'll invite members of the committee to ask their questions. Maybe we'll start with the Digital Economy Council and then move over to the Tech Council.

Ms Goodey: Thank you, Chair and members of the committee. DECA represents businesses working across digital finance, digital assets, agentic commerce, AI, blockchain payments and digital platforms. Our submission marks 10 recommendations covering efficiencies, productivity, regulatory coordination, standards, accountability, sovereign capability and the practical challenges that Australia is going to face when deploying AI.

At the centre of our submission is a simple principle—regulation needs to focus on where AI creates real world consequences, not simply on the technology itself. Australia has already established legal regulatory frameworks. The challenge is ensuring they remain effective as AI becomes more capable, more autonomous and more deeply integrated into every Australian's life. This does bring a focus to DECA on agentic commerce and agentic AI. AI is moving from just engaging with Australians and assisting Australians to making decisions for Australians, particularly in their financial matters.

It does raise the important question: how do businesses and Australians know when they're engaging with an agent or with an AI entity? It then follows on to raise another question about sovereignty and how we deal with that in Australia. We've done this before in Australia when it comes to digital platforms, when it comes to communications, telcos, the internet itself. We're not new to this. We're very innovative here in Australia. We have a massive opportunity. Full credit to the committee, you've got a massive undertaking ahead of you. Good luck with that. I'll hand it over to my members for the questions.

CHAIR: To the Tech Council.

Ms Cornick: Thank you very much chair and members of the committee for the opportunity to be here today. Australia faces a fundamental choice: whether we want to become a global player in AI development and realise the socioeconomic opportunity that comes with that, or whether we want to settle for purchasing and using technology built overseas on terms set by others.

The Tech Council represents companies across the full technology stack from young startups, fast growing scale-ups, Australian-focused global businesses and major international technology firms. Our message today is clear—the opportunity is real, the window is finite and the decisions made in parliament are going to determine which path Australia takes. To best position Australia and our role in the AI economy, action is required on three fronts.

Firstly, we need an outcome on copyright reform. We need to ensure that firms can build AI in Australia on Australian data for Australian purposes. In doing so, we must safeguard our national interest and security and ensure positive outcomes not just for the tech industry but the creative industries and Australians more broadly.

Secondly, policy settings need to attract and retain investment. Australia's share of global AI venture is less than one per cent compared to 72 per cent in the United States. Policy settings in areas like tax, R&D incentives and employee share schemes all determine whether an AI company starts here, scales here and stays here.

Thirdly, regulation and governance must enable rather than stifle innovation. We need an approach that means AI can and will be developed and adopted safely here for the benefit of all Australians. Our sector supports risk-based proportionate regulation, serious obligations for high-risk users and streamlined pathways for lower risk ones.

In closing, I want to emphasise that collaboration between government and industry on matters like these will remain critical as we continue to navigate the complexities of an AI enabled economy, and it will be vital in ensuring safe and responsible adoption by Australians. Thank you for this opportunity to participate in this hearing and in the committee process. Lucinda and I are very happy to answer any of your questions.

CHAIR: Excellent. I might kick off. It flows on from some of the comments you've just made then in your opening remarks around, in reality, the choice that is in front of us and the opportunity that we have to make that choice that may not be with us too much longer. I wouldn't mind getting everyone's assessment on where we feel the timeline is.

But my question really is about the consequences of those two choices. We need to deal with the real and serious challenge that we've got around making sure we land something in the copyright space. But should we get through that—I'm also keen on your reflections on how that might be progressed. Have we got the ability for the training here? I take that point around Australia's opportunity to be a maker, not just a taker. I'd like you to reflect on what that would then mean more broadly, and, in terms of your respective members, what that means for them and our Australian community with a particular focus on the importance of sharing the benefits.

I'd also like to get your reflections on, should we not be able to do that, what in reality it looks like for that second outcome, where essentially our AI capacity is happening not on our shores but elsewhere and we simply continue consuming frontier and consuming that from somewhere else. I'll perhaps start with you, Ms Cornick, and then open it up to the panel. I'd really appreciate some broad reflections on the choice that's in front of Australians and reflections on how we best secure what you think is the best outcome for Australians. Also, frankly, what does it look like for both options that are in front of us?

Ms Cornick: We are strong supporters, unsurprisingly, of Australia becoming an AI maker. What that means is that we will see a flourishing of young tech startups that will hopefully grow into big global industries. It will see our industries being able to innovate here in Australia rather than having innovation imposed on them. It will mean creation of high-skill, high-value jobs. That does not just mean tech jobs. These firms are producing jobs in marketing, communications, finance, operations, sales et cetera. I think it's absolutely fundamental that we do this. The window is finite. As I made the point in our opening remarks, the reality is that we are part of a global movement in AI and we have a window right now to be a serious player in that, particularly in the specialised models, in the orchestration and the application layer. Australia has shown over the last decade that we are well and truly capable of growing global tech firms here. In my own experience, I've spent the last decade running the Victorian government startup agency. We grew the number of startups by 3½ thousand, from a thousand to 4½ thousand. Our sector created 67,000 jobs over that timeframe, and the valuation of those companies grew from $5 billion to $139 billion. That gives you a microcosm in one state of the impact that this sector could have.

In response to your question about what it means if we do nothing, I think that's a very unfortunate outcome for Australia. I think we will be worse off. Our industries will be worse off, and we will just become users of technology that's developed elsewhere. It means we're not thinking about how we're creating for our local context. It means we're not benefiting in the social and economic benefit that can be gained from AI. That's not to say that we must not approach this responsibly, and we very much support responsible development and making sure that the opportunity and risks are appropriately managed while we go about doing this. I'll end there and let others speak.

Mrs Longcroft: As Kate has said, this is a global development, a general-purpose technology that has genuinely transformative potential as well as disruptive ones, which is why that responsible adoption and management is so critical. But, as our research has shown, it has the potential to add up to $115 billion of benefit to our Australian economy and over 200,000 jobs just from generative AI alone. It is critical that we build that AI here and that we are not just absorbing technologies that are created under different laws, cultural values and systems overseas.

Your question, Chair, asked about how we ensure—and it is critical—that those benefits are enjoyed across Australian society, across all industries and across all people. That is a multifaceted question, and it depends upon a real partnership between industries that are developing the AI, both in Australia and globally—that's where strategic partnerships are important—and with the government and with workers and unions and others across our society. The small and medium-sized enterprise sector, for example, is over 90 per cent of our businesses. It is critical that they understand AI and know how to adopt it. At the moment, only a small proportion of them do. But research also shows that, if there were a 53 per cent adoption across SMEs, that would unlock a further $93 billion worth of opportunity across those small businesses, family businesses and regional businesses.

It's also critical that we develop skills in AI training for all Australians, and we are very keen. I know you've heard from Jobs and Skills Australia. We work with them, with unions and with others. There are wonderful

examples globally—in the UK, for example, where the government has rolled out free AI education for every British adult. Similar initiatives exist in Singapore and Finland that we look at and think we can adopt and learn from here to build that skilling across all out businesses and across our population, particularly for marginalised communities. We know that in regional Australia, for people like migrants, people with disabilities and Indigenous Australians, it is critical that they are skilled to understand, to discern their digital literacy and to adopt AI responsibly so that the benefits of this extraordinary technology can truly be enjoyed by all Australians.

Mr Monfort: To further add to that, in my role here as co-chair of DECA, I also work in the field. I am the chief AI officer at a law firm where we are actively adopting AI, and I've seen how that's transformed. What we were doing last year is very different to this year as more and more people get used to the capabilities AI has. In terms of the sovereign AI side of things and even copyright, the one message I want to get across is that we must act now. There is an opportunity, and the unique position that we play in the digital economy space is that we do not just bring AI tools but also bring blockchain tools that can also help in this space.

In our submission, we highlighted a workable test that helps us understand that it's not just the location alone because, for various reasons—and it was mentioned in this committee earlier in the week and last week—that you might have an AI model sitting here in Australia, but, firstly, you might not have the same capable frontier model—Microsoft, for example, was restricting a lot of what Australians had access to to what was called a GPT-4 model when the rest of the world was on GPT-5. We don't always get the best of the frontier models.

Secondly, the types of tests we look at cover various things from authority, continuity, portability, substitutability and recourse. With these things, it is not just where these AI tools are located. A good way to think about AI is as an engine. The AI models themselves are the engine, and there is also a capability for Australians—and we can allude to the examples later on—to build the rest of the car, the roads and the infrastructure that goes around that. There are many great Australian companies both here and overseas that are building all of that infrastructure to make Australia great in the AI space.

Furthermore, there is an opportunity to act upon these things. When it comes to things like copyright, being able to have recording permissions put on chain, preserving evidence and distributing royalties, blockchain technology and digital assets help us do a whole lot of that. I will defer now to my colleague Rick, from EY.

Mr Ross: I'd like to give some colour. You asked about some data points. It's really interesting. If you look back in history, we can always look back and look at the general-purpose technologies. We've been talking about that. General-purpose technologies are things that make economic impact. If you go back to the industrial revolutions, the first revolution was around steam, a significant impact. Then came electricity; then came semiconductors.

What's really interesting around each general-purpose technology as they come around is the pace of the cycle which organisations and technologies have to absorb that change. Between the first and second revolution, where it was steam, it was 100 years. Between the second and third, it was 70 years. Between electricity and semiconductors, it was 30 years. If you look at generative AI by itself, it took 18 months to absorb the capacity in terms of organisations using that. That should tell you something about the pace of change. It's so fast, and to do nothing means that we would fall behind as a nation. We, as a nation, have been inventors in the past. We invented wi-fi. We can look back at where we want to position ourselves. Do we want to be a laggard in this process or just follow other nations and see this change happening around us globally?

The things that concern me and things that should concern the committee are around keeping our talent in Australia. If there are no avenues for our people to upskill and build capability like startups and build within firms, they might go overseas. Providing them training is another avenue for us. People having the challenge around thinking their jobs are being displaced and giving them opportunities to train is a huge thing for them to resolve that concern they have in the back of their minds. So I want you to think about that in terms of those data points and also the opportunities that we have here.

Mr Williams: If it's okay, I might add to that from the perspective of a startup in Australia. I'm CEO of AstraSync AI, which is an agentic identity and trust platform. Effectively, the challenge of a startup in Australia is not a new one. You're out of the room for most of the conversations that occur and you see similar companies with inferior technology or inferior solutions getting greater opportunities because they're in the right rooms because they're in geographically favourable locations. The challenge that AI poses that exacerbates that is—a number of things. Firstly, as Mark mentioned, for availability of frontier technology, it's not just the frontier labs choosing to release models in different locations at different paces. You also have the government, effectively, of the United States banning the release of the Fable model globally. That's not just a corporate and enterprise decision; that is a government decision, and the potential that other governments may make similar decisions in the future and effectively isolate us from that technology is significant.

Secondly, for the surrounding infrastructure, if you look at something like agentic ecommerce, for example, schemes that allowed agentic shopping proof of concept and pilots to occur were whitelisted in North America last year. Here, in Australia, we are struggling to find the ability to get whitelisted. Stripe's solution for agentic shopping has been available in North America for about 12 months. It's still not available in Australia. So there are parts of the surrounding economy where we have companies in Australia that are willing to innovate and willing to participate, but they aren't able to, because of geographic restrictions. That limits us in terms of how we can scale, grow and be a global player, when we're competing against companies that are advantageous by virtue of location.

CHAIR: I'm conscious of time, so I'll pass to the deputy chair.

Senator DUNIAM: Thank you all very much for your testimony and submissions. I want to explore one area, which is how we define 'sovereignty', so sovereign AI and everything else that goes with it. Is there a test, a threshold or a set of criteria that you could recommend to us? Yesterday, I was discussing this with a government minister and some operators in the AI space. Globally, there doesn't seem to be a definition yet, so any free advice to this committee on how we chart that one out would be helpful.

Mr O'Sullivan: Sovereignty is not established by location alone. A workable test covers authority, continuity, portability, substitutability and recourse. Investment should concentrate where Australia holds deep industry expertise; valuable datasets; demanding operating environments, such as our mines; and established international relationships. That means agriculture; resources; health and medical research, which we're historically a leader at; advanced manufacturing; financial services; and the defence industry—sovereignty covers all of these sectors.

Senator DUNIAM: I'm in the business of trying to get, as best as we possibly can, some degree of certainty around what sovereignty is, because the people we represent—and we get this in media interviews all the time—will say, 'When you say "sovereign AI", what do you actually mean?' and I need to be able to say, 'This is what it means.' I understand your point about these being the things that we as a nation are looking at and have strengths in et cetera, but when setting out the conditions for the industry, if I can call it that, in our country moving forward what should we be saying quantifies or amounts to sovereign AI?

Mrs Longcroft: Kate Cornick might be able to address that from a 'what is an Australian or what is a sovereign AI company'—I might then also add to that in terms of sovereign capability, because I fully agree with you. The word 'sovereignty' means something different to every person who uses it.

Ms Cornick: Thank you, Lucinda. The first thing to note is that sovereignty exists across the AI stack, and it exists in the context of supply chain and what does it mean. I think, as Lucinda alluded to, it comes back to the definition of an Australian technology company that is going to build our capability in Australia across the entire stack. So we certainly think about Australian technology companies being central to that—an Australian technology company being first registered in Australia, using technology to deliver their products and services. You could see that existing, for example, in the data centre environment, with an Australian company delivering data centres within Australia. You can see that in specialised models in the context of Australian technology being developed in specialised models. You can also see it in the orchestration layer. You can also see that in the application layer. I'll throw back to Lucinda for some additional commentary.

Mrs Longcroft: Thank you, Kate. We do field this question a great deal also from the Australian tech industry perspective. What is sovereign capability? I tend to frame it as national capability. How does Australia assess its national capability to play where it needs to play in the national interest? This is why the national AI standards that the Prime Minister has announced and that your committee is informing are so important and why we support them. They look to assess what is a capability in AI that Australia needs to be resilient, to be economically prosperous and to serve the needs of the Australian community. That will be a combination of AI developed here by native Australian companies—that is so critically important to our innovation ecosystem—but also through strategic partnerships where there is AI capability that we do not have the resource or the technical specificity to develop here, where we can develop strategic partnerships with global leaders. Australia is uniquely well placed in that regard, given our alliance strategically through AUKUS, through the Five Eyes partnership, through our data alliances and others, as well as our geopolitical position to attract those sort of strategic partnerships and build a local and Australian tech ecosystem that serves the national interests.

Senator DUNIAM: That does sort of bear out the point that Mr O'Sullivan was making; it's not just about location.

Mr Monfort: May I add a further thing? Going back to the analogy—I like to think in analogies here.

Senator DUNIAM: So do we!

Mr Monfort: The AI model as being the car engine is one thing, but all the rest of it—the rules and regulations being how we build out the infrastructure and the roads to make cars possible. When the first cars came about, they were on cobbled roads, and people complained that they didn't work that well. But as soon as they paved the roads, not only did we get cars that work better but we got buses, motorbikes and all of these other things. As we learnt over time with cars, we didn't have seatbelts initially, but we did later on. So the safety that comes from that is important to think about as well.

All of these different elements—it's not just the models as to where Australia can have and play a sovereign role. We can have sovereign AI models; like I said, we certainly need it. In my space, working at a law firm, I am relying on Azure, AWS and even what Google provides in terms of the models, because not only do I need to keep data here but I need to keep the models and the calls that go to the AI models here as well. A lot of people don't look at the underlying things under the hood to really see where these things sit.

There are key examples of where Australians are playing a vital role in the whole stack of AI. The model space is one thing. If you're using open source models, to be able to use them well, you need to train them, and one of the best tools for doing training out there in the world is built by an Australian company called Unsloth. They're based in San Francisco, though, as they've had to move to actually make themselves available.

The other one is that there are great examples of Australia coming together. Our experience from the digital asset space—you may be familiar with the RBA's Project Acacia, which called on the Digital Finance CRC to come together to create something where you have government, industry, start-ups and big banks all coming together to actually test the technology. Perhaps there is a call there for an AI CRC as well. Further on sovereignty, Amy-Rose.

Ms Goodey: I think we need to also think about AI sovereignty and sovereign AI. I think people are using them interchangeably. We might need to distinguish between Australia having capability. You can imagine a hospital using a particular model. What if that model breaks? Does Australia have a fallback to use its own capability, its own infrastructure, to be able to service some of our core infrastructure? That is Australia acting in a sovereign capability, but it has frameworks around it. We write the rules. We govern. It speaks to some of the suggestions we made in our submission. We wanted to make sure that we have the ability to withstand any changes that impact us, so it's a bit of a defence question as well, and I'm sure Defence is well across it.

I might just also draw on my telco background, because the internet is such a prime example of how Australia built this fantastic sovereign infrastructure around the internet itself. It started in 1986 with the au domains, and then we went on to build auDA, which was a representative of the internet sector, and we built and maintained capability and sovereignty around the infrastructure itself. We built a whole bunch of ISPs from the ground up and from backyards and sheds. Now we have some of the biggest service providers in the world—iiNet, for example—and everything is interconnected. We have our cables, and we have built that for 50 years now.

Mr Williams: I'd just add that we've obviously talked about the economic reasons why you would go down this path but, beyond that, there are actually quite strong cultural and value driven reasons why you would do it. If you look at other countries that have gone down the path of sovereign AI, you've got Manus in the EU and you've got a number of Chinese labs, and they've sprung up for different reasons. Obviously, China's response was largely geopolitical, given the risk of being locked out of using frontier models. But the European model is probably one that's more interesting and insightful for us, given that that sprang up to ensure that AI reflected European values and it enabled the EU to continue to operate with a leading model when other labs were starting to leave because of the controls and regulations that were being put in place. So that gave them control over the ability to regulate what was acceptable for AI in their particular jurisdictions. If we don't have sovereign AI, we don't have that optionality, and we're effectively at the mercy of the frontier labs and the frameworks that they'll choose to operate in, which will largely be the United States.

Senator DUNIAM: Thanks very much.

CHAIR: Ms Chaney.

Ms CHANEY: I'm interested in your views on law or legal obligations versus guidance. You recommend that higher risk uses be addressed through regulatory guidance, which is obviously less enforceable or not enforceable, unlike legislation. Which AI harms, if any, does the Tech Council accept should be subject to binding legal obligations with penalties? Where do you draw the line in terms of guidance versus legal obligations?

Mrs Longcroft: Thank you, Ms Cheney. It's a critical question, because we know that the level of Australians' confidence in AI is among the lowest certainly in our region and globally, and indications are that effective governance and management of harms and risk will develop that confidence and enable the adoption and the diffusion of the technology, which will deliver the benefits to Australia. So that governance is critical.

We at the Tech Council represent, of course, a wide range of tech industry, and across that industry there is a broad understanding and agreement that effective, proportionate, risk based regulation is necessary, and it's necessary to develop effective regulation. That is developed most effectively in consultation or by co-design between industry and government to ensure that government and the regulators have full access to the information necessary and accountability for those players to be able to enforce those regulations.

Your question asked about what should be mandatory and what should be guidance, and an effective governance mechanism combines both. There must be internal self-governance and accountability by the companies and then accountability to the regulators. There is an effective and broad range of laws already in existence that regulate all technologies within Australia, including AI. They span a broad range, from privacy to consumer law to competition law, tax and other issues. There are regulators well equipped to deliver on those. What we in the Tech Council say is that AI is a novel technology and there may be gaps, so we propose a gap analysis to identify, in an ongoing way, where those laws may not yet be effective and where they need to be updated. In terms of the management of evolving security risks relating to AI, we say that it is important to identify the different roles that developers and deployers of AI play in the full stack between the development, evaluation and training of AI models on the one hand and the deployment of AI across the economy on the other. There are different players in the stack who should each be held accountable according to the severity of the risk or harm that could emanate in a particular occasion. We have pre-existing harm-assessment tools and systems for reporting, for example, and for accountability to manage that proportionate risk across the spectrum of actual harms.

Ms CHANEY: I'm interested in where you see the line between internal self-governance and accountability to regulators on and things like agentic AI breaches of confidential data sets, AI hallucinations, and chatbots that might aid children in committing suicide. Which things are on which side of the line? As you say, there's an issue with trust, and the Tech Policy Design Institute's research finds 70 per cent of Australians would be more comfortable with AI if strong rules were in place. Where do you think those strong rules are and what's on the wrong side of the line from your perspective?

Mrs Longcroft: I will try and give a fairly high-level answer given the time that is available. You drew on a number of examples that are critically important to Australian society in the online safety space, for example, where chatbots or other AI tools might cause harm to Australian children or Australian families. The eSafety Commission has a very well-developed framework of online safety that is currently being updated in response to those harms to consumers, and they are strictly enforceable laws. In fact, very recently, the Australian government increased the penalties and the capacity of the eSafety Commissioner to enforce those laws. We certainly support that accountability and governance in that frame or that level of harm.

Across the spectrum, this is a rapidly evolving area of technology, and that's why it is so critically important that our institutions are well-resourced. We see the AI Safety Institute, for example, as the hub of an AI safety framework within the Australian government, one we strongly support. That is, as has been mentioned by a number of other witnesses, part of a global network of AI safety institutions. There is work taking place globally. These harms are not just being experienced in Australia, and we would strongly endorse the engagement of the AI Safety Institute and of our government in globally consistent frameworks that respond to the severity of the harms in any particular case.

Ms CHANEY: What do you think the AI Safety Institute needs in order to play that role? Has it got what it needs or are there gaps?

Mrs Longcroft: In our submission we have strongly endorsed the AI Safety Institute as the core of the Australian government's response, and an area where the industry will seek to engage closely. We advocate that they be well-resourced. We wouldn't assume the level of resourcing that is required to do that, but they should be well-resourced and well-connected into industry. We're very encouraged by Dr Conroy at the head of the AI Safety Institute, and her teams work closely with industry through memorandums of understanding and sharing of information to ensure that they are able to do their work.

The more engagement and strategic partnership that we have with AI developers both in Australia and globally, the more effectively those mechanisms in the Australian government can police, manage and monitor, and ensure that Australians are kept safe from harms, and that they can enjoy the benefits that AI technology can deliver to us.

Ms CHANEY: You say we need to do a gap analysis on the laws, because there will be some gaps, but you also seem to be saying that there's already a framework with the eSafety Commission that works well. Can you think of any examples where there is a gap in our regulatory framework to deal with novel risks? Agentic AI is

not something our law has had to deal with before in terms of intent and attribution. What are some examples you see where there are gaps?

