



LOCAL GOVERNMENT, SMALL BUSINESS AND CUSTOMER SERVICE COMMITTEE

Members present:

Mr JP Lister MP—Chair
Mr SM Dillon MP
Mr MA Boothman MP
Mr MPT Healy MP
Mrs ME Nightingale MP
Ms JE Pease MP

Staff present:

Dr A Cavill—Committee Secretary
Mr Z Dadic—Assistant Committee Secretary

PUBLIC BRIEFING—CONSIDERATION OF AUDITOR-GENERAL REPORT 2: 2025-26— MANAGING THE ETHICAL RISKS OF ARTIFICIAL INTELLIGENCE

TRANSCRIPT OF PROCEEDINGS

Wednesday, 29 October 2025

Brisbane

WEDNESDAY, 29 OCTOBER 2025

The committee met at 9.01 am.

CHAIR: Good morning. I declare open this public hearing with the Queensland Audit Office on its *Report 2: 2025-26—Managing the ethical risks of artificial intelligence*. My name is James Lister. I am chair of this committee and the member for Southern Downs. With me on the committee today are: Margie Nightingale, the deputy chair and member for Inala; Mr Sean Dillon, member for Gregory, who is substituting for the member for Townsville, Adam Baillie; Mark Boothman, member for Theodore; Michael Healy, member for Cairns; and Joan Pease, member for Lytton.

This briefing is a proceeding of the Queensland parliament and is subject to the parliament's standing rules and orders. Only the committee and invited witnesses may participate in the proceedings. Witnesses are not required to give evidence under oath or affirmation, but I remind witnesses that intentionally misleading the committee is a serious offence. I also remind members of the public that they may be excluded from the briefing at the discretion of the committee.

I remind committee members that the Auditor-General and officers are here to brief us on findings and observations contained in the report. Any questions seeking an opinion about policy should be directed to the minister or left to debate on the floor of the House. These proceedings are being recorded and broadcast live on the parliament's website, and I welcome those who are joining us online. Media may be present and are subject to the committee's media rules and my direction at all times. You may be filmed or photographed during the proceedings, and images may also appear on the parliament's website or social media pages. Please turn your mobile phones off or onto silent mode for the duration of the hearing.

BROWN, Mr Darren, Assistant Auditor-General, Queensland Audit Office

GUERRERO, Mr Tony, Director, Queensland Audit Office

VAGG, Ms Rachel, Auditor-General, Queensland Audit Office

CHAIR: I now welcome representatives from the Queensland Audit Office. I invite you to make a brief opening statement after which the committee will have some questions for you.

Ms Vagg: Good morning. I would like to begin by acknowledging the traditional custodians of the land throughout Queensland, including the Turrbal and Yagara people who are the traditional custodians of the land on which we meet today. Thank you for the opportunity to brief the committee on my report, *Managing ethical risks of artificial intelligence*. Artificial intelligence offers opportunities for government to improve efficiency, decision-making and its services. AI systems can simulate aspects of human intelligence, enabling them to solve problems and perform functions that have traditionally relied on human involvement or judgement. AI can take a number of forms, including generative AI, natural language processing, computer vision and machine learning. The Queensland government is using each of these types of AI.

This use of new technology requires the effective management of ethical risks, ensuring that the public sector uses AI in a way that aligns with public sector values and our community's expectations. Each type of AI, and the context in which it is used, will impact on the type of risk and the management activity undertaken to manage that risk. There are eight nationally recognised ethical principles that apply to AI which guide the identification of ethical risks: human, societal and environmental wellbeing; human rights and values; fairness; privacy; reliability and safety; transparency and explainability; contestability; and accountability.

This audit examined whether the Queensland public sector has policies and guidelines in place to effectively identify and manage the ethical risks associated with AI systems. We audited the whole-of-government policies and guidelines, as well as how they have been applied to two AI systems used by an agency. The Department of Customer Services, Open Data and Small and Family Business, known as CDSB, is responsible for issuing policies to guide entities in identifying

and managing ethical risks associated with AI. CDSB developed an AI governance policy and supporting tools, and they are appropriate to guide the public sector's ethical use of AI. The policy requires departments and statutory bodies to conduct ethical risk assessments for AI applications they use.

The Foundational AI Risk Assessment framework—FAIRA—developed by CDSB, provides entities with a structured way to assess risks. CDSB could give clearer advice on when and how often to apply ethical risk assessments across the life cycle of AI systems. The FAIRA framework could also be adapted to better support assessment of ethical risks for lower-risk systems, and this will support a more efficient assessment in the public sector.

CDSB advises government on IT-related risk, including those associated with AI. However, it has limited visibility on how AI is used across government and the risks that may arise through its use. By improving its understanding of AI use and related agency identified risks, CDSB can better target its support to build capability in managing ethical risks and improving outcomes across government.

We also assessed how the Department of Transport and Main Roads, which we refer to as TMR, managed ethical risk in its use of AI. While CDSB has a whole-of-government role, individual entities, like TMR, are responsible for identifying and managing ethical risk within the systems that they operate. TMR manages ethical risks of two of the AI systems it uses—QChat and mobile phone and seatbelt technology.

QChat is a generative AI virtual assistant created by CDSB for Queensland government employees. This AI system can draft communications, review texts, summarise information and generate ideas. The mobile phone and seatbelt technology program is an example of machine learning, and this AI system uses image recognition technology to detect possible mobile phone and seatbelt offences.

TMR has not yet established departmental-wide AI policies or governance arrangements to consistently oversee ethical risks for its AI systems. It does not have full visibility of the AI systems in use or under development across the department and, therefore, could not effectively monitor where AI is used, its purpose, or the risks associated with each system. TMR is aware that its governance and assurance mechanisms may need updating