Mrs Longcroft: You're absolutely right. The Tech Council of Australia works with the Australian government's Executive Cyber Council, and through emerging-tech working groups, to look at these questions. Agentic AI and the challenges that they pose to our cybersecurity in particular is an area of particular concern to the Australian tech industry and to our engagement with Australia's security agencies and intelligence agencies.

Mr O'Sullivan: If I may, our submission identified three gaps that warrant targeted reform. The first major gap was in the Privacy Act, which retains qualified small-business and employee record exemptions. Turnover does not reflect a business's capacity to profile individuals, and employee information can be used well beyond conventional record administration.

The second gap that we've identified is in the automated decision-making transparency provisions commencing on 10 December this year, which require specified privacy policy disclosures but do not themselves establish a general right to an individual explanation or human reconsideration. A person may learn from a privacy policy that automation is being used without being able to establish whether an error affected their particular outcome, which goes to the point that some of these frontier labs are closed boxes. We don't know what they're doing. When we made this submission about a month ago, the majority of AI activity was done with frontier labs—closed boxes. That has since flipped to open-weight models, which are available for anyone to inspect the code, which also lends to this sovereign capability argument. If we have data centres onshore, within our jurisdiction, running open-weight models, then this can be a very transparent basis for sovereign capability.

Ms CHANEY: That was two—the Privacy Act and automated decision-making. Was there a third?

Mr O'Sullivan: I do have a third. The third is civil liability principles. They require proof of the elements of a cause of action. If you don't have visibility over the LLM you're using, then how can you have visibility as a board of directors? You simply can't. You need to be using an open-weight model as an entity.

Ms CHANEY: Thank you.

CHAIR: Senator Sheldon.

Senator SHELDON: You mentioned, of course, the opportunities for employment and jobs into the future. It's a touch/feel sort of thing. Where—in what parts of the nation—do you see those opportunities congregating? I'm interested to know if there's anyone from Canberra as well.

Mrs Longcroft: You ask a wonderful question, because at the moment we know that AI is disrupting. It's transforming our economy, but we don't yet—and nobody can yet—know exactly where that is taking us. What we do know from past industrial revolutions and the results of technology is that new jobs, new tasks and new workflows open up. In fact, the majority of Australian workers do indicate that they believe AI will enable them to do their jobs better. In fact, only 15 per cent of Australian workers believe that AI will replace them. That is a very genuine and understandable concern.

How do we identify the jobs of the future? It was interesting; in a recent discussion I asked the head of the automated trucking association whether they were concerned about autonomous driving displacing the jobs of truck drivers, and I was surprised by the answer, which was, 'Currently, we have a shortage of truck drivers.' So in fact this is a problem that needs to be addressed, and it will create new jobs because you will need custodians and stewards of cargo that will be required to manage the loading and unloading and the development of this new work. That was one example, but I think it is incumbent on the tech industry to work better to help equip Australians for the jobs of the future.

The last thing I'd say on that is that Australian workers are gaining, on average, about 187 new skills—that's certainly in the last five years—and they are 2½ times more likely to be human skills. They are skills of judgement, of communication and of customer relationships—the areas where AI can't effectively do what a human can do. We need to lean in to building the ability both to manage the technology and to augment what we do as humans so effectively for rewarding and fulfilling jobs in the future.

Ms Cornick: Can I just jump in here to just call out that our submission proposed close monitoring of the sector. As Lucinda rightly points out, change is happening in the economy and, regarding the impacts of AI, particularly on the workforce, we need to make sure that we've got evidence to make the right decisions. Having that evidence base is really critical, so we very strongly support timely monitoring of the Australian workforce to make sure that we can create evidence based decisions on that.

Mr Cook: My company is a stablecoin company. We've created an Australian dollar stablecoin. We've got 17 staff, two of which, in a training sense, are full-time staff but also doing PhDs by paperwork. We've got seven

staff that are software developers. In, I'll say, a pre-AI world, I would have only two or three developers in Australia and I would have offshored for maybe 10 or 15 developers in India or the Philippines or wherever. In fact, Tim, sitting next to me, is in the same situation, where we're running, let's call them, very specialised development teams in Australia rather than offshoring. Then, the rest of our team, including me, a man in my 60s, are all absolutely upskilling in AI. I feel like I'm back at school, yet really understanding the capability and how we can, in the future, essentially leverage our young people to create more in Australia is just part of my job of being on a board of directors.

We're at the application layer. Just going back to the question on sovereignty, the application we build creates a sovereignty impact around the Australian currency. There's a whole lot of writing about the risk of US dollarisation globally. So we're using these tools to create a business to support education, support local staffing and, in fact, in our case, we think, support a sovereignty issue.

Mr Williams: I would add that my background is that I've run large people teams that have been subject to automation—largely in contact centres and customer admin—and the thought of automation removing jobs is not a new one. I've been commercialising AI in those functions for 20 years and I'm yet to fire someone because AI replaced their job, because there is always more work to do. There is not an organisation in Australia that is lacking work. The greater risk over the last 20 years has been offshoring in lower-cost jurisdictions, and how we've fought that traditionally has been in training and upskilling. If you look at programs of the past, such as RTO based programs, where we invested in organisations upskilling and training their people, oftentimes that led to re-onshoring. If you look at organisations such as Telstra and others, they actually brought jobs back, because, as a people and as a government, we invested in upskilling our people to deliver better service and we saw that the ROI on that would offset the arbitrage of lower-cost jurisdictional resourcing. So this is a challenge we have faced before and can solve again.

CHAIR: Excellent. One final question.

Senator SHELDON: Thank you, Mr Cook and Mr Williams. I think that was really helpful. I've heard your description before of onshoring of jobs, particularly in the AI space, as well as in a number of other companies, which has been enlightening. I just want to go back to the Tech Council of Australia. It's always dangerous asking a former Transport Workers' Union official—an ex-garbo—about trucking, but I certainly respect the advantages that automation brings. There is in the Transport Workers' Union submission a proposal about how to deal with fully automated vehicles. Partially automated has a huge advantage in both safety and opportunity. Fully automated raises some serious questions, with massive job loss, particularly if it starts in the cab or rideshare industry, for example. They've raised some quite serious questions: What are the actual advantages for Australia? What is the tax lost? What is the job loss when you're looking at that segment? What is the transition of that workforce to somewhere else? That is probably at the heart of the concerns across the community, which I'm sure we all hold in some shape or fashion. I hold them very deeply. You might want to take on notice the one on the TWU submission, because I'm quite interested in what your view is, but you may have some immediate views about fully automated vehicles.

Ms Cornick: Senator, I can jump in here. I've certainly seen some of these automated vehicles in action and certainly understand some of the concerns that exist there. I think, as always, it's the trade-off between the opportunity and the risk. In this case, I think there are multiple risks—and this is my personal view on the ledger. I think there's obviously the risk to the workforce that exists, but there's also the risk that this technology gets developed and rolled out elsewhere and then just comes and waves over us. I think we've seen that in the taxi industry, through the creation of Uber, for example, which is now a global firm managing huge swathes of taxi services, including in Australia. My immediate reaction would be that we absolutely have to look at the risks of the workforce. We also have to look at the risks of 'what if we lose control of the sovereignty around that issue such that a global firm comes and imposes on Australia?' That's something that we certainly see can create enormous opportunity for innovation in Australia—to be a leader in that space. But there are also challenges when companies come into market, and we need to be very aware of that. I will leave it there, and we're happy to review the TWU submission in full and come back to you with further thoughts.

Mrs Longcroft: And, Senator, it speaks to the complexity of the job you, and the committee, have undertaken in this regard, because those autonomous innovations and technologies, as you say, do disrupt deeply and cause challenges and problems for a large workforce, and a valued part of our workforce, while those technologies also offer enormous benefits for persons with disabilities, with reducing safety risks, and others. And that broad question applies to the whole of AI: how do we leverage those benefits and opportunities and, as Kate has said, make sure that they don't bypass Australia but do so in ways that manage our workforce, that are respectful and that acknowledge our human rights and our needs, as humans, to drive the benefits for and in Australia's interests?

CHAIR: Excellent. I will need to call time on this session. Thank you so very much for making ourselves available to appear before our committee today and answer our questions. If there's been anything that you've indicated you'd be keen to provide on notice, or any additional information for the committee's consideration, please do so through the secretariat over the next two weeks. Thank you very much once again.

QUINLAN, Emeritus Professor Michael, Private capacity [by audio link]

WILLIAMS, Professor Penny, Professor, Queensland University of Technology [by video link]

[11:38]

CHAIR: Welcome. I understand that information on parliamentary privilege and the protection of witnesses giving evidence to this inquiry has been provided to you. Do you have any comments to make on the capacity in which you appear?

Prof. Quinlan: I worked at the University of New South Wales, but I'm appearing in a private capacity as a health and safety expert and researcher.

Prof. Williams: I am from the School of Management at the Queensland University of Technology. I'm also the Director of the Centre for Decent Work & Industry, which is a research centre based at QUT.

CHAIR: Thank you very much to you both. I'd like to now welcome some opening remarks from both of you, should you wish, at the conclusion of which we'll move to questions.

Prof. Williams: Thank you. Because our submission was a joint submission provided by both Michael and myself, I'll make just one opening statement on behalf of both of us, if that's okay.

CHAIR: Sounds fantastic. Thank you.

Prof. Williams: We would first of all like to thank the committee for the invitation to give evidence at this public hearing. And as outlined in our submission, Professor Quinlan and I have each been involved variously in different recent research projects that explore the adoption and impact of AI and algorithmic management on workers. This has included things like a national survey mapping participation in digital platform work in Australia, which is a context in which algorithmic management is prevalent, Industrial Transformation Training Centres supporting the adoption of collaborative robots in manufacturing—here, we note the integration of AI is dramatically increasing robotic capability—an ARC funded project to explore how AI and algorithms are reshaping human resource management and a related project on the use of AI in hiring processes.

Currently, Professor Quinlan and I are working together with international researchers on a project based out of the Karolinska Institute in Sweden. This project looks specifically at the occupational health and safety implications of algorithmic management, and it aims to address its negative impacts. Our submission and evidence today draw from this research base with a particular focus on how AI impacts on health, safety and wellbeing of workers.

To this end, I'd actually like to clarify what we mean when we speak of AI and, relatedly, algorithmic management. Much of the focus has clearly been on generative AI, but earlier forms of predictive or rules based AI that use algorithms have been changing workplaces for some time, and we've been monitoring this. They underpin robotic automation and the new methods of management that we first saw emerge in the gig economy but are now commonplace in Australian organisations. That is where algorithms match workers with tasks, monitor and track their performance and automate decisions that were once the role of a human supervisor. Generative AI has vastly extended the capabilities of these technologies to significantly transform work.

We would say that research on these changes is limited, but our early findings suggest that, while AI has many potential benefits, such as improving safety, it can also have profound adverse effects on work conditions and the health and wellbeing of workers, including exacerbating existing psychosocial hazards. For example, research suggests that heightened surveillance and tighter management control increase worker stress and work intensification, especially when linked to performance and reward systems, that workers may take physical risks to achieve AI driven targets or that new physical hazards may emerge from software malfunctions or poor AI judgement.

Our submission draws attention to the need for a national approach to regulation with stronger integration between any national AI legislation and existing labour protection regimes, as well as a reconfiguration of occupational health and safety regulation to incorporate digital work systems. We also note that, while the EU AI Act operates as a model for Australia, it is limited when addressing occupational health and safety. Secondly, while Australia can learn from the evidence gathered through international projects such as our work with the Karolinska Institute, more local research is urgently needed to understand the risks specific to Australian workplaces.

I'd just like to conclude by saying: we believe that Australia has, through its national workplace relations system, a unique opportunity to proactively address the risks AI poses in workplaces, learning from the gaps that are in the EU AI Act but also leveraging institutions such as Safe Work Australia and the Fair Work Commission to establish robust, industry-wide safeguards for the responsible adoption of AI. Thank you.

CHAIR: Excellent. Thank you very much for that. I'd like to get your reflections on something, given your work in this space. We had evidence provided to us the other week by representatives of the SDA. In their submission they gave an example of an Aldi logistics warehouse where efficiency gains were made around an algorithm for the most efficient way to pack within the warehouse. It was devised by, essentially, an AI algorithm and was mathematically perfect—fantastic—but didn't take into account what it was like to actually do the job on the ground, in the actual warehouse. So I just wanted to get your reflections a bit further around this concept of the imagined job versus the job in real life and how we make sure we've got the capability to look at ways in which we can genuinely see AI deliver greater efficiencies and greater safety for workers. We're seeing more and more of these instances where the algorithm has done this remarkable job and notionally delivered these efficiencies but then not taken into account the human doing the work, in terms of health and safety, psychosocial risks and things like that. I just wanted to get your reflections on this concept of the imagined job versus the job in reality.

Prof. Williams: I'll go first, and then Michael might like to add to this. Certainly in my own research I've seen similar issues reported by workers that I've spoken to. What appears to happen in many cases is that the algorithm determines at what pace a worker should work or how many tasks should be completed, and it monitors that and reports on that often in a way that creates a leaderboard for managers to look at. The problem with that is that it can adjust the required level of performance and gradually push it up. The evidence that we're seeing emerge is that that has a number of effects.

Firstly, it can increase work intensification. The constant monitoring that comes with that makes workers feel that, firstly, they're not trusted. They also increase the pace at which they're working because they need to meet this higher standard. In the worst-case scenarios—and this has certainly been reported in warehouses in the US—it also results in contracts not being renewed or workers being terminated if they fall below par. The difficulty with that is that some workers might have a bad day. They might be unwell. There might be legitimate reasons why their performance is not up to standard. Where algorithmic management is at its extreme, there might be very limited opportunities for workers to have a voice to explain why they're performance is subpar at that point in time.

The other issue that's problematic with this is that workers reported to us in our research that often the quality of their work seems to be overlooked. So the shift has moved from balancing not just work output but the quality of that work output to measuring volume of work output and ignoring the quality of that. So some workers feel that they can't produce work to the standard that they did before—certainly retail workers were reporting that to us—or that they can't complete a task before they need to move on to the next one.

Where I've seen this work effectively versus where it seems to work less effectively is where there is a genuine consultation with workers and a review process. Some of my research has suggested that, while many organisations have introduced various forms of AI and algorithmic management, fewer of them have done a review to determine what the effects have been of that on workers and what the unintended consequences of that have been.

However, certainly the research that we've been doing with the Karolinska Institutet has shown that, where consultation has occurred early on and workers have had a voice on how this can be best managed and where the implementation of it is consistent with a culture that supports worker safety and meaningful work, there are fewer issues associated with the use of algorithmic management. I don't know if you want to add to that, Michael.

Prof. Quinlan: I'll just add a couple of things. The first thing is that risk assessment and consultation prior to the introduction of new technology is limited. It's in the legislation, but my research in mining and other industries indicates that it almost never occurs in practice, even though it's a legislative requirement. Secondly, I'd reinforce Penny's point that in reality the introduction of these systems is tending to place increased pressure, mental and physical, on workers, not just in warehouses but in retailing and in transport. It doesn't take account of very experienced workers in transport, who may have their own routes best worked out. It will dictate a route to them, which they find very annoying.

Finally, I would like to add that there are actually physical risks associated with these regimes. Let me take the case of autonomous vehicles, which was raised in the last session. The International Association of Transportation Regulators has raised a number of serious safety issues in relation to autonomous vehicles, and I'll just give you three. There are more. One is that they raise the issue of uncalibrated and unapproved aftermarket modifications of equipment, including changes of windshields, which affected the operability of those vehicles. Software changes can affect the braking and lane-keeping performance of this equipment, and it's coming in next year. And there can be sensor calibration oversight. So, once you introduce these systems, there are a whole series of maintenance and oversight issues, and it is naive to pretend that they're all going to be dealt with effectively.

The second area is blocking emergency vehicles when AVs get into a funk. California has actually had to pass a law dealing with that. The third one and, I think, the most frightening one that was released on 12 October—that's just a few days ago—is that the US government is now doing a review of Tesla's robotaxis because it's found that they were speeding through school zones. The other frightening aspect of this is that I'm the chair of the New South Wales SafeWork advisory council. I'm not appearing in that capacity, but my inquiries of that and other bodies indicate that no OHS regulator was actually approached or consulted in the introduction of autonomous vehicles.

CHAIR: Just on that, where then do we have the ability for the government to have the most impact when it comes to AI and algorithmic management of work?

Prof. Quinlan: It comes in three areas. One is in the AI legislation that will be introduced, which should clearly seek to interact or align with OHS legislation and set broad duties in terms of design of products. That's the sort of approach that legislation will take. But it needs to recognise that it needs to integrate with other legislation. The second body of legislation is occupational health and safety legislation. At present in Australia, while the mining act in one state refers to remotely operated operations, only one state has legislation dealing with digital work systems at present. That is New South Wales. It's a good starting point. It has a good definition of digital work systems, and it allows authorised officers to access the information in the algorithms when they have due concerns and when there's a guidance note in that area. This is in some ways parallel to what's happening in Europe with their guidance material for inspectors. So that's the second level.

The final level is the workplace industrial relations system. You can't negotiate over these changes at enterprise level and even at occupational health and safety level. They can produce codes, but there's a need for industry-level negotiations about the transition to new technologies, work systems and their safety implications. Australia is particularly advantaged. We have the Fair Work Commission, which can make rulings in this area and which is already dealing with health and safety issues through some of the legislation it's introducing now and where all the parties get a say, and you can come up with sensible industry packages that the industry can lead with. If you try and do it at enterprise level, the company is going to reasonably say, 'Well, we can do this, but we'll be undercut by all our competitors.' So those are the three levels, I would argue, that we've got for addressing this issue.

CHAIR: Senator Sheldon.

Senator SHELDON: Thanks for joining us. I'm not sure if you heard the evidence from one of the previous witnesses, Mr Alec O'Sullivan, Head of Operations at the Digital Economy Council of Australia. He was asked to name a number of gaps, and one of the ones that struck me was this. The first major gap he identified—and I'll use his words—was in the Privacy Act, which, he said, 'retains qualified small-business and employee record exemptions'. He went on to say:

> Turnover does not reflect a business's capacity to profile individuals, and employee information can be used well beyond conventional record administration.

It goes to that question—the gap that he identified. But also, is there a gap in the case of the occupational health and safety act? And are you able to venture any potential gaps in the workplace relations act?

Prof. Quinlan: With regard to the health and safety act: as the New South Wales legislation indicates, you have to get access to the algorithms. If you've got concerns, you need to be able to access the algorithms. This will mean authorised officers or union officials will require the expertise to do that and have grounds. The guidance material in New South Wales is setting out the grounds for which that can occur, and that's accepted. Why do they need that? Because, in Italy, where they got access to an algorithm—I think it's in our submission—they found the algorithm was tracking people in their non-working time. This is not just a privacy issue; this is a health and safety issue. If you're tracking people and they're not aware of it—or even if they are aware of it, worse—this is going to affect their behaviour.

You can't really deal with these arrangements unless you can get inside the algorithms themselves. The danger with AI, as Penny knows, is that they can rewrite algorithms as they go. So that becomes even more opaque, over time.

With regard to Fair Work, I think they are already developing the expertise in this space. They're looking at the evidence. We have a project currently running. I'm a director of TEACHO, a transport safety body. We're running a research project on autonomous vehicles right now, looking at their safety implications, and we're running another one dealing with the impacts of gig work on the mental health of workers. This is the sort of evidence that the Fair Work Commission would use in making decisions. As I said, the OHS regulator can do important things.

It can set up guidelines. It can set up codes. But there's a level of specificity and a level that needs to be negotiated at industry and sector level, and that's where that third tier comes in.

The other point I would make, which is in our paper very much, is that what the AI act shows—and this is very consistent with what I said about autonomous vehicles—is that it doesn't deal with the post-implementation stage; it only deals with the design stage. The trouble is: that's too generic. Risk cannot be anticipated entirely at the design stage. They're also specific for industry and workplaces, so you can get unanticipated responses that encourage workers to take risks. For instance, some of the surveillance equipment on trucks on mining sites actually gets drivers to look in a particular direction all the time, and they're arguing that that is actually distracting, because they know they're being recorded. Also where you put the emergency or the safety related material, on a set of ergonomic screens, can also have safety implications, if they're right down at your feet or right above and people are operating multiple screens. So there are all these post-design, introduction-phase things. Then companies can modify their systems after they have already been introduced, or they can bypass the systems even if there are safeguards built into them and that places the onus then on the regulator in identifying those bypasses and actually doing something about it—or until somebody gets killed. And, as we point out in our submission, software flaws have already killed people on worksites in Australia.

Prof. Williams: I would just add that I do believe that there are gaps in workplace surveillance regulation when it comes to these types of technologies. The surveillance and monitoring component, as Michael mentioned, is one of the key problematic aspects. That contributes to unsafe work practices or psychosocial risks and is heavily linked to how that data is used and managed. So I think there are multiple regulatory levers that need review and need to be considered in managing this.

The other point that Michael touched on, which I think is really important to consider here, is that it is not just the technology that's problematic but the way it is used. In many cases, various forms of AI have been introduced to the benefit of workers, but it can also be used punitively. So it isn't a straight line to say it's the developer's responsibility, for example, or that it's the technology and the algorithm itself that is the issue. It's also the responsibility of how the organisation and managers within the organisation use that technology and if it is used in the way that it was designed or supposed to be used.

Senator SHELDON: I know we have to go to somebody else shortly. Professor Williams, you said there are regulatory levers that could be considered to be used. Are you able to outline some of those potential regulatory levers? You may want to take a little bit of it on notice. I know you've touched on it in your submission. I will follow up part of that from the comments you made before. Where do you think there are some shining examples of where we are thinking about only the technical part of the AI rollout? When I say 'we', that could be we as a committee or we as a community. Are we too focused on the technical piece, and do we have to be more focused on the humanity piece? Is that something that we need to be considering as well when making these judgements?

Prof. Williams: I will take your question on notice in relation to the regulatory levers. I think there's quite a bit to say there, and Michael might want to jump in there. But, in relation to the focus on the humanity piece, I wouldn't say that we are too focused on the technology. I think it's quite critical that we focus on that right now, given the rapid expansion of the technological capability. But we can't lose sight of the human role in applying and adopting that technology and the human implications if it's adopted incorrectly. So I think we need to balance what we can do about regulating or controlling the development of technology with how it is used by organisations. I think—to Michael's point—occupational health and safety provides us with a nice avenue in which to balance those two perspectives—you know, what are the outcomes and impacts on workers and who's responsible for that, versus what is the capability of technology. Michael, perhaps you want to—

Prof. Quinlan: Yes. Apart from actually bringing Safe Work into the conversation in aligning AI legislation nationally—it is already doing a national uniformity report, I understand, if it isn't already submitted—the second thing is that another lever would be that we need a national framework of OHS legislation which is integrated with this area. New South Wales's was a very good starting model, and we should never level down. The regulator in various states can produce guidance material at industry and sector level in relation to this. In Sweden and Finland, they're issuing inspectors with a set of tools to conduct their inspections in relation to AI. We could learn and borrow that sort of idea, and that's part of our research project. That's going to be very informed. So that's really important. We could do that too, and we should be doing that too, because inspectors need the skill set to know what to look for, and they can ask and do the surveys. We've actually got the tools there, and I'd be happy to provide more advice in that terrain.

The third level is the one I recommended with Fair Work. Fair Work has considered technological change in the past. In fact, there were some important New South Wales cases on engineered standards in the warehouse industries about 20 years ago. So it is a body that's able to consider these things. Again, we need that negotiation

going on at that level. Amongst other things, that could come to agreements about the transition into new technologies, about the pace of introducing, for instance, autonomous vehicles or automated systems into particular industries and about how that would be managed—also to protect the working conditions and other aspects of that and to deal with performance pressures. So that's another avenue for dealing with that.

We've already got, as I said, some evidence that that can be used in that regard. We're very advantaged. Other countries don't have that body that can do those sorts of things. The Fair Work Commission is very evidence driven, as Penny knows, so it's an ideal body to take the interests of the industry and unions and other parties and come to some really flexible and fluid agreements that enable us to address technology without incorporating some of the costs, which will happen if we just let it happen.

CHAIR: I'll give the final question to Ms Mascarenhas.

Ms MASCARENHAS: I am interested in social licence in particular. I know that the mining industry has had some difficult lessons but has improved frameworks over time. My questions are: How do you see the maturity of data centres and the AI industry with relation to social licence? Do you have a suggestion of the top two to three things that the federal government should do to make sure that we see minimum standards?

Prof. Williams: Could I ask you to repeat the question. Is it about minimum standards for data centres?

Ms MASCARENHAS: Yes, it's about minimum standards for data centres and artificial intelligence, but from the perspective of making sure that we maintain and earn social licence, and lessons from other industries, such as the mining industry.

Prof. Williams: Michael might have some views on mining because he's worked more deeply in that than I have. I do have to say that I think your question about social licence and data centres is beyond the scope of the research that I tend to focus on. I also think social licence is something that has to be earned. I think starting with the need for some kind of guardrail is important, but I couldn't answer the question on minimum standards at this point in time.

Prof. Quinlan: If we take the mining industry as an example, there are environmental and other standards which the mining industry has to abide by. It doesn't always, but it does have to abide by them. So I guess, with data centres, there could be a framework set up nationally and at state level which sets a series of criteria that must be met for data centres to be established in a particular area, and that could involve a degree of community consultation and also place an onus on those who want to establish data centres to be able to indicate some of the benefits that will come to the community as a result of that. They're opening up everywhere, and it doesn't seem unreasonable that, in a risk assessment model, the data centres would have to indicate to some degree what they're going to do and what their costs and benefits are going to be to the community which they serve.

Ms MASCARENHAS: Thank you very much.

CHAIR: We'll bring this session to an end. Thank you again, both of you, for making yourselves available to provide your contribution to our committee's work. If there's been anything in particular that you have been asked to take on notice or anything further that you would like to provide the committee, please do so through the secretariat over the next two weeks.

HAWKINS, Ms Zoe, Co-Founder and Deputy Executive Director, Tech Policy Design Institute

WEAVER, Ms Johanna, Co-Founder and Executive Director, Tech Policy Design Institute

[12:10]

CHAIR: Welcome. I understand that information on parliamentary privilege and the protection of witnesses giving evidence to this inquiry has been provided to you. I invite you to make some opening remarks before we get underway with questions from committee members.

Ms Hawkins: Good afternoon, Chair, Deputy Chair and members of the committee. Johanna and I welcome the opportunity to provide evidence to you today. As a non-partisan, independent think tank, the Tech Policy Design Institute exists to shape technology for the benefit of humanity. One of the ways that we do that is by developing independent research that can serve as an evidence base for important policy processes like this one. I will focus our very brief opening remarks on the dual concept of AI sovereignty and AI agency. I know there's been a lot of discussion about AI sovereignty before this committee and rightly so.

All nations are sovereign, but their power to exert that sovereignty on the world stage surely differs. Consider, for example, the power of a Pacific island nation and the power of the US. Both are sovereign but clearly wield different levels of power and influence on the world stage. At TPDI, we've developed a methodology to define and measure AI agency, which is the power of a country to exert its sovereignty on the world stage and to shape its future, particularly as it relates to AI. You can think of AI agency as the way that Australia can adopt AI whilst making sure that it aligns with our national interests and doing it on our terms. That's the concept behind AI agency. Sovereignty is often framed as a binary: you either have it or you don't; you're either self-reliant or you're dependent. When it comes to AI, that's not a very helpful way to frame the policy problem, the global complex supply chains that sit behind artificial intelligence and the large number of distinct AI capabilities that make up a thriving AI ecosystem.

Our AI agency tool expands the concept of sovereignty from a binary to a sovereignty spectrum. This spectrum recognises, firstly, that international access to world-leading technology does give a country power as a starting point. Secondly, we must have sovereign control over key capabilities. That's a must. Thirdly, having resilience actually comes from reliable and diversified choice between different sources of that capability, both international and sovereign. Fourthly, and in a best-case scenario, a country actually has such sophisticated levels and surplus of a particular capability that it has export leverage, so other countries actually depend on it for that capability, giving it strategic advantage. When you zoom out and apply this frame of access, control, choice and leverage across the full list of 103 AI capabilities that we've defined in this AI agency tool, it allows you more levers to pull. It gives Australia and the committee more options and visibility from a policy design perspective. It gives you the full chessboard so you can see the next strategic move to make based on our unique circumstances as a country.

We have used this methodology, which was developed in consultation with more than 250 experts around Australia, and applied it to see where Australia is today, mapping Australia's maturity, sovereignty and agency against all 103 AI capabilities. This structured assessment can be used as a basis for measuring and monitoring Australia's sovereign AI capabilities going forward. We hope that this framing of AI sovereignty as a spectrum and AI agency as the policy goal is useful to the committee's considerations. We'd be happy to answer questions about the assessment and any other issues of interest to the committee.

CHAIR: Thank you very much. I might just kick off flowing on from that. It's a really interesting concept about AI agency and the spectrum of sovereign capability. We have heard from previous witnesses, and I think this is consistently raised in the national debate, around the opportunity that Australia has with regard to the frontier model hyperscalers that are keen to come here, keen to invest and to bring a proportion of their compute for training their frontier models to Australian shores. Can you provide your assessment of that connection or of taking that opportunity. What does that mean? To what extent does that flow into the work that you are doing at the moment around the assessment on sovereign capability and agency? I like the idea of the 103 AI capabilities. I also suspect that that's not a ceiling that we've reached. I suspect there'll be more as this technology continuously evolves. I'd really be keen to get your views on what's in front of Australia right now, that concept of securing these frontier models for training here and how that relates to our capacity for agency and sovereign capability.

Ms Weaver: Obviously, there's a lot of focus at the moment on frontier capabilities and Australia securing those capabilities onshore. That is one of the 103 capabilities the assessment looks at. We made the determination that it doesn't make sense for Australia to develop our own frontier models. The investment that that would require is enormous. Just to put it in context—it is hard to get reliable figures—roughly speaking, the amount that Anthropic spends on compute for nine months of training is roughly the amount that the Australian federal

government spends on public order and safety. Of course we could do it, but it comes at a significant opportunity cost. The idea of securing access—if we think about the elements of agency, sovereign control being one and international access being another—does give Australia a significant advantage to have these capabilities onshore, particularly from a national security perspective, and access to those models early, but that's not something that will happen unless we negotiate for that access. I think that's something that I have seen in the committee's letter that has been recommended. We definitely endorse that approach; that if they are on our shore, our intelligence community should have early access so that we can be having, for example, patching of security vulnerabilities and other aspects of that.

I think, though, when we look at the frontier capabilities, we also need to look at not just the large language models but also the other national security benefits that come from that. There are, of course, significant national security benefits that others have spoken about extensively to the committee, but a lot of Australia's other defence and national security capabilities will also depend on advanced AI capabilities. The analogy that I like to think about with this, particularly with the frontier labs, is a port analogy. In the 19th and the 20th centuries, securing a port meant that you not only had the foreign ships and the Australian ships passing through that port but also the economics, the whole businesses and industries that pop up around those ports. I think there's every indication that these frontier models will be the equivalent of a port of the future. There's a real national security advantage of having those models, but there are also economic and broader benefits for Australia to be secured from that as well.

CHAIR: I think that captures quite a lot. I hadn't thought about that concept of a port. We've been getting some really good analogies. It's been fantastic. In terms of that, and in terms of the work that you're doing in your space, where do we need to be getting the settings right in order to, as you say, build up that capability and agency, but do so in a way that continues to bring a sense of trust and confidence from the Australian people in this technology and all of the infrastructure that sits behind it?

Ms Weaver: The research really indicates that it's obviously not going to be sufficient for Australia just to have access to frontier models. We need to develop our own industry and our own capabilities. For us, the assessment is prioritising an investment in public interest compute. I know there's been quite a lot of conversation about compute reserve, and I'm happy to speak about, perhaps, some other alternatives or complements to that if that's of interest to the committee.

I also think we really need to be leaning into—sorry, when I say 'I think', I mean the assessment reveals that the areas we should be leaning into are also skills. This is expertise for AI-specific skills but also general skills and skills to ensure that we have inclusive and discerning adoption of AI across the ecosystem.

Also, separate bodies of research, which Kate Chaney, the member for Curtin, was referring to earlier, around trust really indicate that Australians—there's a lot of conversation at the moment around Australians not trusting AI. We looked at 'well, why don't they trust AI, and what would it take for Australians to trust AI?' The resounding response—this was a national, representative survey—was that they want to feel that it is well governed. Seventy-five per cent of Australians have said that they would adopt AI if they felt those guardrails were in place. So, I think, regulation is so often presented as red tape or a barrier to adoption, but, actually, to the contrary, what the research is showing is that, if we can demonstrate that these guardrails are in place, we will encourage and enable adoption.

They're just some of the areas that I think we should be leaning into.

Ms Hawkins: I'll just add to what Johanna mentioned about that enabling nature of regulation as a tool for building Australians' trust and that willingness to adopt. It's just recognising that that is another capability that's listed in the 103—that concept of whether we have the regulatory frameworks—and whether we have the capability to enforce them as well is a separate capability. So it's recognising that, actually, to deliver on that sentiment requires all these different components, and it's an area that the assessment revealed warranted additional attention—looking at what Australia's regulatory settings are and what our capacity is to enforce. Clearly, that's been front of mind in the public debate over the last year and brought to the forefront over recent developments as well. So I think it's breaking down that typical narrative of the idea of 'innovate or regulate', and I think what that national survey and research really found is that that's a sort of false binary and that one can enable the other and is an important foundation ingredient to the whole AI stack.

CHAIR: Yes. I appreciate that. Deputy Chair.

Senator DUNIAM: The 103 capabilities you've got in your assessment tool—how did you come up with that list?

Ms Weaver: It's a very good question. We actually started this body of work because Zoe had done some work with partners at Oxford, looking at mapping global compute. It was around the conversation that we were having before the National AI Plan. I said, kind of blithely, 'Oh, surely we could articulate the elements of national AI capability,' and Zoe kind of blanched.

Senator DUNIAM: No pressure!

Ms Hawkins: I fell off my chair, actually!

Ms Weaver: Yes, exactly! We embarked on a specific research project to do that. We held consultations around the country—13 different roundtables—and consulted, as Zoe said, with over 250 different experts. It covers everything from infrastructure through to data, models and applications and then also looking at skills, governance, and diffusion and innovation. So it really covers the full stack. That was the result of those consultations, which also helped us to develop the methodology, the assessment method and the maturity rankings.

Ms Hawkins: Just to reinforce that, the findings of this assessment are not purely our own. Certainly, both the framework development, that list—it was 101 and then became 103, which I was disappointed about because 101 sounded so much better! But we followed the experts, and they said, 'There are two missing.' It was not only to define that structure, which I think, as I said, is a repeatable method that other countries have indicated interest in as well, because every government and country is struggling with this question of 'what are the ingredients that we're looking for?' The experts also then helped us do that assessment of where Australia stands today against that 103. So it's certainly not just TPDi's assessment; it is shared by the multistakeholder expert group.

Ms Weaver: For show and tell—

Ms Hawkins: It's always good to do!

Ms Weaver: this is an example of one page of that assessment.

CHAIR: Wow.

Ms Weaver: What is particularly relevant, I think, for the committee is that it's not only defining those 103 capabilities. We're also giving examples of those capabilities and indicators of those capabilities. It's really breaking it down into making concepts like a data centre—we're often talking about a data centre. This breaks a data centre down into 27 component parts. That's not because we're trying to make it difficult; it's because each one of those component parts is owned usually by different organisations. If we're really looking at a question of sovereignty or agency, which of that do we need access to, which of it do we need to own, which of it are we good enough to export and where is our choice? You need to have that level of nuance, but we also wanted to make it accessible. This is something that maybe a nerdy and interested person off the street could pick up and, we hope, actually understand, whether it's via the example column, the indicators or the definitions.

Senator DUNIAM: Is that work you've done available online?

Ms Weaver: Yes.

Senator DUNIAM: I'm very interested in this. There have been these nebulous answers provided to most questions asked around sovereignty. I know it's a spectrum. It's not a hard and fast 'this is sovereign, this is not'. It's all of those nuances you talk about, which is very helpful.

Ms Hawkins: That's wonderful to hear because that is literally the policy problem that we spotted. Why we did this work was both because AI was being spoken about like it one thing despite it being all these different capabilities that come together and because that binary of sovereign/not sovereign doesn't leave you with much nuance for real policy solutions. This is intended to do that. Beyond the beautiful PDF which Johanna just held up, we also really wanted to make it usable for policymakers, so we've digitised it into an interactive dashboard.

Senator DUNIAM: Is it on your webpage?

Ms Hawkins: Yes, ai.techpolicy.au.

Senator DUNIAM: We'll find it there.

Ms Hawkins: Look for that interactive dashboard.

Senator DUNIAM: On where the rubber hits the road, in terms of a government trying to catalyse a sovereign industry, are there any outputs or recommendations that might arise out of the work you do or from the model itself that lead to how we might structure a procurement process to, as I say, catalyse an Australian industry?

Ms Hawkins: The assessment shows us where Australia is already leading and the opportunities where we can double down on areas of competitive advantage. Out of the 103, Australia is world leading in eight of those capabilities. That includes critical minerals—which might be of no surprise to the committee—but also on our

datasets. Australia has some of the richest datasets in the world, everything from medical datasets, environment and resources, geospatial, demographic and infrastructure data. Five of our eight world-leading capabilities actually sit in the data layer, which is really interesting, as well as in expertise that Australia has in computer vision—so the ability of machines to interpret and visualise the world around them. Lastly, Australia is a world leader around international engagement and global norm shaping. That gives you a starting point. Then what you can do with that is say, 'If we zoom out, what are the complementary combinations of those leading capabilities that could guide where you'll get the highest ROI for government investment—where we're already strong.'

As an example, think about the fact that if we have really rich medical datasets, and we're really good at computer vision, that positions Australia well to be a world leader in, say, AI medical imaging. That's the value, again, of seeing the full picture rather than looking specifically at those capabilities. We can speak more about that.

Ms Weaver: That's a starting point.

Senator DUNIAM: On the assessment model, as far as you know, no-one else has replicated this or done equivalent work elsewhere in the world?

Ms Weaver: To our knowledge, this is a world-first methodology to assess the nation's agency but also its sovereign capability. We've had a lot of interest internationally in the model. It is building off the work that Zoe and her colleagues did out of Oxford.

Senator DUNIAM: I'm very excited about this. This is absolutely—

Ms Weaver: So are we!

Senator DUNIAM: Yes, I can tell. Thank you very much for that.

CHAIR: Ms Chaney.

Ms CHANEY: I want to start by thanking you for the amazing work you do. I think the AI Agency Toolkit is a useful and nuanced piece of work. I commend you for that. No doubt you read our letter to the PM this week. I'm interested in your views on whether there are any safeguards that are missing from that letter that the committee should be covering.

Ms Weaver: I think we definitely endorse all of the safeguards that were included in the letter, particularly the additional funding for the safety institute—I think it's important for us to have on the record that we support that—the independent evaluation, the safe harbour for incident reporting and also the legacy systems review. I think some of the things that were missing for us are some of the issues that have been raised today around privacy reform. I actually think this is the single most important piece of legislation that can pass and have a huge impact on AI regulation. The digital duty of care, which of course was introduced to parliament this week, is also very, very important. And I also think the automated decision-making in government legislation, which we are waiting for—I know, Kate Chaney. you've introduced a private member's bill on that. I think it's really important that we're also looking at government use of artificial intelligence; that is really important.

One other area, which wasn't included, that I think is a very important safeguard for Australian democracy is looking at the use of AI in elections. I appreciate that this is a challenging space. We're effectively asking you to regulate yourselves. But I do think it is important that we have a commitment from political parties and from all candidates to have transparency in the use of artificial intelligence and. in particular, the labelling of any products that you're using that are intentionally deceptive. Effectively here I'm talking about deepfakes. Some people call for the banning of deepfakes during elections. I think that would be an ideal position. If we don't move to banning, then there should be, at least, a commitment, if you are doing a deepfake as a political party, to label that. I think that is a bare minimum that we should be having in place.

I also think we should be using the forthcoming election campaign—which I don't want to think about yet because we've still got plenty of time to do really good work in this term of parliament—as an opportunity to have a widespread campaign teaching Australians to spot authentic content. We've moved past a point of being able to teach people to identify fakes. It's too hard; the AI is too good. Let's teach them where to go and how to identify reliable, verifiable information, whether that be through a marking from a political party or verification through trusted public interest media. There are plenty of ways we can be teaching people to spot the authentic content, rather than spotting the fakes. So I think they are some of the safeguards that we're missing.

The final thing I'll say on that—well, actually, I have maybe two points. One is if we are relying on existing laws—whether that's privacy or whether that's workplace health and safety—to regulate artificial intelligence, we need to make sure that we're funding and resourcing the regulators to be able to do that. I think we've got a lot of

focus on the fact that we need to use those existing laws. We haven't had much focus on funding the regulators to be able to do that.

I also think we should be specifically looking at the question of liability, as well, around AI agents in particular. Some of the aspects of liability can be managed through existing provisions—for example, saying to a company director: 'Set the expectation that, if your company is using an AI agent, you have an obligation to make sure that agent is doing the right thing and obeying Australian law.' This is quite similar, I think, to the provisions and the expectations that were set around cybersecurity of companies. Directors have an obligation to ensure cybersecurity is in place. I think we need to make sure that we have that expectation when it comes to AI agents.

I think there's also a question about liability for companies that are developing these AI agents, as we've seen recently through the OpenAI and Medicare incidents. And I think there is, finally, a unique opportunity for Australia, potentially, around AI agent identity, as an area where Australia can move forward.

Ms CHANEY: Can you talk a little bit more about agentic identity.

Ms Weaver: I think what really focused my mind on this was when we became aware that the Medicare site had been compromised. That was because OpenAI had a notification. If Services Australia had identified that themselves, rather than it having been brought to their attention, it would have been very difficult for them to have understood that it was an AI agent versus a standard compromise, and it would have been very challenging, if not impossible, for them to have been able to identify whose agent had compromised the system.

There are a number of very nascent proposals being discussed globally around whether or not it is possible to establish identities for AI agents, and it is interesting, when you think back to the early days of the internet, that there are analogies here to when we first started having the domain system, for example. You would have the spread of different .com domains, for example. When it first started, it was researchers to researchers. They knew each other, so they knew who those domains were being issued to.

Then, when it became more commercialised—similar to what we're seeing right now with AI agents—some domain administrators made the decision that they would attach identity requirements. For example, in Australia, if you want to register the .au domain, you actually have to provide some evidence that you have a connection to Australia. Usually, that's by proving that you have a business name register, for example. That means that the .au domain in Australia is one of the most trusted domains globally. We're not a domain where huge amounts of scams or spoofing comes from because we have that identity requirement.

It's not privacy invading, because it's not saying that this is Jo Briskey's website; it's just attaching it to a particular company, for example, similar to the way that a number plate registration registers a car, but you're not necessarily sure who's in the car or who's driving the car at any particular time. It provides a mechanism for accountability, and I think that, as AI agents become more pervasive through our economy, we will need to be able to verify who the agents are acting on behalf of. We will want to do that anyway in the transactions that we're having. But I think there's also an important accountability element to that, and Australia has been a leader on this front in many spaces before. So I would recommend that, if Australia does not establish an AI agent mechanism, the committee recommend that Australia actively look at establishing this.

Senator DUNIAM: I reckon you could be guaranteed of that.

Ms CHANEY: Thank you. That's a whole area we haven't really talked about much. I just want to go back to the deepfake issue in elections. You say: ideally, ban AI use in political ads but maybe a compromise position is labelling. I've met with psychologists who say, basically, we've been trusting our eyes for 10,000 years, or longer than that, and even though you might be told something is AI, it still actually changes how you feel about it. Why is there a positive use case for AI in political advertising that means we can't aim for a ban on the use of AI in political ads?

Ms Weaver: I think there are lots of ways that AI is used in political advertising every day. Just a general Canva image involves artificial intelligence. So I don't think you want an all-out ban. But that's not really what you're asking. You're talking about its nefarious uses. I think that's where the labelling does become important. There is a lot of research to support the point that, once people see something, they think that it is real. But I think that is why the education campaign is really important. We are trying to have people understand now that the content that you see online is not reliable unless it comes from a verifiable source. That requires teaching Australians that they need to look for markers of trust and authenticity rather than look to try and spot fakes. It is a total switch in the way that we have been engaging with the public to date, but I do think an election campaign where you have a trusted entity like the Australian Electoral Commission, which has played that neutral role for a long time, provides an opportunity to educate the Australian population of that issue during an election campaign,

but then, also, for that message and those lessons to have much longer—hopefully—positive benefits across the economy.

Ms Hawkins: To add to that, I think that it's also an opportunity to demonstrate leadership—that government is expecting to see this level of transparency and labelling from industry. I think it's a really great way to match that expectation and say, 'We expect that from industry, and here's how parliamentarians, politicians and government actually reciprocate that expectation of transparency for the public.' I completely agree about that concern of being convinced even if it's labelled. I think it's maybe an example of us marrying the policy solution with the pragmatism of what is going to be politically viable as well. You don't have to look far across the online safety and misinformation debates to know that there are going to be large pockets of what society expects. Artificial intelligence, as Johanna mentioned, comes in many gradients of converting an image—sometimes it's a background remover; it even goes all the way through to political satire. If someone made the argument that this is the new form of political cartoon, some people might object to an all-out ban. So I think, if you're looking for a pragmatic solution that would move the dial but maybe not get caught up in some of those more complex debates around political expression, that might be another reason to consider that as a middle ground.

Ms Weaver: Yes, pragmatic optimism.

Ms Hawkins: Exactly.

Ms CHANEY: Thank you very much. I just have a last quick question. It's easy for Australia to feel like all these decisions about the big picture are being made overseas. What do you see as being Australia's role in shaping global norms and standards? Are we kidding ourselves in thinking that we can have some influence, and is there any evidence to the contrary?

Ms Weaver: As Zoe said, one of the eight areas where Australia has the highest competitive advantage globally is our ability to engage and shape global norms on these issues. I know that there is a lot of scepticism among the general Australian public about that, but Australians have been engaging on these issues for a really long time. So—whether it's Aurelie Jacquet or Lyria Bennett Moses, who've been leading discussions on standard AI standards at the OECD or at the international standards organisations, or whether it's people like Briony Daley Whitworth, who led negotiations on the cybercrime treaty, or whether we go as far back as 2013, when we had Deborah Stokes, who was the chair of the group of governmental experts at the UN, where Australia led the negotiations that established that international law applies to state conduct in cyberspace—Australia has been leading conversations around the rules and norms in the online environment, traditionally, in the context of cyber—and you can add an adjective to that. That means that countries are looking to Australia to provide leadership in this space.

So, yes, the US and China are leading in terms of technology development in this space, but, if you travel internationally and ask other governments who they look to to be the sensible middle-ground voices, especially on tech issues, it will be Australia. I think Australia was the first country to have a cyber ambassador, which evolved into an Ambassador for Cyber Affairs and Critical Technology. That position has largely morphed to become focused on cybersecurity issues. I want to disclose my interest here. I did work for the first ambassador for cyber affairs, and back in that time we deliberately conceived it as broad—as in 'cyber affairs' in the same way we talk about foreign affairs. I think the subsequent cyber ambassadors have very deliberately, at decisions of government, done an excellent job but have narrowed the scope back to cybersecurity. I think there is a huge opportunity for Australia to be playing a leading role, and countries are looking to Australia to play a leading role in this space, but we need to either resource the Ambassador for Cyber Affairs and Critical Technologies more to be focused on these AI issues or perhaps establish a position of an ambassador for AI affairs. Dare I suggest another position?

But the point is that countries are looking for Australian leadership, and there are areas where we do need international cooperation on these issues. Talking about incident reporting, that will be much more effective if we do that in partnership with a coalition of countries. Let's learn from the mistakes that we made during cyber incident reporting. I was in government negotiating some of those agreements. If we can get it in place across multiple countries at the same time to be having the same AI incident reporting, it makes it so much more efficient. Let's get that going from the start. AI liability, for example, will be much more effective if we have a coalition of like-minded countries agreeing on how and when we will hold companies to account. So there's real demand for it, and countries are looking to Australia. I think that's an area which I would recommend that the committee look into more. I just want to emphasise I'm not being critical of the two very excellent cyber ambassadors, Jess Wilson and Brendan Dowling. They have focused very deliberately in that area. But I think there's scope for us to be doing more.

Ms CHANEY: I appreciate that.

Ms Hawkins: I will just add to that. I also think that it's not only the interoperability and the efficiency of that moment that Johanna's describing. I think it gives middle powers the strength in numbers to not feel deterred by the either implicit or explicit threats of retaliation from the US government for imposing regulations that might be perceived to disadvantage American companies. So there's a strategic reason to do that as well as a great opportunity for leadership.

Ms CHANEY: Do you see the AI Safety Institute playing a significant role in that work? If so, what resources do they need to do that?

Ms Weaver: It could be something that becomes a pillar of the AI Safety Institute's work, and, if so, they'd need a team to be equipped and resourced to do that. That fits naturally, in the sense that the AI Safety Institute already operates within a network of safety institutes, variously named, internationally. I would encourage us, though, to think about it as something that needs to sit within the Department of Foreign Affairs and Trade, because the AI Safety Institute is focused on technical issues. This is a strategic issue. This is about positioning Australia so that we are leveraging our full suite of strategic capabilities, and that does require diplomatic statecraft to be able to do that—which, of course, could be housed within the safety institute. But I think it's a complementary focus and skill set, which would work very closely with the safety institute.

In general, I would look at it through the various different cyber negotiations I was involved in, in my diplomatic career, and I would say that, when you have the technical people at the table, it's a very, very different conversation. It is a maturing of the recognition of the strategic impact when you see it sit within a foreign affairs portfolio. That's why it was so significant that Australia was the first country to appoint a cyber ambassador sitting within DFAT.

Ms CHANEY: This is a very quick last one. We do need that international cooperation around things like incident reporting. International cooperation can take a long time. How do you think we see that balance? Do we get on with having domestic requirements now, and then feed them into international discussions, or do we wait until there are international standards before we do anything in Australia?

Ms Weaver: We have to do both simultaneously. But I would push back on international negotiations taking a long time. After the Christchurch terrorist incident in New Zealand, within one month 15 countries had come together with all of the US social media companies and agreed the provisions of the Christchurch call—that we will not accept abhorrent and violent content online—and agreed the measures to take it down. If countries are motivated—and I believe they are motivated, and I think Australia should be using the Medicare compromise as the catalyst for Australia to be leading a conversation to say: 'This is not acceptable.' We need countries to come together, and this doesn't have to be something that takes a long time. Some negotiations—particularly ones that sit within the UN, for example—do take a long time, but there are things you can do in minilateral environments that can move very, very quickly.

Ms CHANEY: Thank you.

CHAIR: Senator Sheldon.

Senator SHELDON: Thanks very much. You spoke before about the deepfakes in election times and some of the potential solutions in a policy sense, though I appreciate you spoke quite broadly. Have you had an opportunity to look at the South Korean example—in, I think, two election cycles or thereabouts—as to how that played out, and whether that is an example that we should give a lot of weight to or give some consideration to?

Ms Weaver: You are stretching my memory a little bit on this one! Correct me if I'm wrong, but I understand the South Korean is a direct ban on deepfakes. It's really going to Zoe's points before, around the pragmatism. I think that is a good measure, but it hasn't and doesn't, in Korea, prevent other parties using deepfakes. For me, it comes back to this: if we're not banning it, then label it. Label it as satire. Label it as: 'We've used AI to do this,' but also teach people: 'If you want to make a decision about who to vote for in the election, go to the relevant political sites to get that information. Don't rely on information that's being fed to you elsewhere.' So it's that flip, in saying: 'We're teaching people to look for the authentic content, to look for the verifiable content, rather than teaching them to spot fakes.' That would be the key lesson that has come out of the Korean example, for me. I also think this is not something that Australia needs to make up. The 2024 Munich Security Conference brought together a number of countries and a number of companies. There's a declaration. There are protocols in place. We're not reinventing the wheel here. This is something that tech companies have come together for that has been rolled out in other jurisdictions.

Ms Hawkins: Within the 103, we actually considered, which is what Johanna was just mentioning, making sure that citizens are equipped with that ability to live with AI. That was one of the things that came through really strongly in our national consultation. We talk a lot about national capability. I think we often have very

institutional visions of what that means. But really it's actually thinking about national capability at that individual level. Societal and democratic resilience will often come down to how equipped individual Australians are with that critical thinking and with that AI literacy. That's thinking through how this needs to start in schools and public awareness campaigns around elections. There are so many different ways that we need to start building that national capability, but it's at an individual level as well.

Senator SHELDON: I'm not sure if you heard our previous two witnesses. They spoke about this. It goes to this question of general skills and probably that point you made, Ms Hawkins, on how we build up capability across the community. They went to two examples of particularly the Occupational Health and Safety Act and the Fair Work Commission. My impression from commentary on this is that, in particular, the Occupational Health and Safety Act has quite a settled sense across the Australian community, business community, workplace and workforce that you abide by those laws. There's an ambition and almost inherent knowledge that, if wires are hanging out of the wall, something's wrong. There's something a lot more elaborate than that. People are aware to have a conversation if something looks like it may be or has a possibility of being dangerous or has a possibility of being safer. 'So how do we not only risk-manage with the risk management I have done but consult with the rest of the workforce or the rest of the company or the rest of the industry on those things that have embarked with government guidelines built by industry et cetera?'—is that the frame? Do you think one of the frames would be an area where we've got a really cutting-edge, already built in hot-wired system of community consultation, engagement and actual voice? Is that an area that has some dual ability for the sorts of things you're talking about?

Ms Hawkins: Yes, absolutely. I think that it's really picking up on that duty-of-care concept as well. That's why the relevance of the duty-of-care reform is really putting that expectation on service providers to say, 'In the space you create, you're responsible for the risks that may arise and for anticipating those and preventing them as well.'

Ms Weaver: I think also both the Fair Work and WHS provisions are relatively technology-neutral. I think we want to avoid a situation where we're continually updating those pieces of legislation whilst at the same time making it clear that AI is just a part of most workforces now and those WHS obligations extend to the use of AI and the implications of AI—potential harms but also the benefits—and then making sure that we're equipping workforces and regulators.

I thought the most important thing that I'd underscore from the comment that was made in the previous session was the example, I think, from Finland where they've provided guidance and information about what that means. I think most people accept that WHS would extend into artificial intelligence and the use of artificial intelligence. What Australians are grappling with is: How do I comply with that? What do I do? So I don't think it's necessarily a situation of new laws. It's about making sure we're equipping the regulators to provide that guidance.

Senator SHELDON: Similarly, the Fair Work Commission just recently assisted with the guidelines, but also the commission assisted with the now infamous Commonwealth Bank redundancies replaced by AI. But, to the credit of both parties, after that argument and fight and with the assistance of the Fair Work Commission, they came up with a settlement to retrain and reinvest and also deal with the productivity questions the company was trying to deal with. I understand that they've just come to an agreement on some wider initiatives regarding how to engage with AI with the workforce, going to the education and skills development piece. So, again, that question is about how we adapt existing institutions rather than have something separate.

Ms Weaver: I think it's also bringing an element of humility. I hope this isn't insulting. I know that might be hard in politics, but the reality is we're going to make mistakes. We haven't done this before. So it is recognising that maybe there are circumstances where things are done wrong and that cause damage. But how do we then make sure that we are recovering from that and applauding when companies and regulators recognise those mistakes and move forward rather than pointing the finger? I think that's actually going to be a really big part of creating a resilient Australian AI ecosystem.

Ms Hawkins: Just to add quickly to that, it's one of the reasons that one of our national capabilities is discerning adoption. Importantly, our framework is not starting from the assumption that more AI everywhere, all the time, for everyone is always the right answer but actually that sophistication and maturity look like giving organisations, workers and different people meaningful critical skills, choice and strategic consideration about when AI is actually going to add value. I think that that will help with that trajectory that Johanna was mentioning.

CHAIR: Excellent. That brings us to time. Well finish up there. Thank you so very much for making yourselves available. It's been a really useful contribution to our work as the committee. If there's anything further that you've been asked to provide on notice or there is anything further that you think would be helpful that you would like to contribute, please do so through the secretariat over the next two weeks.

Proceedings suspended from 12:56 to 13:40

KEAY, Dr Sue, Deputy Director, Kingston Artificial Intelligence Group [by video link]

TEASDALE, Dr Hayley, Head, Science Advice and Policy, Australian Academy of Science

WILLIAMS, Ms Sally-Ann, Fellow, Australian Academy of Science

CHAIR: Welcome. I understand that information has been provided to you on parliamentary privilege and the protection of witnesses giving evidence to this inquiry. Thank you very much for making yourselves available. We can kick off with brief opening statements, and then we'll get underway with questions. Over to you.

Ms Williams: The Academy of Science is an independent, not-for-profit, non-partisan organisation. Thank you so much for the opportunity to appear today and contribute to the proceedings of this committee. Australia is at a critical moment in time. We will be either a creator of AI or merely a consumer who hosts infrastructure for others. We must make strategic investments to build sovereign AI capability, as we are already falling behind. We are also a country in high demand for data centres, hyperscalers and AI providers to operate and build their infrastructure. These companies want our land, our energy and our stability. In return, they must contribute to our sovereign capability and not just extract value.

The academy's position is that Australia should channel that demand into something that will bolster the Australian economy. That's why we've proposed a levy of 2.5 per cent on AI providers' and data centre operators' revenue that would be used to build public AI infrastructure for use by government, researchers and Australian businesses. This mechanism has been referred to as a public interest research payment by the Office of AI in their recent consultation on AI standards. This kind of investment would propel Australian productivity in a way that is urgently needed. Not investing now will not only stall productivity but lead to the Australian government paying significantly more to access this infrastructure over the decades to come. Without that investment, Australia risks becoming a host for global data centre infrastructure while lacking any true sovereign AI capability of our own.

We were very pleased to see the interim advice to the Prime Minister include many critical areas which we support, particularly recommendations on AI safety and the need to maintain funding for public compute. The important gap that still remains is a mechanism to fund the scale of infrastructure we require. In times of global instability or simply extreme commercial demand for compute, Australia must have its own capability. We cannot leave our scientific advancement and national security at the mercy of surge pricing or foreign trade controls. Relying entirely on offshore AI capability is an escalating national risk. But, if we are bold enough to demand reciprocal investment, we have a profound opportunity to secure Australia's economic resilience and a scientific future for generations to come.

Dr Keay: The Kingston AI Group is a collective of Australia's most senior AI researchers from multiple universities. We'd like to challenge the idea that Australia can only ever be a consumer of artificial intelligence. Our researchers publish in AI's most competitive venues at around six times the global per capita average, ahead of the United States, Europe and China, and our top AI researchers are more likely to hold a PhD than those of any country except the United Kingdom. We have the talent; we're just not backing it up, and that forces our talent offshore. This committee has spent the past fortnight hearing why that matters.

When an AI agent accessed non-public data on a Medicare statistics portal, the government found out 84 days later, by email, presumably written by AI, but at least the company told us what had happened—eventually. In June, a United States export directive cut Australian users off from the most powerful frontier model. Both of these examples show how little Australia can do when another country controls the AI that our businesses and services are increasingly relying on, and that is true regardless of whether those services are located on Australian soil. Sovereign AI capability means being able to build, adapt, evaluate, operate and sustain AI systems on compute and data infrastructure that Australians either own or control, including when access to foreign services is disrupted. Today, we are clinging on to the people part of that capability and little else.

We recommend a targeted response, coordinated through a national sovereign AI coordination vehicle. First, data centre and AI infrastructure developers seeking government facilitation should be required to reserve a defined share of installed artificial intelligence compute for a national pool on the same principle as gas reservation. That should be backed by a dedicated computing fund and by Commonwealth procurement that builds local capability—the same way that NASA's purchasing underwrote the semiconductor industry. Second, major AI companies doing substantial business in Australia should maintain a genuine research and development presence onshore, with the IP staying here. That obligation should attach to the companies that train and control the models, because that is where the value lies. And third, the Australian AI Safety Institute needs more funding, and a formal way to draw on Australian researchers in evaluation, explainability and machine unlearning so we can independently evaluate foreign systems.

Australia doesn't need to beat the United States or China at frontier scale. We do need to own the capabilities that matter most to our security and our economy so that Australia is, as the Prime Minister puts it, 'much more than a data warehouse for AI products made overseas'. Australia has no shortage of AI talent or ideas. What we lack is the confidence, investment and coordination to back our own ability. Thank you.

CHAIR: Thank you very much for those opening remarks. Something that has come up, as you quite rightly said, in the last week of the hearings and also what we're considering and included in our early iterations of recommendations to government is around sovereign capability and compute. I want to get further clarity on what that could and should mean in terms of reserving compute. Do we reserve compute capacity from the hyperscalers if they are to come here? What is the difference between that and actual investment in public compute infrastructure? Could you give the committee a bit more clarity on what we should be looking at to secure, as you rightly say, the ability to have sovereign capability in our compute capacity for public benefit?

Ms Williams: I think it's a case that we need to be really clear that it's not an either/or; it's a yes/and. That's the first point. We have to be crystal clear on that. I think the second point is understanding the types of things that are done from both a research perspective and the users that we have. When we consider public compute, and if I think about Pawsey and NCI—disclosure: I am the chair of NCI, so I'm intimately familiar with both supercomputers right now—and if we look at the users of that, it's not just researchers. There are actually around 21 government agencies and others doing their modelling on there as well, and those things don't necessarily make sense to do in a commercial environment. That's for two reasons. One is around cost and the longevity of that. The other really critical one is around the datasets and where they're stored.

So, when we're talking about large datasets for climate modelling, for bushfire predictions and things like that, we actually need those to be housed both in a sovereign compute infrastructure location but also close to where the modelling is going to be done. Moving data from one place to another makes absolutely no sense commercially, and there's a huge cost to it. So there's an element of, 'What do we need for public compute?' That's both in delivery of government services and government work and in research. Then I think it's, 'How do we work together?'—so the 'yes, and'. If I think about reserving compute, we need to understand a little bit more about what that looks like. There needs to be a coordination of effort. Not every AI model can be deployed and not every AI infrastructure stack is built the same, and they're not always appropriate and easy to shift from one to another. So we need to have a coordination effort to help with anything in that to really seamlessly harness the benefit both from a research perspective and from an industry capability perspective as well.

CHAIR: Dr Keay.

Dr Keay: I will just add to what Sally-Ann has just said. Like Sally-Ann, I wear a few different hats. I'm also the director of the UNSW AI Institute, and we know from our institute's perspective that researchers in Australia are not getting access to the compute that they need. There hasn't been the public investment to help to shore that up. We're fortunate that UNSW is one of the largest universities in the country. We're working together with the University of Sydney to invest in additional computing capacity for our researchers. However, we're only taking on the burden of that cost because we are not able to access the public infrastructure that we require. We're looking to access GPU as a service with a local Australian neocloud provider, but that will cost us a lot more than it would cost if we were able to access public infrastructure. However, it is cheaper than using services from the hyperscalers, so it's a reasonable middle ground.

But the idea of reserve compute capacity is only necessary if the government decides that there is going to be no public investment in artificial intelligence, so, as a backstop, a few of us in the research infrastructure community have tried to look at what that would mean and how it would work. As Sally-Ann said, it requires very good coordination, because having access to stranded compute is not at all helpful. The way the reservation could work could take a number of forms. It wouldn't necessarily be that each company would have to give us GPUs from their facilities. You could create a marketplace where they invest the equivalent amount that it has been determined that they should allocate and it is given to that national coordination body to determine what that looks like, where it is housed and how it can most effectively be used by research and industry.

CHAIR: Just to clarify then, is the concept that you're looking at the importance of having—I don't know if 'flexibility' is the right word—the ability to be a bit more directive about the compute, what it is and what it's used for? As opposed to just a straight-out condition where part of hyperscalers coming here is that they must reserve or provide access to a certain amount of compute for research or for public use, what you're talking about is the extent to which, yes, we do that but we do it in a way that gives a sense of flexibility to how, where and when it's used. Is that what you're saying?

Dr Keay: Yes, exactly. Those decisions should be up to us, as opposed to up to the supplier of the compute.

CHAIR: I appreciate the point about the importance of public investment in compute, but, in your view, is there an opportunity here, with the leveraging of these hyperscalers wanting to come here and invest and significantly increase the investment in compute onshore here, to provide a sense of scalability for us in that space if we get it right, rather than simply determining how much compute there is or what it looks like?

Dr Keay: I think it's definitely an opportunity. Sorry, I'll let Sally-Ann talk.

Ms Williams: I think it is an opportunity. I think the devil is in the detail, always. What the academy has proposed is a 2.5 per cent contribution, and I think there's a really great example that we have had operating in Australia for decades—which is the RDC model—of how we align the investment side of industry—so the hyperscalers and others coming to Australia—with our research and government needs and our public needs to come together and provide clarity. If you think about the GRDC, the Grains Research and Development Corporation, what that enables them to do is to bring together public and private investment or public and private contributions around the long-term needs of that sector. So they invest in long-term research needs that have 10- to 15- or -20-year horizons. Then they look at innovation in the short term as well. I do think that, when we think about contributions, regardless of whether they are in the form of a levy or reserved compute, that coordination and that alignment of strategic intent and bringing both sides together is actually one of the critical features that we need to get right. If we leave it to the decision-makers of hyperscalers, they have an obligation under the public markets to benefit their shareholders, and that's not necessarily always going to align in the long term with what Australia's needs are.

So I think we have a way to come together. We've got a proven model that works. We've got a way to actually come together and say we can make the win-win on both sides of the equation here if we actually bring that coordination together and facilitate that in a way that has the long-term expectation of benefit for Australians, for Australian industry and for productivity uplift across all sectors, not just AI.

Dr Keay: The Kingston AI Group put out a joint statement with the Australian Council of Trade Unions that is very much along the lines of what the academy has said: it needs to be an investment that not only has reserved compute capacity as part of it but also has funding to enable us to build AI models and to build up Australian capability in the form of the people that are required to actually make use of that computing capacity. The RDC's is a good model for how to do that. I would actually go a lot higher than 2½ per cent, but that's something I'm sure could be worked out.

CHAIR: Finally before I move to the deputy chair, you pointed to a framework that we could work on, but is there anywhere else overseas or in other jurisdictions that is or has done this well?

Ms Williams: Thank you for that question. I think it's really relevant. There's an increasing number of countries that are actually looking at how do they get the settings right around both public and private investment and align them for sovereign capability and needs. So we're seeing that. We can provide more information afterwards, and we'll take that on notice, but we can look at the European Union, Germany, South Korea and Singapore—all of these settings—and see if they're actually accounting for their sovereign needs in the future as well as enabling businesses to get a productivity uplift not just by being consumers and users of that and adopting AI in their business but also by looking at the long-term settings. I think that's a really good model to follow, because the other point to make around this that I think is really critical is that we've got to get our risk appetite around it right. If we are solely placing our requirements, our needs and our services in one company, one country or anything like that, we create an unnecessary risk in the long term that we need to think about mitigating. We need to make sure that we're future-proofing for our needs.

CHAIR: Deputy Chair.

Senator DUNIAM: Thank you very much for your time today. I wanted to clarify something. Dr Keay, you may have said that the reserve compute is only necessary if you don't have publicly funded compute or purchase it. Can you explain that to me? I got the sense from other parts of what's been said that it would be good to have both. I just want to understand that because, obviously, at the end of our committee's work, we're going to make a recommendation to government about what that public good element of this will be and how best to structure it. There's obviously sectoral leadership here. I'd be keen to understand that.

Dr Keay: Australia hasn't invested in upgrading its supercomputing facilities since 2018. There have been calls for the amount of public investment in that space to not only be guaranteed but be significantly increased so that it is in line with peer nations. In terms of whether there are analogies with other nations, there aren't any other nations that are copying Australia in having almost zero public investment and expecting the private sector to fully take up the burden of providing AI infrastructure. I think that's the main point to make. I don't think there's any suggestion that it should be all public or all private, but, at the moment, we're actually looking all private.

Senator DUNIAM: On that point, what sort of quantum of investment are we looking at here to catch up with peer nations? I take your point that we are aeons behind.

Dr Keay: Last year the Australian Academy of Technological Sciences and Engineering put forward a report that suggested we should be investing at least $1 billion per year over the next five years to catch up. At this point, that would be an underestimate compared to what other countries are putting in. You might be aware that Korea has invested or is looking to invest the equivalent of $212 billion and they have already acquired the rights to get 260,000 GPUs. In Australia, our researchers can access at most 2,200 GPUs, and we can't pool them all together. To actually do trains on large clusters of compute, our researchers have to go offshore, as do our businesses.

Senator DUNIAM: Dr Teasdale?

Dr Teasdale: The Australian Academy of Science is working very closely with the Academy of Technological Sciences and Engineering. The number that Dr Keay used was more looking at AI factories and AI infrastructure specifically. The number that we have used specifically for NCI and Pawsey, as the two largest public supercomputers in Australia, is $200 million per year over 10 years, or $2 billion, just for those facilities to keep pace and so that we could start with the next generation of computing that many other countries are already enjoying the luxury of: exascale computing. That would help us basically get a pathway to exascale, but that would not pay for exascale. It's been about two years since we created that number. That number has probably gone out of date. If we were to try to restore those computers right now, it would probably be more, but that's our most-up-to-date estimate.

Ms Williams: I can't go into details, but I can speak specifically to that around procurement and refresh. I believe you heard from NCI and Pawsey last week. There are refresh cycles going on at the moment. When we look at where the market is in terms of cost, we are not going to get a like-for-like in the current refresh cycle. Part of that is that other countries do model refreshing. They don't have to wait for a grant procurement pathway to go forward; they actually have, every two years, a guaranteed refresh cycle, which you would do in a commercial setting too to make sure that you are continually maintaining, uplifting and giving certainty and confidence that you can guarantee the capabilities from a facilities perspective and give certainty and predictability to your customers. In the case of NCI and Pawsey, that's not just the research sector. It is the Bureau of Meteorology. It's Geosciences Australia. It's our research, our medical infrastructure and our health research units. We give them certainty and confidence that they can then plan alongside of us for both research and the delivery of services that they need to do as well.

Senator DUNIAM: And it stands to reason that refresh with this sort of technology is going to be necessary over time.

Ms Williams: A hundred per cent.

Senator DUNIAM: So you're speaking to a change to procurement. In my mind, procurement in government investment is necessary to catalyse a sovereign industry here. You're talking about a change to government procurement and how that relates not just to R&D in the science world but also across government to give effect to what you're talking about.

Ms Williams: Yes, I am. I think it's really important to note that this isn't just about AI. When we think about it and we think the large number of departments that actually use the supercomputers in Australia, that is around our weather modelling. It is around our bushfire predictions. It is also going into agronomy data and actually supporting our farmers in how they are thinking about making decisions and changes around producing food. It touches every single area of the Australian economy, and we need to think about those sorts of things in the long term to give certainty and stability to the system as well.

Senator DUNIAM: Dr Keay?

Dr Keay: This is why digital infrastructure must be recognised as critical infrastructure. At the moment, it is only recognised as critical infrastructure under the SOCI Act, so for security purposes. If it were recognised as critical infrastructure, it could be looked at in terms of long-term budget and planning.

Senator DUNIAM: I am interested in that as well because I think the rigidity of procurement rules does prevent government from sometimes making perhaps the best or most appropriate decision. We've just talked about case studies where you've got standard-form procurement, which doesn't allow us to get the best out of it. Yes, I think the recognition of critical infrastructure is a good one. Because AI and all that goes with it does broadly—if you look at what happened a few weeks ago with the publication of the unauthorised access of government databases by an agent, there is a national security component here which might be a vehicle for altering the procurement rules like we're talking about. Is that something that you've contemplated?

Ms Williams: I think it's a really critical thing to look into, and I think it goes to not just the hardware and the compute side of things but also data and storage. The other piece to this is, when you look at the modelling and the work that is being done on our supercomputers from a research and government services perspective, is that those datasets are also really critical. We need to think about what the value of those datasets is and not just how we protect them but also how we see them as an asset and ensure that we are giving longevity to the ability to do things upon them.

Senator DUNIAM: Thank you very much. That's very helpful.

CHAIR: Ms Chaney.

Ms CHANEY: Building on that, and this is maybe displaying my ignorance, what's the impact of recognising digital infrastructure? Firstly, do you mean public interest digital infrastructure or broader as critical infrastructure, not just for the SOCI purpose—for protecting it—but also for long-term budget planning? What actually changes if we recognise that?

Dr Keay: You've got dedicated bodies like Infrastructure Australia and state based bodies like Infrastructure NSW that report on these things, so there'd be someone who, instead of reporting on how many tonnes of concrete are poured every year and how many tradies get trained, is actually more broadly looking at the health of our digital infrastructure. I think it directly feeds into budget considerations as well. It kind of forces more long-term thinking about this rather than it being a surprise every year that we need to fund this material.

Ms CHANEY: I'm trying to understand the substitution potential between the public-interest supercomputers that we have and the compute that could be reserved on an LLM training model. I know that both national facilities are oversubscribed. I think we heard three-to-one demand versus supply last year. How much of the demand for national compute facilities could be feasibly met by reservation in those training data centres? Is the main challenge data transfer in and out of those, or are there other particular types of AI use for public interests that don't work for LLM training models?

Dr Keay: I think the point is that you wouldn't reserve the capacity in those facilities. You would put the levy in place, and then they would meet that levy by either contributing into or supplying infrastructure that then gets coordinated. Whether it's housed at NCI or Pawsey is a completely separate issue; it depends on where it makes sense. Sally-Ann might want to weigh in on that. But I think it's not about having GPUs where those frontier models are being trained. It's about asking for a percentage of reserved capacity and then having a marketplace where those companies can meet that obligation, but it's not necessarily by trying to get the GPUs that are in that facility because that's almost impossible to be able to do.

Ms CHANEY: So it's more they provide an in-kind contribution of chips or infrastructure or something. Or does it come down to a cash payment?

Dr Keay: It could be either—if they want to donate GPUs, if they want to provide cash.

Ms Williams: If I may, I'd like to take this one on notice to come back with some very clear examples for you of where this wouldn't work in practice, because it isn't like for like. You mentioned LLMs; that is one AI model. When we're talking about AI models in health, weather, astronomy, critical minerals or other sorts of things, we're actually talking about different ways of configuring and servicing those needs. It's not an apples-and-apples comparison. So I think it would be really helpful for us to come back to you with some examples of where that would and wouldn't work.

I do think we've been having an extensive conversation across Australia, and where I think the coordination needs to come is: what are our high-performance computing needs? In that includes the part that is public investment and public infrastructure and also private infrastructure. What we'll see is what industry needs and what business needs for an uplift. They're looking at something quite different from what our agencies are looking at and from what our researchers are looking at. It is about understanding the nuance of that, and going, 'We want to be efficient about this, and we want to get the wins on both sides of the equation.' So those who are coming here to set up and set up data centres, and, actually, the hyperscalers as well, do need to win on a business front, but how do we actually align those two things so that it isn't extractive on one side and we're balancing the equation in the right way. I think we can come back to you with some very specific examples of which ways would work to help illustrate that, because it is complex technically to get into the weeds about that.

Ms CHANEY: I think that would be really useful. In the public mind, you're saying, 'Well, we're going to build these big data centres, and Anthropic will do training in them, but they'll cut a bit of it out that we can use for our purposes.' So explaining that it actually looks as though, in order to come here, you have to make a contribution to our public interest compute, which is going to sit somewhere else and has different requirements to what you're doing in that shed.

Ms Williams: Yes. There's a real challenge if you are a hyperscaler or somebody operating. They don't want to be doing deals one to one with everybody around how they actually access this sort of thing. There is a coordination piece and an ease of use on the other side of the equation that I think can actually help industry. They all want to make a meaningful—well, they say they all want to make a meaningful contribution to Australia in the long term. I think we can provide a much better way to help them do that efficiently and align it.

Dr Teasdale: I'll just add to that. Beyond the access to compute itself, one of the benefits of a model where we might be able to deliver merit based access to compute across private and public is actually the skilled staff that are there to help researchers, government and business actually figure out what compute they need to be using, so they're using the right product, their project is ready to use and they're all set to go. So those things are not guaranteed. And, when it comes to really niche science topics, specialised staff are incredibly important.

Ms Williams: To give you a metaphor for that, there are many ways to bake many types of cakes. Do you want a vanilla sponge cake or do you need a chocolate gateau? I know that might sound ridiculous, but that's kind of what we're talking about when we're actually working alongside people to help them figure out what it is that they need to do. You can overengineer something where you could do it much more simply, and there's a cost-effectiveness thing to that as well. That's kind of what we need to do.

The final point I would make on this is that we talk about AI a lot and we talk about LLMs. There are a growing number of researchers and disciplines that are now in this space that would not have been here five or 10 years ago. They wouldn't have seen the opportunity for the research that they're doing. So, if we think about the social sciences, if we think about other disciplines as well, they're able to participate and actually provide a huge amount of public benefit in analysing things in new capacities, but they don't have the capability or skills necessary to do the technical side of that. So working in partnership with those skilled staff is really critical to deliver on those needs.

Ms CHANEY: You want data centre expectations to be enforceable with measurable reporting and independent assurance. I'm interested in what you think should be measured and who is well placed to do that assurance. Is it around power use and water? We probably have more capacity to measure those things. What's happening there in the data centre? What other things should we be measuring and tracking?

Dr Teasdale: Our key asks are that energy use and the types of energy being used—water and emissions—are measured and data released. Those are the main things. Obviously, reporting on the skilled workers that they're employing in Australia would be preferable. That's a separate kind of issue. One of the reasons for this is that the academy has been piloting ways that we could look at our energy transition in one spot, and to do that is actually very difficult. You're trying to find every single dataset, on a very different website, in a very different format. Obviously, those other datasets are all in different formats. But as long as we have something that is accessible, and we all know where to find it, and it is released on a reasonable schedule—preferably annually—that data could be used by bodies like the learned academies and others to ensure that our energy transition is not disrupted by this big explosion of construction in this space—and that, if that does happen, we make sure that we can respond in a timely manner.

Ms CHANEY: What about hardware verification technology? Do you think we need to understand what the models are being used for? If we're going to have training data centres in Australia now, what sort of oversight will that require? Is that beyond your scope?

Dr Teasdale: It is slightly. But we do support regulation and independent monitoring of the types of things that are being trained onshore and regulation of the AI tools that are being built, particularly when it's on Australian infrastructure.

Ms Williams: I'll just add that I think there are some really good parallels that we can look to. I spent quite a lot of time working in, supporting and developing medtech companies. We've more recently had this development of software as a medical device become part of our regulation for medtech companies. So we need to think about what's being done in these data centres and on these AI models is already going to be captured by other regulation—health will already have some level of that—and understand and get a measure of that. It's the things that will interface with human health and human wellbeing that I think we need to give some consideration. If there is an existing legislation appropriate, how would we do that? On a social level, I do think it's really important—and this is a personal view—to consider inclusion. We do want to have models in Australia that are not causing harm or making us a less cohesive society. So we need to consider what some of the inputs are to those and understand how the oversight is being done to make sure that we are doing no harm.

CHAIR: I'll just flag that we are getting very close to time. After Dr Keay contributes, I'll throw to one more question from Ms Mascarenhas. Dr Keay.

Dr Keay: With some of the things around how the models are tested, you don't need to have a look inside at the hardware. So trying to put in place mechanisms to look at the hardware may be challenging. But that also demonstrates why there is a fallacy that just having frontier models being trained on Australian soil imparts some sort of control to us. It's hard to see how that would actually work short of us trying to take over a data centre facility and its GPUs. Even then, that wouldn't give us control of the model. All we can really do is test them once they've been developed.

CHAIR: Ms Mascarenhas, the final question goes to you.

Ms MASCARENHAS: This is a question around type of compute. It was great to hear the suggestion of having a body that coordinates how that 2.5 per cent is donated to public research interest. When talking about the cake analogy, I know that different types are required. Do you have a sense of how we would make the mechanism or the market work to get the different types of compute? Could you give us a sense of what that actually looks like? I don't have a good sense of the different types. Dr Andrew Rohl started talking about this. But I have to say that, as members, we don't have a good feel for that yet.

CHAIR: That's fair.

Ms Williams: I'm very glad that Andrew did have a go at it. I will follow up with him to make sure that we come back to something that is digestible with a good metaphor.

Ms Williams: We will take it on notice so that we can provide you with a really nice description, just in case I butcher the one I'm about to give you. We have a system already between our two big computers we've talked a lot about today, NCI and Pawsey in particular. It's called NCMAS. It's a national computational merit allocation scheme. It works kind of like a grant program would. There's a certain amount available, and that's the amount that's reserved for researchers to bid for essentially. Researchers can apply for a certain amount of compute for what they need—maybe it's CPUs or maybe it's GPUs, depending on what they're trying to do, and, depending on when they need it and when it can be given, it is then allocated. It's essentially like a booking system, but it is on a merit based process. So the idea is that we would use something like expanding that to then include allocation to private compute.

Ms Williams: Just to add to that, I think it is an understanding of what is sitting on the private compute side and what that looks like. I think it's best placed to be able to technically allocate that and the datasets. If the datasets they want to do the research on are already existing in one particular location, then actually it makes more sense to do the compute there. You don't want to be moving large datasets between things; it's cost prohibitive and time prohibitive. So I think there's a coordinating effort that can be made to actually do that. I do think there's also a translation pathway. When research does hit a certain point, if people are looking at what is a commercial outcome from that, then there's actually a really good, long-term way of thinking about how that actually then moves fully into the private sector in a way that can take them forward and help lift productivity in the economy in that way as well.

CHAIR: I'm going to use the chair prerogative and ask one last question—sorry, Zaneta, you go.

Ms MASCARENHAS: I was just going to quickly say the thing that will be tricky for government is the different types of compute and what the ratios are against them in the whole pie that we allocate. If you've got views on that, it would be great to get that on notice. I'll stop there.

CHAIR: Just a final thing, also on notice. Just to pick up on the cost of moving data around, do we need to be also thinking about the connection as part of public access to compute?

Ms Williams: I believe we do. This is again where I step back and go beyond access and look at—we're talking about outputs right now, whether it's public compute or private compute, and where that's at. We actually need to come up a level and go, 'What is the national strategy for high-performance computing, both public and private?' and have an understanding of that. Again, this comes back to that kind of coordination and centralisation, which an RDC-type model could support. It's a knowing of where everything is and what it is, and then actually being able to efficiently allocate, co-ordinate and consider the long-term needs both publicly and privately and plan for that. That actually starts to help on both sides of the equation. It's not just about helping government and the public infrastructure and researchers; it's actually helping private companies over the long term to get their long-term settings right, which is advantageous on both sides.

CHAIR: Excellent. Thank you very much. We are just over time, so I'll end that there. Thank you for making yourselves available and providing us with your contribution to our inquiry. The additional asks that we put to you on notice—if you could provide that through to the secretariat in the next two weeks, that would be greatly appreciated.

BASSI, Mr Justin, Executive Director, Australian Strategic Policy Institute

BROWNE, Mr Bill, Director, Democracy and Accountability Program, Australia Institute

DENNISS, Dr Richard, Co-Chief Executive Officer, Australia Institute

WROE, Mr David, Head, AI and Security, Australian Strategic Policy Institute

[14:23]

CHAIR: Welcome. I understand that information on parliamentary privilege and the protection of witnesses giving evidence to this inquiry has been provided to you. I'd like to now invite you to make some short opening remarks before we get underway with questions from committee members. I'll start with the Australia Institute.

Dr Denniss: Thanks; Bill and I will both make some brief comments. Thank you for the invitation. It's good that the Australian parliament is looking into these issues. A lot of them have been well ventilated. We're happy to take questions on our submission.

I'd start with a general point, and that is that, unlike any technology we've ever dealt with, the business model of AI is to generate mass unemployment. That is not an unintended consequence; it is the business model. As an economist, I understand productivity; people lose jobs when technology changes. But the valuations of these AI companies are literally based on the assumption that hundreds of millions of people will lose their jobs and employers will buy AI services instead of labour services. Again, technology change drives labour market change, but the valuations of AI companies are based on the very rapid uptake of these services, and the economic consequences, I think, have been barely scratched at.

I would urge this committee and parliaments more generally to reflect on a few questions that to date have just not been considered. For example, how many jobs would need to be replaced with AI for AI companies to become profitable? We can calculate this, but I don't know that we've looked into it yet. How will the claimed benefits of AI be shared between Australians and foreign owned AI companies? What will the impact of the AI rollout be on electricity prices, water prices and construction costs, and what will it add to inflation? Will the government be spending more money to equip public servants with AI tools, or will the government be funding AI rollout through job shedding? Finally, how much will it be to retrain workers to take advantage of this new technology? As we shed huge numbers of workers, if AI goes according to plan, who will fund the retraining? I'm not anti AI. I am not a Luddite. I'm not anti technological change. But no other industry, no other proposal and no other policy would avoid scrutiny of these kinds of questions, and I suggest this inquiry is a great place to consider that.

Mr Browne: Generative AI has opened the floodgates to political propaganda. It has never been cheaper or easier to create false and misleading sound and video, and these are being used in fraud to incite hatred and to ridicule and misrepresent politicians and others in public life. It is not just in the political arena that the big tech companies have caused harm. Management consultants are fulfilling lucrative government contracts by handing in error strewn work cheaply written by large language models. Chat bots are helping teach the next generation of psychologists in Australian universities. The fight against AI generated spam was used as a justification for restricting freedom of information rights in Australia. AI search summaries are intercepting the ad revenue on which investigative journalism depends.

Artificial intelligence is doing harm now and it is doing harm at the behest of humans. There are understandable concerns about rogue AI doing enormous damage, even threatening human life. But the harm that we are experiencing now is a result of human weaknesses like greed and laziness, though facilitated by AI. In some ways this is a comfort. AI did not invent political misinformation, even as it makes it easier. AI did not invent poor-quality services, even as it lays a high-tech veneer over them. Addressing this will be difficult, but the solutions are at least familiar ones: more accountability from government, better investment in services, strict labelling of AI generated content so people know what they are getting, protections for people who have had their likeness misused and law reform that is human centred, not done for the benefit of tech companies.

Mr Wroe: I suspect there'll be some contestability on this panel. Thanks very much to the committee for giving us this opportunity. One of the chief risks to good policymaking on AI is failure of imagination. All countries face this risk, Australia included, and the work of this committee is critical in putting us on the right track. We were very pleased to see such strong interim recommendations delivered to the Prime Minister this week, and we congratulate the committee on those. It's hard to imagine something smarter than humans. It's hard to imagine something that can build successive versions of itself that are smarter still. Yet the people at the frontier are telling us that this is exactly what we will be facing in the next few years. The number of science fiction scenarios that have become reality this year alone—some exciting, some deeply troubling—should acquit

the industry of any charges of hyperbole. And it's not just the industry. The head of Britain's MI5, Sir Ken McCallum, said this week in his annual threat update as he warned of the risk of misaligned and uncontrolled AI:

… human history has given us no experience of facing an actor which is smarter than we are. Is capable of out-thinking us. This is genuinely novel.

Australia needs to prepare itself for the advent of broadly superhuman AI. An indispensable part of that will be carving out a role for ourselves in the global industry. Attracting US investment in training compute is a singular opportunity to build a foundation for our future sovereignty and strategic options. Our ally is the leading nation on this technology, and, as we have done in other areas, we should capitalise on that, which in turn means making a contribution. AI will be so decisive to the futures of nations that we should think about a kitchen sink strategy, one of 'everything and'. Our goal should be to capture the benefits of AI while influencing the direction of global standards to reduce catastrophic risks and to reinforce liberal democratic values. You have our submission, which contains our views in fuller detail, and we're happy to answer the committee's questions on that or any other matters.

CHAIR: I'll kick off. Despite perhaps diverging views, I think you gave us some interesting contributions on the overall issue facing not just this committee and this inquiry but this parliament and the country. This is a sort of amorphous technology for which, as you said, the companies that are at the frontier are now so big that they are eclipsing nations. My question is to both of you because I think you raised some really important questions that we should be asking and looking into, Dr Denniss, and, Mr Wroe, you talked to the strategic and defence imperatives that are in front of us at the moment. Given the fact that AI is here—it is happening; there are frontier models that are leading, and that is going to happen with or without us—what is the opportunity facing Australia, taking into account the really important questions that you raised, Dr Denniss, of actually playing a legitimate role in shaping that by leveraging off the interest that the frontier have in coming here and training here? How do we, in your view, truly capture that as an opportunity that takes into account some of those significant risks, not just to our defence and not just to our cyber capacity or those issues but also to our society, to our general way of life and the values that we have as Australians around looking after each other? You talked about job loss. How do we potentially use this in a way that truly benefits us as a nation? Is that a naive assumption, or is that something that is genuinely open to us at the moment?

Dr Denniss: I think it's absolutely open to us. To be clear again, for the avoidance of any doubt, I think AI is with us. I use it. People at the Australia Institute use it. I'm not afraid of a new technology. I think that in my lifetime I've never heard so much hype about a particular industry and I've never heard FOMO, the fear of missing out, being used so aggressively to suggest we can't talk about things before they happen. To be clear, I think the defence and intelligence community will obviously see opportunities and threats from AI that I can't imagine, but there are plenty of technologies that the military use that I do not have ready access to. The idea that, as a nation state, we might want our intelligence services to invest very heavily in AI is, I would suggest, unrelated to the question of whether data centres should be built wherever we want to build them. But this rush, this sort of generalised assertion that if we don't say yes to everything now we will miss out, is, I would suggest, unprecedented.

Do we have agency? Of course. Do other countries share our concerns? Of course. Are there people in the US who share these concerns? Of course. But, again, the way it's often packaged up is, if we, if this parliament, if the public are anything other than welcoming of every proposal, then we will be left out. I think this is unhelpful. I disagree with my colleague's submission. I do think we have to wait for public trust. I think to overtly say we can't afford to wait to build public trust is actually one of the problems not only on the issue of AI but on so many issues that our democracies face. I think it's a mistake to think we don't have agency. I think it's a mistake to think other countries and other communities don't share our concerns. For all of the talk about frontier models, from an economic point of view, it's not even obvious that the frontier models will be generating the economic benefits that we're so afraid of missing out on. The frontier models might be the absolute thing the intelligence community needs, but the idea that companies and organisations are going to replace their AI products every six or 12 months as exciting new models are released is fanciful. It would be the opposite of productive for all companies to be using everything at the frontier all of the time. Again, what makes sense for the intelligence community is unrelated to how Australia can maximise the benefits it captures and minimise the harms that that bill referred to.

Mr Wroe: I might pick up on the national security side of it first. Clearly Australia having strong national security is good for us. I would argue that it's also good for the world. Without waving a flag, we are a good international player. We are a liberal democracy who supports a rules based order. I think a strong Australia is good for global security as well as national security.

With respect to some of what Richard was just saying then about the value of frontier, I certainly think at the national security level we want frontier technology. We want the absolute best. If your adversary has Qwen3.8-Max, you do not want to be going up against them with GPT-3.5, because you will get smoked while you're still putting your boots on. Frontier is absolutely critical to national security.

You're probably sick to the eyeballs of the word 'sovereignty' at this point, but there is an important sovereignty discussion here, and that means building as much as we possibly can in Australia whilst recognising we can't do everything. Ultimately, it's about preserving our options and agency in areas of strategic importance. AI is as strategically important as anything the world is looking at at the moment. I would argue it's the most strategically important technology in human history. As I say, it means doing as much as we possibly can here, but, for those things that we can't do here, we want to form reliable partnerships with countries and partners that can actually provide those. Obviously, we are in an advantageous position in that our chief ally is also the leading country in providing these technologies.

The benefit of building training compute here in Australia is that it meets both of those objectives. We are doing something here for ourselves. We can build an industry around that. That will improve our domestic capability, and we are also securing access to frontier capabilities from a reliable partner but with some assurance underpinning that in the form of us making a contribution to that relationship. To come back to the question, from a national security perspective, I would say that that ticks all of the boxes that we want to be ticking from a national security point of view.

Mr Bassi: If I could respond a little bit to one of Richards points, I agree that there is a distinction between what the intelligence community needs and what the public needs. Where I fundamentally disagree is on the issue of public trust. I completely agree we need public trust, but we can't wait for public trust when we're dealing with something that is a geostrategic issue and is a national security issue. It is government's responsibility to build the trust and not wait until the public has that trust. In no other field of defence or national security should we be waiting to get the social licence before government acts. That is shown in history. If we look at World War II, the Americans had to wait to enter World War II until there was social licence, and it took Pearl Harbor. We don't want to do that again.

Dr Denniss: To add to the agreement, to be clear, I don't think the intelligence community need to wait to get public trust, because they'll never tell us what they're doing. I'm okay with that. But that is not the proposition that is generally being put. I'm not trying to put words in ASPI's mouth, but generally it's, if we don't say yes to everything that the AI industry wants, the intelligence community will miss out. I don't disagree with a word Justin just said, but I think that the needs of the intelligence community for access to compute are unrelated to the claimed productivity benefits and other benefits for saying we can't regulate this technology either in terms of the untrammelled rollout of data centres or restrictions on the use of fabricated content. I could not agree more I just think that that is a small part of a much bigger conversation and the fear of missing out is being used, I would suggest, to railroad an important public conversation where trust is key—about where data centres are built and how they're built and for what.

Again, we've always separated the intelligence and Defence establishment from those decisions. Imagine nuclear technology was just invented. Wind the clock back. We've just split the atom for the first time. Would we say, 'Well, let's just let everybody play with it and we'll see how it goes; we don't want to impede innovation'? Of course we wouldn't. But would we let the Australian nation-state play an active role in it? Yes. We have been here before, and I think the biggest threat we face in regulating this is conflating all these different elements of what is an important and complicated issue.

CHAIR: Yes, I would agree. I've made the point consistently through this, and in some of the questions that I've been engaging previous witnesses in, around the critical nature of social licence and how we bring people, the Australian public, with us in terms of how it is directly affecting their lives. Obviously there's the importance of having access to the frontier technology for our defence, which is also fundamentally about building trust for Australians as well, knowing that we've got the capability to keep ahead of our adversaries and making sure we're not missing out on that. But to your point—and this is part of the conversation we're having —what do the settings look like? What should they look like? How do we make it, as we have previously said, about shaping this and not just letting it happen to us? And, with all of that in mind, how do we build that social licence? Because fundamentally I don't think we get anywhere near the benefits proposed, whether they're real or not, if people don't want it or trust it or understand how it's used. I really respect that and take the point of the juxtaposition between the two but also the extent to which they are concurrent as well. I'll throw to my deputy chair.

Senator DUNIAM: Dr Denniss, you drew a distinction between national security and Defence and the rest of society. What's your view on organisations—whether they be government or private sector—that are custodians of data being able to access and use AI? What's your view on that?

Dr Denniss: I think most of them already are. And if they don't think they are, they probably don't understand how their data is being managed. AI is such a broad term. It's being used by organisations, often when they don't even realise it. To be clear, I like technological progress. It drives a lot of productivity growth and in turn can significantly improve the wellbeing of individuals and countries.

Senator DUNIAM: Just on that, the point's been made around trust. Say, for example, I'm an individual who's got my health data with my private health fund and it is storing my data overseas as opposed to onshore. I'm wondering how we align that with what you've said around it being okay for security and Defence to be utilising it and necessitating compute build-out here in Australia and data centres and all that goes with it. When it comes to a bank, a health insurer, the department of health or social services or whoever it might be and doing the same for them here in Australia, your contention is that we shouldn't be building things all over the place. We would then be sending things in a big cable underwater to the US or China or wherever it might be stored. I'm wondering how we deal with trust in that context.

Dr Denniss: Sorry, I misunderstood; yes, again, I think it's very important. Data centres exist now. There never was a cloud. People have just figured this out. There was a giant network of sheds. Anyone who's been talking about the cloud for the last decade has been talking about large sheds somewhere else. I'm 100 per cent in favour of some of those sheds being in Australia, and I would prefer that information—

Senator DUNIAM: How many? Do we have a cap on how many we should have in Australia?

Dr Denniss: Yes, but I was going to go to category. Do I think that Defence and intelligence should have access to the best models and the fastest compute? Yes, absolutely. Would I prefer Australian citizens' tax data be stored here in Australia if possible? Yes, absolutely I would. But, again, we need to think categorically about the uses to which things are being put are, because I put it to you that every data centre proponent is effectively suggesting that without it there will be no national security and all of our tax data will be stored overseas when actually, maybe, that's where the cat videos are getting made.

I think it's really important to build trust and public understanding about what is important new infrastructure and to be far more categorical and admit some of this computing is being used for national security, some of it's being used for research, some of it—and some of the storage—is being used for what we would think of as confidential, if not top-secret, information. A lot of it is being used to store 10-year-old emails and to whip up fake political images for ads. I think it's okay for us to be upfront about that and say, 'You know what? We have enough data centres to store the ATO data and the health data. If we want to build a lot more for those other reasons, we'll get in the queue.' There are competing uses for land. There are competing uses for labour. There are competing uses for electricity. Those cat videos don't have to be made here in Australia.

But, again, the debate is so overwhelmingly; it's so totalising. If you don't like data centres, you don't care about national security or the confidentiality of Australian information. I think a lot of sensible people can see all of these categories quite differently and I think the parliament and the public would be well served if we broke it up like that.

Senator DUNIAM: I'm struggling to think of another witness in this entire hearing who's actually made the same point that we don't need to try and catch the train and do what a lot of proponents, from the national security witnesses we've had through to the tech providers, news organisations, academia—the academy of sciences said exactly the same thing—have said. If your contention is to be pursued to its conclusion, I'm wondering where you draw the line and say, 'No, you can't have that data centre because it's going to be used to make videos of cats,' versus, 'You store tax data.' How do we recommend to government how we delineate between cat videos and tax data, according to the Australia Institute?

Dr Denniss: Yes, that's a fair question. I'm surprised that no-one else has attempted to make these points. I don't mean to be flippant when I say this: I actually think it's kind of the job of the committee to think about how you might draw those lines. I'm not here with a concrete proposal on how you should do it, but, unless we can, I think people will not only lose faith in AI but, increasingly, lose faith in the ability of our governments and our parliaments to stand up for us, because people don't actually believe that all the data centres are mission critical for national security. They do not, and if we—

Senator DUNIAM: No, not when you tell them they're being made to make cat videos.

Dr Denniss: Where the lines are drawn is, I agree, difficult, but unless some lines are drawn, people will not accept the—well, let me phrase it in more concrete terms. There is a real risk the intelligence community will

miss out on what it wants and a real risk that the banks and the insurance companies will miss out on what they want, unless the public believes that there is some form of stratification going on here. I'm not suggesting it's easy, but, unless we're honest with people that not every data centre is a critical piece of national infrastructure with significance for defence and confidentiality and unless we can find a way to have that conversation, I fear the lack of public trust will actually harm what it is that ASPI are here telling us we really need to do. I'm not suggesting it's easy.

Senator DUNIAM: Is it your view that data centres are used for one purpose exclusively, or do they have multiple applications through their compute load?

Dr Denniss: That's a very good question. A lot of people with far more knowledge of it will meet with me and say they have no idea, because we're not allowed—

Senator DUNIAM: I'm asking you though—

Dr Denniss: Well, I guess I'm saying that there are people that own data centres that I've spoken to that said, 'We have no idea what's going on inside our own data centre.'

Senator DUNIAM: The point you've been making is that we can have them if they're used for certain things, but we shouldn't wave the green light around for them if they're not used for those purposes. A lot of evidence has suggested they're used for multiple purposes. You might have 50 clients doing different things through that centre—one might be a cat-video maker, the other might be ASD. I'm just trying to figure out how we deal with that.

Dr Denniss: At a technological level, I am the last person to ask. I'm not pretending I came here with an answer to that question. But, again, I think it raises exactly the point that I'm trying to make: no-one's in charge of this if no-one even knows what's going on inside the centres that they literally own. Is it possible that the cat videos and the ASD computer are sitting side-by-side? Of course it is. But, again, if we're going to tell the public, 'We are going to need to multiply the number of data centres tenfold and the electricity demand tenfold'—if we're going to take people into that project—we might be able to get away with it today, but as this rolls out, as the electricity price rises, I think if we don't confront these categories now we will be having much more difficult policy and political discussions in three or five years time. I agree it's difficult. I really do.

Senator DUNIAM: Mr Wroe?

Mr Wroe: I'd like to make the point—thanks, Deputy Chair—that no-one I know who believes in Australia carving out a strong role for itself in AI has the faintest interest in data centres for the purposes of storage of cat videos or ten-year-old emails or anything like that. The interest is in data centres that are optimised for artificial intelligence and, particularly, for training.

Yes, there are different types of data centres. There are different types of chips. I'm not a superexpert, but there are CPUs, there are GPUs and there are TPUs. GPUs—the high-end things that NVIDIA designs—are very sought after. They're about the most in-demand products on the planet at the moment, which is why NVIDIA is the most capitalised company on the planet. TPUs, tensor-processing units, made by Google, or Alphabet, are what we are talking about here. That is really what we're talking about now. Yes, they can be used for different things, but you would be wasting a GPU by using it for storage, for instance.

We're very much talking about attracting a very specific type of data centre investment here with a very specific purpose. One is to build local innovation clusters around that, and I think history is pretty kind to the idea that, if you build an anchor point, then an industry tends to flourish around it. Look at Silicon Valley, for instance. The other one is that we build those partnerships. Some people are talking about it as leverage. Some others prefer to avoid terms like leverage. I don't think it should be a perpetually adversarial position, where we are constantly saying to the Americans, 'If you don't give us the latest version of GPT-6 Astra or whatever, we will cut you off.' It's just as with Pine Gap. We don't tell the Americans every second year, 'We will close Pine Gap unless you give us the latest whiz-bang radar,' for instance.

The ideal framework is to think about it as an evolved partnership where we have an understanding and we share those two things and we each make a contribution to the relationship. That is very much what we're talking about. The broader public concern around data centres, for all of the local reasons that people have concerns about them—power prices, they are aesthetically displeasing in some places, they can be noisy, they can consume water—are all legitimate and should be managed. The discussion I've seen from the government and the parliament so far, I think, is very encouraging, but this is an AI data centre proposition that we should be focused on.

Mr Bassi: Just quickly, if I may, I do agree with what Richard was saying around data centres and critical infrastructure. I think it's incumbent on us to know what we view as critical infrastructure versus just

infrastructure. I think one of the lessons we can learn from the last 20 years is that we did not have those lists. We did not understand what we viewed as critical, whether it be technologies or infrastructure. I think that empowers us to be able to separate a building that is being used for a particular purpose and a building that is being used for something that is nationally sensitive or nationally critical. I do think that it's incumbent on us, whether it be parliament or the different sectors, to be able to say, 'This is what we mean by critical infrastructure.'

I do think a subset of AI data centres should be viewed as critical infrastructure and then, to have the flow-on impact of that, that it shouldn't be owned or operated by a foreign company or a foreign adversary—those types of issues. It would also mean that, if you're not classified as critical infrastructure, then—I think this is part of what Richard is saying—you have to prove why you might be proposing to set up a bit of infrastructure if it's not going to be determined to be critical.

Senator DUNIAM: Thanks.

CHAIR: Ms Chaney.

Ms CHANEY: First, Richard, I'm interested in understanding your perspective on this more because it seems like you agree that there are some zero-sum domains where small capabilities make a difference, such as national security and cyber, where enduring access to frontier models will be important. I guess I'm interested in testing: do you also agree that we don't control those models and we can't replicate them at any feasible cost? So we have a broad national interest in ensuring that we have enduring access to those. Do we then need to have a bargaining position or, as I think ASPI put it, a partnership where we both bring something to the table? Wouldn't holding a key part of the supply chain give us that bargaining chip? Isn't the binary nature of access to frontier models in the context of potential recursive self-improvement a big enough win that means it justifies building a part of the supply chain in Australia so we continue to have a seat at the table, and also have a seat at the table on determining global standards so that we have a little bit more power to actually control the direction that AI might be heading in?

Dr Denniss: There's a lot in that question; I might not get all of it in my attempt at a short answer. To be clear, do I think that Australia should facilitate the construction of some data centres for some critical reasons? Yes, I do. Do I think there are benefits to Australia in trying to enmesh itself in the AI industry? Yes, I do. I'm not an expert, but I've been reading a lot for a long time about this. Where I think there's a lot of confusion, which perhaps the committee members don't have, but I do, is about what this supply chain is. I think there's a lot of confusion about what AI investment is. I think there's a lot of confusion about what it is to train a model and what the lasting benefits of those models are. Again, if the committee doesn't share my confusion on those key terms, then that's okay—

Ms CHANEY: I think we've heard evidence on all of those issues—

Dr Denniss: I've heard lots of it.

Ms CHANEY: and there is uncertainty around them as well. They're not all settled conclusions.

Dr Denniss: The reason I say all that—I understand the question you're asking, but I'm limited in my ability to answer it—is it's not at all obvious to me that the ownership of a concrete slab and a shed on which some GPUs that may or may not be owned by Australians, that piece of physical infrastructure, gives us really anything at all. It may or may not. We might describe the mere act of laying the slab and installing those GPUs as placing ourselves in the supply chain of AI, and it might feel like it, but without far more information than I'm able to get my hands on about the nature of contracts and the nature of the economic relationships between all the people in this supply chain—the metaphor of supply chain. If we think of it in terms of how a fighter jet is made and the role Australia plays in the supply chain of a fighter jet, as an economist I understand that quite well, but we're using metaphors like 'supply chain' in AI when it's not at all obvious to me that the economically significant parts of the industry are in any way linked to who owns the slab and where the GPUs are. That doesn't mean we can't do it—

Ms CHANEY: I suppose there are the geostrategic and the economic arguments, which most people have treated as being separate. The position we find ourselves in is we are able to impose some conditions on that happening. So if we've got something that someone else wants, we can actually say, 'You can do it on these terms,' and determine the terms in which it is good for Australia too. I think that's what the committee is trying to get to. What are those conditions that make this a good move for Australia, both economically and strategically?

Maybe this is a question for ASPI. You're talking about hosting giving us a stronger bargaining position or an evolved partnership to get frontier model access. I'm interested in your views on the extent to which you think hosting would have protected a country from being subject to the export control that the US imposed in June.

What is the bargaining power then, and can we imagine it in real terms, if the thing that we're concerned about is the US government saying, 'We're withholding access to this model'?

Mr Wroe: Common sense tells us that, if we are instrumental to the development of current and future models, we are in a better position to say to both the frontier companies and the US government that we are entitled to privileged access, early access and some sort of parity position with US users. I think that is a pretty logical position. The exact mechanism through which that would be done is a discussion that a lot of people are having at the moment. ASPI is actually doing some research precisely on that at the moment, which, unfortunately from the committee's point of view, we will report in early December. We are looking into precisely what those government-to-government agreements would look like, what the commercial agreements would look like to make sure that there is some substance behind that. It is a reasonable, logical statement to say that's the case.

Can I just come back to your previous question, Ms Chaney, if I may?

Ms CHANEY: Yes.

Mr Wroe: My answer to your previous question was just a resounding yes in the sense that the role of frontier—and you mentioned recursive self-improvement as being relevant here—is going to become immensely powerful in a pretty short window ahead of us. We're talking about 12 to 18 months. That's the prediction window at the moment coming from the frontier about when recursive self-improvement might become possible. There is new data coming out; there was something that came out just this week that showed what is called research taste, in which frontier models, rather than just writing the code and doing the mundane routine stuff, actually have the new insights to say, 'Hey, what new ideas can we apply to iterate current frontier models to design the experiments, to test the efficacy of that and then to analyse the experiment results afterwards?' Frontier models are getting much better at that, and it's just in the last 10 or 12 months that that has occurred. 2026 has been a remarkable year for AI progress, particularly at the frontier, and I think that in 12 months from now we are going to be seeing models that we would think of as science fiction superpowers this year. I do think that's immensely important. We've talked about FOMO. I could be wrong about this, but let's say there's a 50 per cent chance that I'm right. I think we should feel FOMO about missing out on the opportunity to be part of something that is going to be that powerful in such a short window from where we're sitting today.

Dr Denniss: I largely agree with all of that, but, again, to your question—'should we try and get as much out of it as we can?'—of course. Conceptually, what I'm trying to suggest is: Who is it that we would be having this negotiation with to get the best we can? Is the Commonwealth of Australia talking to Anthropic? Is the Commonwealth of Australia talking to the US government? Are we assuming that all the frontier models in five years' time will be coming from the US? Are we assuming that Anthropic won't be broken in three years and a company we haven't heard of yet will be the company that we'll wish we had? If your question is 'should we try and get as much out of it as we can?' yes; it's a resounding yes. I just think it's harder—conceptually harder and technologically harder and politically harder—to answer the question: is having the slabs and the GPUs going to get us our fair share? We have to have a really sophisticated understanding of 'fair share of what?' And 'from who?' If the US can renege on its free trade agreement with us in a heartbeat, what deal is it we think we're going to write that will lock in the most powerful country in the world after it's got the most powerful technology in the world? What piece of paper is it we think we're about to write that is going to protect us and ensure that they share fairly with us in five or 10 years? It's just bigger and harder than we're talking about.

Ms CHANEY: I take your point. It's pretty hard. Unfortunately, we don't have an alternative in the circumstances that we're trying to make decisions in. I've another question for ASPI. In your recommendation 2, you talked about funding the AI Safety Institute to do independent frontier evaluation, cyber and bio assessment, incident analysis and pre-deployment access. What sort of capability uplift do you think is required for it to fulfil that role?

Mr Wroe: In dollar terms? How would you like me to quantify it?

Ms CHANEY: Whatever terms you're happy with.

Mr Wroe: Sure. I've seen some of the other figures that have been mentioned. The UK is the gold standard here. I think matching the UK in some sort of proportional GDP terms is a pretty good way to think about it in dollar terms. People are going to be absolutely critical here. It's a very select set of skills. The reason the UK institute has been so effective is that they have been able to convince people, fresh out of the frontier companies, to come and work for them. People of that level of experience and talent are really crucial, so I think we should be aiming to get some of those Australians who are currently working in Silicon Valley. There is a bit of an Australian mafia at the moment—I mean that in a kind way—working at the frontier labs. Some of them are

possibly about to cash out with significant shareholdings if some of the frontier companies do go public. Perhaps, out of a sense of national duty, they can be persuaded to come back and work for something like the safety institute. I think focusing on people is absolutely going to be critical.

And then what can we do? In terms of the particular work they do and what the uplift is there, some of the things, like agent governance, we've already established and got runs on the board for, and that is only going to become more and more relevant. If you're working in agent governance at the moment, then you're doing really important work. There are organisations, like GovAI in the US, that have done some great work on this, so I think the Australian AISI should look at partnering with organisations that are doing a similar kind of work. That should include the UK AISI, and it should include organisations like GovAI as well. That's where I would like to see them put their priorities. Like I say, I think something roughly comparable on GDP terms with the UK institute would be useful.

CHAIR: Sorry to cut you off just at the bottom there. We are beyond time now, and I know Senator Sheldon has one question that he'd like to be able to ask. Perhaps a response that's short and sharp, but a fulsome response on notice, would be helpful. Senator Sheldon.

Senator SHELDON: It's good to see you, Richard. This is to Dr Denniss, and, by all means, to others who might want to make some comments as well. I might go to, first of all, that last question, just briefly—the question that you raised as a figurative question to the committee. A lot of the evidence we've received has answered, to a degree, that proposition you put, so I'd strongly suggest having consideration of that evidence we've received on the committee because it does go to a lot of those questions. But I want to go back to what you initially said, right at the beginning of your presentation, about—I'm saying it in my words—not only looking at the technical questions but also looking at the humanitarian questions and the impact on society and workplaces, amongst other areas of deep interest. Can't we walk and chew gum at the same time? This is really a matter of the data centres having their place. The important issues we've got with copyright have critical questions that we have to resolve there. Data centres have critical questions. Workplaces have critical questions. Next year, or in 2028, we're going to have, supposedly, driverless vehicles being tested more substantially in Australia.

There are critical questions about 125,000 taxi drivers and rideshare drivers. How does the economy possibly absorb that? You raised the question of whether we are preparing ourselves for where there are impacts on parts of community. Have we got the training and retraining? Who's paying for it? Do we want these companies here in the first place? In some of the surveying that I'm seeing, most Australians don't want to go into a driverless vehicle, and I don't want to see all the profits and taxes go over to the US whilst we lose all the jobs, opportunity, income and taxes in of our economy. That's a personal view. So there are lots of questions to be answered, and I agree with you, but every piece of this has its own innate importance. I particularly liked that you're highlighting the fact that there is a critical piece that we need to be constantly thinking about: what impact is it actually having not on building a data centre, as important as that is, and not only about what happens or does not happen with copyright, as critically important as that is, but on what happens to the rest of the community, whether we have data centres or not or whether we have copyright or not? Is that the point you're also making to us—don't get fixated on lots of issues all at once?

Dr Denniss: Absolutely. Again, to be clear: I'm not anti AI. I'm a user of AI. The Australia Institute's initials are AI! I'm not anti AI. I think that the public not just in Australia but around the world are being railroaded through this sense of FOMO. That doesn't make the technology bad, but I do think that, as Bill alluded to at the beginning, we are going to have elections where fake depictions of real people pretending to say things they've never said before are going to dominate our analysis of what's happening in an election. That's a national security risk, I put to you. Our democracy. And it's the AI industry saying, 'Oh, there's a 10 per cent chance we might end humanity.' If I came in here and said that, people would suggest I'm being hysterical. I have no idea. Maybe it's only a five per cent chance. So I'm not anti AI, but I'm deeply concerned.

Senator SHELDON: I'm not putting to you that you are. I'm just saying that you're raising, in my mind, a legitimate series of questions that the committee should come to grips with about what we do about it now in any interim report and also what we recommend for future steps. I think the broad subject matters you're raising is that question about needing to walk and chew gum at the same time.

Dr Denniss: Yes, absolutely. We've been rolling out AI for five years. We're talking about speeding it up. We've been rolling out data centres for 10 years. We're talking about speeding it up. Again, the way the debate is often projected is that there's a binary: Are we in or are we out? Should we start or shouldn't we? But the reality is we've been doing a lot of this. It's already having consequences. While understanding the limitations of what any one committee can do, I urge you to look at those big questions that I put at the beginning. I think part of it is to talk to other parts of this parliament. Other parts of the parliament need to be getting their head around this pretty

fast. Again, I'm taking the AI industry at their word. If these companies are worth a trillion bucks, it's because a lot is going to happen really fast. If they're right, a lot is going to happen really fast and we're not ready. If they're wrong then we shouldn't be rushed into doing everything they're telling us. So this is a real challenge for our parliament. It's a lot harder than 'is it good or is it bad?' But, literally, those valuations of those companies suggest a lot is going to happen in a very big hurry, and I don't think that most of the public is aware of that. I don't think most of the parliament is aware of that. I don't think most government agencies are aware of that.

I don't think you'll be able to answer all these questions, but I do think that what you're in a great position to do is to tell everybody that they've got to get their skates on—not come out with the answer or decide whether we're pro or anti but start literally thinking about, if Anthropic is going to be that big in five years time, if their revenues will be that big in five years time, where the money is coming from. It comes from not paying workers. The only way the Department of Health can spend $1 billion on AI is if it spends $1 billion less on humans or if you guys increase their budget by $1 billion. I don't know what the Department of Health's plan for AI is. Maybe there is one. What is the ATO's plan? What are the state government plans? Again, all that money that Anthropic is betting on is coming from somewhere, and that's going to have very real consequences for individuals, for communities, for the economy, for the way Commonwealth budgets are drawn up. At the moment, most of the people talking—it sounds like most of the people giving evidence to this—are people who are really interested in the tech. Well, I'm really interested in the macroeconomic consequences. I'm really interested in the democratic consequences. Again, thanks for inviting us along.

Senator SHELDON: It's that thing about scaffolding and how we deal with ongoing challenges, which that scaffolding should give us the capacity to do whilst we're trying to predict, prepare and respond.

Dr Denniss: Very much so. And, if you guys don't do it, the question is: who is?

Mr Bassi: For me it's not so much a fear of missing out. I don't think we will miss out. It's a question of what we will get anyway. If we don't do it, in my view, we'll get the Chinese technology if we choose to do nothing. I know that's not what Richard is saying at all, but we're not going to miss out on this technology. The technology is going to come here. We just have a chance to determine what that technology is and how much control we have of it.

CHAIR: I will bring the session to an end. We are quite a bit over time, but these are really important conversations and important contributions, so thank you very much to both of you. If there is anything that you've indicated that you would like to provide further on notice, please do so through the secretariat within the next two weeks.

Proceedings suspended from 15:16 to 15:39

BOLAND, Ms Marie, Chief Executive Officer, Safe Work Australia [by video link]

MACDONALD, Dr Fiona, Policy Director, Industrial and Social, Centre for Future Work [by video link]

CHAIR: Welcome. I understand that information around parliamentary privilege and the protection of witnesses has been provided to you. Thank you so much for making yourselves available. I'd now like to invite you to make an opening statement, at the conclusion of which we will get under way with questions. Would anyone like to nominate to go first?

Ms Boland: I won't be making an opening statement. I'm happy to go straight to questions.

CHAIR: Thank you. Dr Macdonald?

Dr Macdonald: I have a brief opening statement. I thank the committee for the invitation to participate in the hearing today. The Centre for Future Work is a not-for-profit independent research centre that conducts and publishes research into a range of labour market, employment and related issues. Since its establishment in 2016, the centre has undertaken research into technological change, labour markets and employment, workplace surveillance, and digital platform work in care and transport sectors. Our most recent work, which is the basis of our submission to the committee, concerns artificial intelligence, algorithmic management, workplace surveillance and data protection at work. Our submission addresses the terms of reference (g), (i) and (k) regarding the workforce implications of AI, the adequacy of existing laws and the risks and harms from AI.

The centre recognises the potential for positive outcomes of AI innovations in the workplace and we welcome these. We also know that many of these innovations come with risks for organisations, for productivity, for job quality and for worker wellbeing. Our research into AI and algorithmic management practices—which is drawing on Australian and international academic and workplace practice literature to understand current and emerging practices and the consequences of these—we are undertaking to explore whether and what regulatory reforms are needed to protect workers from harms caused by algorithmic management. Our research into these matters also draws particularly on my and other colleagues' research experience over several years on digital platform work.

The promise of algorithmic management, as with the promise of AI for all kinds of jobs and tasks, is that it will improve efficiency and productivity. However, the evidence suggests mixed results in regard to the achievement of these benefits. At the same time, there can be little doubt that algorithmic management practices are generating unacceptable risks and harms, with significant negative consequences for workers. The growing body of evidence includes consistent findings of negative consequences on working conditions, health and safety, workers dignity and freedom to exercise rights in the workplace. The regulatory frameworks most relevant to algorithmic management are workplace relations, privacy and work health and safety laws. The gaps and shortcomings of these frameworks and their current operations have been identified and documented, including in the final report of the recent inquiry into the digital transformation of workplaces.

We believe new regulatory responses are urgently needed and should draw on some of the approaches taken elsewhere, including the European Union's Platform Work Directive and the very recently passed Californian workplace technology accountability laws. Both of these regulations contain prohibitions and limits on AI practices in certain circumstances and for certain purposes. They mandate processes for accountability and transparency and for impact assessments. They require human oversight and mandate worker participation and involvement, including ensuring that workers have the information and resources needed and an understanding of AI to be able to have input. Thank you for the opportunity to make that statement, and I'm happy to take any questions the committee has.

CHAIR: Wonderful. Thanks, Dr Macdonald, and thanks Ms Boland for your appearance before us today. I think we've heard some initial evidence. I expect we will be hearing a bit more next week. We've got a number of unions and their members appearing before us around the impact of AI in the workplace, particularly regarding the rollout and the combination of, as you said then in your statement, algorithmic management and the extent to which driving efficiencies is coming from an application of AI over the top, or in, perhaps, the risk of little to no consultation with the workforce in terms of how it actually works on either the shop floor or wherever people are. This is obviously something that's quite important to this committee for its considerations of how we look at the deployment of AI and the safety of that for people, both in terms of their working lives as well as more broadly.

At Safe Work Australia, Ms Boland, you have been doing quite a lot on how we could potentially use workplace health and safety as the frame in which we can work through, ensuring that there is decent and meaningful consultation with workers on the introduction of AI into the workplace. Can you expand on some of that thinking for the committee, and what you've been looking at. And perhaps following that, from a Centre for Future Work perspective, what does it look like in your capacity to make sure we are getting the systems right so that we are able to see in which ways we can potentially improve not just efficiencies but also improve people's

working lives. How can we think about this in a way that amplifies or augments people's work, and makes it more meaningful, more enjoyable by utilising AI to do some of the more mundane tasks. Those are two questions that I'm just going to throw out there, and I'll throw it to you Ms Boland first, thank you.

Ms Boland: Thank you, Chair. I guess the first thing we'd like to reinforce is that the model work, health and safety laws are principles based and technology neutral. They were specifically designed to be flexible and capable of dealing with changing work, changing ways of work, changing technology over the years, including the introduction and use of AI. The existing primary duty of care, which is section 19 of our Model Work Health Safety Act, requires businesses and employers to ensure the health and safety of workers, regardless of the source of the risk. This includes the physical and psychosocial risks arising from the use of AI. It's important, I think, to remember that, regardless of their level of autonomy, AI systems are tools that are designed and used by humans to undertake specific tasks. Under our current framework, businesses and employers are responsible for managing the work, health and safety risks of AI, so they can't shift that liability to AI systems. The starting point is a work health and safety framework that establishes the duty.

I listened to some of the previous people, and there was that question of safety in design but also safety in implementation. It's important to also recognise that the work, health and safety duties recognise that journey, that chain of responsibility. There are duties on upstream duty holders who are the designers to ensure as far as reasonably practicable that their products are without risks to work, health and safety, but there's also a recognition then that there are duties on the businesses and employers that implement those systems, and all of those duties as they apply mean that the entire framework in our model laws apply. That means that all the rights and all the obligations contained in the model framework apply. That includes the duties to consult with workers, all the health and safety representative rights, and the union right of entry. It all kicks in.

The other thing I think is important to emphasise as well is that really someone says, 'How do we meet this duty in practice with AI?' It's just the fundamentals. We know what we need to do, and we know how to do it because we've been doing it for years. AI is new, but the obligation to consult with workers has been around for decades. The approach is really that same basic risk-management approach. You consult with your workers to identify the risks, assess them, control them, and constantly review and monitor them.

That issue came up earlier today about a job imagined as opposed to a job done. Again, the key to that is those consultation duties, and we know in risk management for any health or safety risk in the workplace that the people who ultimately know best how to manage risk and identify it are the people on the shop floor, the people doing the work. I obviously can't emphasise enough the importance of that consultation piece not only being good practice but in terms of the legal obligation as well. That's a big-picture overview. I'm happy to be more specific as the committee requires.

CHAIR: Thank you.

Dr Macdonald: In relation to that first question and in response to Ms Boland's answer, there are three points I'd like to make in relation to the current work health and safety approach and legislation. Ms Boland has emphasised the importance of that consultation piece, and I think the difference between algorithmic management and other technologies is the potential for it to absolutely change workplace relationships, to change power in the workplace and to increase information asymmetries. Workers may have no idea of what algorithmic management processes are in place and how they're being used. The capacity for the reliance on identification of problems is pretty significant, I believe, in our workplaces.

My experience is in looking at work, health and safety in the home-care sector, which in some ways has some parallels in that workers don't have opportunity for consultation, don't know where decisions are made, don't know how decisions are made, don't know on what basis decisions are made, don't even know what the decisions are and don't know how they're evaluated. Those workers are actually subject to algorithmic management practices in these now.

But, where workers don't have opportunity and have less access to managers—which is a feature of algorithmic management in the workplace—less access to peers and less general peer communication, they're subject to communications around their performance that may be very individualised, such as: 'You attended four meetings last week. That seems like a lot.' The ways in which algorithmic management systems give individualised feedback to workers actually puts them in a much more vulnerable position than where those kinds of communications are made more generally and are discussed between managers and colleagues. I think there's that particular problem.

There is also the problem of managers' understanding of algorithmic management systems and how they operate and what AI systems are collecting and what information they're collecting and how decisions might be

made where algorithmic management is used to substitute rather than augment managers' decision-making. That is something that a recent ILO study of human resources practitioners suggests; increasingly, with AI, all algorithmic management is AI based and is seeing much, much more replacement rather than augmentation of management functions. So I think there are some significant changes that may impact on the ways in which the current frameworks can be applied and amplify the problems that already exist in the practice, if not—

CHAIR: It's almost like anything. AI is the tool which amplifies existing challenges or difficulties that were already there. It's just, with this new technology, that it enables a scalability and amplification of it. We do quite rightly and importantly focus on the significant concerns many people have, in particular about the rollout, adaptation or deployment of AI across our workplaces and how it's starting to shape and change in adverse ways people's experiences. Do you have any examples of where it is working well—of where there are good cases of proper consultation with workers and there's genuinely this idea of augmentation as opposed to replacement or displacement and making working life better? Are there any examples that you can point to where it's working well?

Dr Macdonald: Can I take that question?

CHAIR: It's for both of you, yes.

Dr Macdonald: The international research is suggesting that there are good examples. Those examples of where it is working are where it does involve a lot of consultation, where workers are involved—I was speaking to a German colleague this morning—where works councils are involved, where it is being used to ensure that AI applications are actually augmenting work and, where data is being collected, it is being fed back and used in positive ways that involve workers. Fundamentally, the problem with algorithmic management is that it involves the collection of masses of data and it involves a lot of surveillance. But where there are limits on that surveillance and where workers have trust and confidence, are involved and don't have any reason to be concerned about what's happening with that data, it seems that there are positive outcomes. There is a growing number of small studies. Some of them are not yet published in academic literature, but practitioners in Europe who have been involved in those are talking about those.

Ms Boland: Those issues that Ms Macdonald raised around power dynamics, lack of agency and unheard voices are consistent across workplaces no matter what, sadly. I think there's always been an obligation both within the Fair Work system and obviously in our system to consult when there's any workplace change. The places where I think the success is, as has been identified, are where those consultation obligations are taken seriously and are almost like moving from consultation to co-design. The element of trust again is just so huge in this. But I think that's why our framework is important too; we have, as I say, all those rights and obligations within it around electing health and safety representatives to be the voice. I know that that doesn't work in every industry and every workplace, but I suppose the message I want to give, too, is that there are very strong foundations within that work, health and safety framework that we can build on, around support for workers, and support for businesses, as well, who are saying, 'This is all new technology. What do I do?' but kind of clarifying for them that the approaches that we've been talking about for years are the same approaches.

Some of this is a compliance and enforcement issue as well. I think there's a real role for work, health and safety regulators, but also Safe Work Australia, to build the capability and the knowledge and provide the certainty about how our laws do apply, because I think, for a lot of us, as is becoming clear even just from dipping in and out today, it's kind of, 'Oh my God, what do we do?' instead of stepping back and thinking: 'Okay. We have all these systems in place. We have all the good frameworks to build on.' For Safe Work Australia, we need to be on guard, as Dr Macdonald has said, to identify risks and act quickly—or not risks but, rather, gaps; if there are gaps, we need to move. But at the moment we feel that the foundations are strong, particularly around [inaudible].

CHAIR: What I'm hearing from you is that we've got some pretty robust foundations, when it comes to workplace health and safety in particular—that duty—but it's then how we apply it within this seemingly evolving or novel way in which AI is changing people's workplaces and work experiences. Is it about how we better apply it or being consistent across—

Ms Boland: It's more about reinforcing that how we apply it is back to the basics—the fundamental risk-management basics. It's not even about, 'Oh, there's a new way of applying it.' It's about educating businesses about their obligations, educating workers about their rights and also making sure regulators are clear about how they will enforce it. Within that, there are challenges that we need to address to make sure that we're clear to people. I think, in terms of the work, health and safety framework, our first step is to be clear about how they currently apply, then support people to do their risk assessment and their consultation within that framework, but then, at the same time, monitor closely if there are any big gaps.

CHAIR: Excellent. Ms Chaney.

Ms CHANEY: That's been useful. I think you've pretty much answered this one, but I'll just give you a chance to add to it. On consultation: I'm just wondering if there's anything at all that's different about rolling out AI from any other change that means there are different requirements in consultation. You mentioned that it involves gathering a lot of data, for example. Are there any characteristics that mean that employers need to be tweaking their existing processes?

Ms Boland: From Safe Work's point of view—if you're talking about the algorithmic, which I think our researcher is—the thing I would identify there is that it's probably more of a stronger focus on the psychosocial risks, as has been highlighted with those kinds of management systems. The risk is the increase in job demands and some of the examples of the expected pace of work. There will be a lot of implementation risks that will need to be identified and consulted about, firstly, around job control—that's a key one; secondly, organisational justice—around workers feeling they have no capacity to question; and thirdly, surveillance, as has been identified. While we wouldn't say there's a different approach to consultation, it would be that businesses and employers would need to be very aware of the risks that they need to be looking out for. We do have a set of psychosocial regulations within our framework too. That's quite clear and specific about step-by-step approaches. I would suggest that that would be a key area that businesses and employers would need to look at in consulting with their workers in that context of algorithm management in particular.

Dr Macdonald: Can I just add to that. In some of the European—for example, the platform work directive requires that employees have access to an expert who can explain the AI application and what it's doing, acknowledging that that's very difficult for workers to actually know about. In some cases, it's very difficult for managers to explain.

This is slightly a different point, but I think there are other aspects of algorithmic management, similarly, which are dealt with in the European legislation, that actually put in limits and prohibitions in relation to psychosocial risks and health. They actually ban the collection and use of data that is, I think, often called emotion engagement or that makes inferences about people's feelings and attitudes through taking metrics from people's facial expressions and the way that they speak. There are some very well-known risks of harm from some types of data collection and in some contexts. Rather than that being dealt with through a risk assessment approach, that is actually dealt with through recognising those, banning those at the outset and putting limits on them.

Ms CHANEY: I have one other question, which may be outside your remit. There's obviously a fair bit of uncertainty about impact on jobs over the short, medium and long term. I wonder whether either of your organisations has done any scenario planning about those different scenarios and whether that longer term planning is useful or if you know how much of it is being done or who's doing any in Australia.

Ms Boland: I know that the Department of Employment and Workplace Relations have a chief economist in there who's doing a lot of work on this. That's certainly the area we look to for that kind of information.

Ms CHANEY: Okay. I think they're doing stuff based on historical data rather than future scenario planning. You're not aware of anyone who's doing that forecasting work?

Ms Boland: No, I'm not. The only data information I've been aware of is that work that DEWR are doing.

Dr Macdonald: I'm not specifically aware. I'm only keeping up with the research on how jobs are being impacted and how exposure measures are not really being very helpful in terms of who's going to lose their jobs. But I think that the present—

Ms CHANEY: As in they're not predicting it very well?

Dr Macdonald: Not forecasting it very well. For example, the recent research into call centres shows that the lower skilled workers are actually benefiting more and are more likely to keep their jobs than some higher skilled workers because the AI assistance is actually upping their performance to levels that mean that those workers can do the job that other workers had previously done. I'm suggesting it's tricky, and it appears to be changing constantly.

Ms CHANEY: Thanks.

CHAIR: Senator Sheldon has some final questions.

Senator SHELDON: To follow through on Mr Chaney's question, Dr Macdonald, do you think there's a gap in the amount of research that's been done on predicting future scenarios of employment?

Dr Macdonald: As someone who's tried to keep up with reading them, not someone who's tried to do any of it, yes, I would suggest so. Although, I think, the evidence for scenario planning is potentially not there yet.

Senator SHELDON: So there's looking at what's happening now and seeing what the scenarios are and what the outcomes are, and it's quite different, from industry to industry, to get that baseline?

Dr Macdonald: Yes. And it's changed in the last couple of years. Ideas about what's happening in different areas where there are levels of exposure have gone in different directions in different industries.

Senator SHELDON: What policy approach do you think should be considered in dealing with what future scenarios might look like? I appreciate you've just mentioned the foundations are still being built. What the scenarios could be based on—is that what you're saying to me? If you are, what sorts of timeframes are we looking at to get more substantial future scenarios, potentially?

Dr Macdonald: I think the scenarios are going to depend. What's going to happen is going to depend on how we deal with it. It's going to depend on training. It's going to depend on working to develop. I think there are some economists that have done some thinking around the idea of pro-worker AI that actually is AI that is used in ways that support the creation of new jobs. I suppose that's not a very helpful answer. It depends. But I think—

Senator SHELDON: The Brookings Institute has done a number of studies in this space about AI augmenting work and how we classify work and being engaged with the use of AI and technologies rolling out. So it's almost that need for a forward thinking of what we want from AI—

Dr Macdonald: Exactly

Senator SHELDON: rather than what AI is going to do to us. If we work out what we want from AI, then we can start to have a better chance of working out what it's going to do and what the scenario should be that the society either accepts, disagrees with or goes in a different direction or different route. Is that the challenge?

Dr Macdonald: Yes, I think that's the point I was trying to make. I think that, if we think about AI that is going to support the creation of good jobs through augmentation, that's the work we need to do to be able to plan—to look at where those jobs might be and how we might get there will determine our future.

Senator SHELDON: Earlier today, Alec O'Sullivan, head of operations at the Digital Economy Council of Australia, said one of major gaps is in the Privacy Act. He said:

… the Privacy Act retains qualified small business and employee records exemptions. Turnover does not reflect a business's capacity to profile individuals … and employee information can be used well beyond conventional record administration.

He's identified the Privacy Act, in his view, as an area where there are deficiencies. Have you got particular acts where you feel there may be deficiencies? I think, in your opening statement, you touched on some concerns.

Dr Macdonald: There are definitely deficiencies there, and the ongoing review of the Privacy Act over the long period doesn't seem to have addressed this—hasn't done anything about those. I think the proposal that we would put is specific workplace regulation that actually places prohibitions on data collection and data use from data in the workplace, as the platform work directive does. If there's some barrier to addressing those deficiencies in the Privacy Act, which are really critical—employees need different protections. You can't use consent in a workplace. You can't just because workers aren't able to freely consent in the same way that a consumer might be able to. That is a very significant area that needs to be tackled and—

Senator SHELDON: Dr Macdonald, I have another question. Earlier this morning, Amazon appeared and weren't able to answer questions in a particular line of questioning. I appreciate you haven't got the documents in front of you, but let me just give you the scenario that was given to them. Atlas is a system that Amazon uses worldwide to categorise workers and whether they're likely to organise around human rights issues in their workplace and having better conditions, safe conditions, better pay and the right to a voice for corporate decisions. The evidence that was handed to the Canadian courts just a month ago said that Atlas was ingesting hundreds of data points on each employee, including external data points from Reddit and other places trying to identify worker activity on seeking rights. Does that sound surprising to you or is that—

Dr Macdonald: Not at all. That's been found in other examples in the US. There has been the example of a large retail chain that classified their workers according to their propensity to unionise. I think the risks of algorithmic management—beyond the health and safety risks that are identified in our legislation—to collective organisation and to representation are really significant. That's part of the prohibitions that have been included in the platform work directive, which are about the collection of data around people's personal interactions at work, around communications at work and outside work on certain platforms and in certain circumstances. It's to prevent the use of that collection of data to enable organisations to take actions by using that to prevent workers from collectively organising.

Senator SHELDON: Ms Boland, on some of the challenges in the OHS space and regulation, I would say it's one thing having laws; it's another thing having laws, rights and obligations actually being exercised. I like the

way you described OHS as an opportunity for employers and workers but also as a guideline about how to cooperate on getting better outcomes, sometimes more productive, on safety.

A submission we received from Quinlan and Williams cited an Australian survey which said that 83 per cent of HR professionals said they were familiar with AI tools but only 20 per cent said they had a good understanding of them and 44.4 per cent of HR professionals reported that employees often did not know AI was being used in HR. Those are the professionals themselves. I'm concerned that in many of those circumstances the employees weren't aware that AI was being used with regard to them. In light of this—what can be said to be an absolute vacuum—aren't we at some sort of cliff? We're seeing AI being placed into workplaces where it's gathering more momentum. Our professionals don't know how to use it, and, most importantly, those on the receiving end of this surveillance don't know what's happening to them. Are the laws sufficient? If the laws are sufficient, then how do we not just say, 'We should all know more,' but how do we actually make sure we do know more?

Ms Boland: There are two aspects to that. From my point of view, what we can say is, yes, we've got a strong foundation, but I see four key challenges even within that foundation, some of which you've touched on. First, I think many AI systems are opaque and change, as we know, as they learn. This makes risks harder to foresee, decisions harder to interrogate and compliance harder to enforce.

Second, we have third parties developing and managing the systems used by businesses. That can obscure who's actually controlling the risk and can complicate those parallel and upstream duties. Ultimately, it makes it harder to determine what's reasonably practicable in the work health and safety space.

The third one is about that bias inaccuracy in data collection and automated decision-making that's been highlighted. This can create systemic safety risks, discrimination, surveillance, reduced autonomy and, as we've identified, threats to job security.

Fourth, regulators, health and safety representatives and entry permit holders must retain—and, as we would say, are retaining—their suitable powers and functions as technologies and working arrangements evolve, but we need to make sure that that keeps up to date.

I see the task for us is to keep the law effective as AI evolves, ensuring that businesses understand their obligations, regulators know how to enforce them and workers are able to understand and exercise their rights.

In that context, I didn't get a chance to do so before, but it's probably worth signalling to the committee that we're just about to publish our report of a substantial review, which occurred over the last 12 months, of the laws, and that will be published on Monday. We can provide a copy to the committee. One of the key recommendations coming out of that, and a priority for ministers, is that SafeWork is tasked with developing a national AI strategy. Listening to the questions to Dr Macdonald today, that's a great opportunity to pull together some of these things that don't quite fit under our work, health and safety framework but are relevant. We'll be modelling it on the Asbestos National Strategic Plan where, as you'd be aware, there are many levels of government and many legislative frameworks that touch on asbestos. I think that's a really good example of everyone working together to highlight gaps, to see who can fill them and to make sure we're all consistent in our approach.

I think there's a great opportunity to use that strategy to respond to some of those issues that Dr Macdonald and others raised and others and make sure that we've got a coordinated approach to filling any gaps that we identify.

Senator SHELDON: I know we're about to finish, but that connection between asbestos and the use of AI has a ringing endorsement from me because I think it's an apt parallel.

CHAIR: I am going to bring the session to a close. Thank you so much, Dr Macdonald and Ms Boland, for your time this afternoon and your contributions to this committee's work. Please send through to the secretariat responses to any questions on notice—including that report once it's released—and additional information within the next two weeks. That would be greatly appreciated.

ABKIEWICZ, Mr David, Acting General Manager, Digital Markets Branch, Australian Competition and Consumer Commission

BUNFIELD, Ms Megan, Chief Information, Data and Digital Officer, Digital and Technology Division, Australian Competition and Consumer Commission

RICHMOND, Ms Jayde, General Manager, National Anti-Scam Centre, Australian Competition and Consumer Commission

[16:22]

CHAIR: Welcome. I understand that information on parliamentary privilege and the protection of witnesses giving evidence to this inquiry has been provided to you. I'd now like to invite you to make an opening statement, and then we'll proceed with questions.

Mr Abkiewicz: We don't have an opening statement for today.

CHAIR: No worries. We'll get straight into it. I want to get your assessment from the outset on where you see your role in this everchanging and fast paced environment. On the impact of AI across our economy and across our society, what market and comparative dynamics do the ACCC see emerging because of competition between AI models and in the data centre developing market?

Mr Abkiewicz: The ACCC is Australia's economy-wide competition and consumer regulator. Our general role is to promote competition and protect consumers across Australia by protecting Australians from harmful products and practices, championing stronger competition in the economy and delivering benefits through right-sized regulation. We understand the benefits of AI in spurring innovation, potentially disrupting markets and bringing online high-quality new services and products to Australians, but we're also highly aware of the potential to surface, magnify or introduce new competition and consumer harms. Based on the committee's interest, I'm happy to get into any of that in more detail.

I think the question was on our views on competition, both at the AI model stage and on data centres. The thing to emphasise is that we've been observing and inquiring into digital markets for a number of years now, in particular through our digital platform services inquiry and more recently through our 2025 AI industry snapshot. Competition in AI markets continues to be a strong focus of the agency. Our work in digital markets has demonstrated that in these kinds of markets there is a propensity to tip towards strong, dominant players, and, if intervention isn't made early enough, it can be really hard to unwind that market dominance. Here I'm talking about in particular what we might call traditional digital platforms in terms of search, social media and app stores. I wouldn't say that competition in frontier models is anywhere near that point at the moment. We're still extremely early days. But it is something that we're aware needs very close and regular scrutiny to make sure that competition remains strong in these markets, particularly when we're talking about consumer-facing services. The better the competition in these kinds of services available to Australians, the better quality and potentially better price of services we'll have access to.

In terms of data centres, we're aware that there is a lot of work going on across government on this issue in terms of trying to incentivise investment while making sure that data centres are built in the interests of Australians. This is largely being done through the Office of AI in the Department of the PM&C. We're working closely with them and other relevant government agencies, particularly looking at the standards that have been recently consulted on.

CHAIR: Are you seeing any emerging risks to consumers? What emerging concerns are front of mind for you?

Mr Abkiewicz: The general proposition is that our existing laws in consumer protection apply regardless of the use of new AI tools. Where businesses are using AI, if they're making false or misleading representations the general law would still apply. That said, we're obviously very alive to emerging risks of new AI enabled services. In particular, things like chatbots used directly by consumers or used by businesses when representing to consumers are one issue we're looking closely at. Our main focus in this space is largely on the emergence of agentic AI and, in particular, the potential for AI agents to be doing shopping and transacting on behalf of consumers. That was first identified in our snapshot last year, and it's a continued focus of considering what that looks like and making sure those agents are working in the interests of consumers themselves.

CHAIR: Yes, there's been some thinking about agentic but also personalised pricing, especially in the online and the digital market space. Is that something you're alive to?

Mr Abkiewicz: Yes, definitely. Personalised pricing, or surveillance pricing as it's termed in this context, as well as agentic commerce are major parts of the government's consumer safety priorities, as I'm sure you're aware.

At ACCC we're doing ongoing work into that. As I think AISI indicated earlier in these hearings, they've been doing some work contributing to that as well. I should say that there are positives and negatives to personalised pricing, when you term it as that. But the real harm that can come about is where that pricing gets extremely granular and gives prices to individual consumers that could potentially establish their willingness to pay completely or potentially on the basis of personal data that might indicate vulnerable consumers.

CHAIR: So it's potentially used to expose or to take advantage of vulnerable consumers?

Mr Abkiewicz: That's exactly right—to take advantage of vulnerable consumers.

CHAIR: I appreciate that work is currently underway and ongoing, but are there any early indications as to where, if any, gaps exist in terms of our ways to respond to that or regulate against that?

Mr Abkiewicz: I don't think we're at the point of recommending any change. I will also note that any change would be a matter for Treasury, as the relevant policy agency.

CHAIR: Of course.

Mr Abkiewicz: I will say, though, that the recently passed unfair trading practices prohibition may be useful in combating some of that activity. That's one thing we're thinking about in that context before that new obligation comes online.

CHAIR: I appreciate that. Ms Chaney, do you have some questions?

Ms CHANEY: I'm interested in the concept of collusion when it comes to agentic AI and the extent to which there are different principles, and the ACCC has a role to play in thinking about how that works.

Mr Abkiewicz: That is certainly a focus of a lot of the recent economic literature. For quite a while, I would say, the concept of algorithmic collusion was of concern even before the recent emergence of generative AI and the huge popularity of these kinds of services. The broad point is that, as I mentioned before, existing laws still apply. If collusive activity is being carried out through a service, it doesn't mean that it's not still illegal in the Australian context. However, there are obvious issues with the ease with which you can attribute liability where this conduct is carried out through a third-party agent. That's another thing that we and other agencies around the Commonwealth are thinking about.

Ms CHANEY: If there's one entity that's creating and controlling 700 agents that are then colluding in a way that could influence markets, is that collusion under our current definition? 'Collusion', I assume, means collusion between different people or entities, but where those are agentic agents, which don't have any identity or a separate identity, is the law fit for purpose there?

Mr Abkiewicz: I don't think I can comment on whether the law is exactly fit for purpose in that circumstance, but I would say that, yes, as you mentioned, collusion is usually where there are multiple different parties who should be competitors working with each other rather than against each other. In that example, if it were the same party who owned a number of different agents, it would probably be similar to if it were a business that owned a number of different branded companies working together, which wouldn't necessarily be a competition problem.

Ms CHANEY: Although, for banded companies, we have a concept of separate legal identity when it comes to bodies corporate but not when it comes to agentic AI. If you could, confirm that this is still an unresolved issue in terms of ensuring that our laws are fit for purpose in dealing with agents and liability and identity.

Mr Abkiewicz: Yes, I can definitely confirm that there are a lot of emerging issues with agentic that require further consideration here.

Ms CHANEY: To the National Anti-Scam Centre: how many of the scams that you're seeing through that involve deepfakes or voice cloning? What does the Scams Prevention Framework require specifically in relation to AI? Is it currently fit for purpose given what we're seeing with the use of AI?

Ms Richmond: I'm happy to take that. The National Anti-Scam Centre is definitely seeing the use of AI in scams. We're not seeing it reported in significant numbers, but we are seeing things like AI voice cloning, the AI deepfake videos and AI being used, potentially, to manipulate search results. We know that artificial intelligence is enabling scammers to operate at scale, to automate and—similarly, as you were saying before—to personalise scam contact. The Scams Prevention Framework itself will provide some protection from some of these. It's technology neutral, so it's going to be more relevant as to how that scam reaches the consumer. For example, if it is through a social media platform or it surfaces through paid search, then that's likely to be covered by the framework. The framework itself doesn't cover every part of the economy, as we know. It is starting with banks, digital platforms and telecommunications companies. We understand there are gaps, but the framework has been designed in a way that enables flexibility so that, where there are gaps identified in the future, other services or sectors can be designated and brought into that framework.

Ms CHANEY: In terms of what you're saying in AI and across the ACCC in all its functions, which are the areas where you're feeling that you need to upskill the fastest? Is it in technical skills to understand what's coming down the line? Is it in international negotiations or discussions? Where are the areas that are changing the fastest?

Ms Richmond: I'm happy to start, and then I'll hand over to David, who'll talk a bit more about the work that the digital markets team undertakes. In the scam space, we certainly find that having a program like the National Anti-Scam Centre that focuses on collaboration and coordination means that we are engaging continually in terms of cybersecurity, cybercrime and emerging threats. We also have significant engagement with international organisations involved in scams. As part of that, we are continually able to get information about emerging scam threats and the use of AI. That engagement happens through a range of forums and, as I said, is both domestic and international. I'm happy to hand over to David to talk about the other work they do.

Ms Richmond: In terms of upskilling, I think it's less a skills issue in the realm of AI. We are lucky enough to be supported by a number of very expert staff in our data unit, and Megan might be able to say more about that. The main challenge is, potentially, capacity in following the extremely fast moving developments. The state of the market seems to change from one day to the next, and it is, yes, a lot to keep on top of.

Ms CHANEY: This is conjecture, but presumably that would be the same across other regulators as well. Is there any reason that the ACCC would be in a different position to other regulators in terms of that capacity challenge, with things moving so fast?

Mr Abkiewicz: No, I think it would probably be a common issue.

Ms CHANEY: Thank you.

CHAIR: Ms Mascarenhas.

Ms MASCARENHAS: This follows on from when we were talking about agentic AI polluting and what we saw with the Medicare cases. We didn't know that agents had broken in and taken data. The thing I'm trying to get my head around is: at the moment, is there the possibility for agentic AI to significantly participate in anti-competitive behaviour, does the ACCC have the tools to detect that and, if not, what kind of investment would be needed for us to be able to detect that AI? I'm concerned that it could change quite quickly and that we will probably be reactive rather than proactive.

Mr Abkiewicz: I would say there's definitely the potential for this kind of conduct but would emphasise we haven't seen any evidence of this kind of conduct in the competition realm by agents in Australia at the moment. However, we definitely agree that they're theoretically possible and, yes, very potential. I wonder if Megan wants to talk about tools or any investment that might help.

Ms Bunfield: We are currently experimenting with different technologies to detect these kinds of activities—as in, proactive detection across a range of current and potential use cases. We see huge potential for AI to assist with web scraping at scale and in a few other areas, allowing us to capture information for public websites. It is moving at a pace, and therefore there is still a lot to be done. So I would say at this point that, yes, there is a lot of work to be done, but certainly we're trying to keep in front of it.

Ms MASCARENHAS: I'm thinking about some of the ways that we've dealt with climate risks internationally. The Financial Stability Board introduced the Task Force on Climate-Related Financial Disclosures to try to create more certainty so investors could understand the climate risks. The thing that I think is fascinating with AI risk is it's much more difficult to quantify. I'm wondering whether you've got a particular view on how we ensure that markets remain competitive without collusion on a global scale. This might be out of your remit, and that's okay, but I thought I'd ask.

Mr Abkiewicz: I think the general answer to making sure markets remain competitive is continued scrutiny, which is one of our major roles at the ACCC. As you've identified, AI can accelerate these kinds of risks to competition, and that's something we are doing our best to keep a watching brief on, and we're doing the proactive work, as Megan mentioned.

Ms MASCARENHAS: Thank you.

CHAIR: Senator Sheldon.

Senator SHELDON: I have just a couple of questions. Maybe there is an obvious answer to this. Back on the question of pricing and individual pricing for people, what's the potential breach if a different price is offered to you for the same product than is offered to me? Where's the potential breach if agentic AI is used? It relies on a series of information and data inputs that tell it the price it should offer you and the price that I should be offered. Where does that differentiation in price, all other things being equal—we live as far away from the store and the delivery time is the same—go amiss of the consumer protections?

Mr Abkiewicz: There might not be any consumer harm in that scenario, which is why I mentioned that there are pros and cons to differential pricing as a general concept. It happens all the time in markets we're familiar with. Airline pricing can be quite different depending on when you book or where you book. Even discounts at the movies are a form of differential pricing for particular groups—seniors or students. In some circumstances, it can help consumers achieve cheaper prices, which is a positive. Where the harm comes into play is if and when it becomes extremely granular and targets individual people in manners which would be considered unfair or potentially abusive. I should clarify this isn't a breach of current laws. This is just a potential consumer harm that is talked about when we're considering the potential for extremely—

Senator SHELDON: So, at this point, unfair or abusive is not a breach in the case of an AI making a decision? I've got grey hair, so they've decided to charge me a higher rate because I've got grey hair, and you haven't—congratulations!

Mr Abkiewicz: Under current laws, that example of giving a discount or giving certain groups of consumers different prices isn't against the law.

Senator SHELDON: I'm trying to equate it back to a pricing mechanism. In the case of the US, there are a number of labour hire companies that operate, providing people to work as nurses in hospitals, and they do an analysis on their credit cards. If they have a high-cost credit card, they are contracted at a lower rate of pay because they're more desperate to get jobs. If someone were to actually equate a price to me for a particular item and that price was based on a discriminatory matter because I had a low income or a high income and based on what I'm more likely to actually purchase as an individual, is that at the moment not a breach of the consumer legislation?

Mr Abkiewicz: No.

Senator SHELDON: They'd have the capacity to do that on an individual basis? I could be singled out on multiple products that I use every day under the existing laws?

Mr Abkiewicz: That is definitely a concern.

Senator SHELDON: That's an issue about the societal view on what I would consider as unethical. I'm not asking for your opinion, but I'd consider it as an unethical behaviour.

Mr Abkiewicz: That's correct.

Senator SHELDON: Are there any examples around the world where there's either proposed legislation or actual legislation that you're aware of that's been put in place or considered to be put in place to deal with these challenges?

Mr Abkiewicz: Yes. This has been a focus in the United States recently. I understand several states have proposed or potentially even passed laws that make certain forms of surveillance pricing illegal in those jurisdictions. The US Federal Trade Commission, which is the ACCC equivalent, has recently issued guidance that they see doing this kind of discriminatory pricing or surveillance pricing without disclosing it to consumers as illegal under their settings.

Senator SHELDON: I'm going to another track, back to another part of the conversation. It's a parallel question to the other conversations we've had on this already. With the existing laws, if there were a breach of activity by an agentic AI, who in the technology stack would be held accountable for that and what player would be held accountable for that? Is that clear yet, or it's not clear?

Mr Abkiewicz: That's a very perceptive question. I would say it's not clear yet. It would depend on the circumstances, and it's an emerging and untested area.

Senator SHELDON: If we had a series of agents working together from multiple sources, which were partially, as I understand, but maybe I'm incorrect, from OpenAI—multiple agents were used from different programs.

Mr Abkiewicz: If those agents were colluding to set prices or conduct other illegal activity on behalf of the people who are running those agents, that would be a problem. It would come down to the circumstances.

Senator SHELDON: Are you aware of considerations that have been given in Australia or elsewhere—I quite like the occupational health and safety space for lots of reasons. Part of it is because there's supply chain accountability, multiple accountabilities, and in a sense there's a sliding scale based on the actuality of what's going on. However, in the factory, if somebody gets injured—if the guard is taken off because the worker is told to work harder and faster and they take the guard off to get the work done more quickly and they lose their arm—there is a high responsibility on the employer. There's some abrogated responsibility for the worker depending on a whole series of factors, but the employer is definitely accountable. In the question of what the penalty is, there's

absolute responsibility on the manufacturer and the owner. It's a sliding scale depending on accountability—did they train them, did they skill them, did they tell them not to take it off, were they properly supervised and should they have actually forced them to work to a production level that was unrealistic?—but there's accountability nonetheless. Are there any suggestions about there being an automatic fault responsibility on anyone with regard to AI? At the moment, I can't see how anything other than that works.

Mr Abkiewicz: I think we would definitely agree with the general principle that there should be accountability regardless of the tool used, but the thorny issue of attributing legal liability, as I mentioned, is still an emerging area and would have to be tested on a case-by-case basis under current settings.

Senator SHELDON: Thanks.

CHAIR: That brings us to the end of this session. Thank you so very much for making yourselves available to come before us and contribute to this work. If there's anything further that you committed to providing or are keen to provide in terms of additional information to the committee, please do so via the secretariat. The committee has agreed that responses to questions on notice should be provided in two weeks. That now concludes today's hearing. Thank you to you and every witness that appeared before us today and to Hansard and Broadcasting for their ongoing assistance.

Committee adjourned at 16:51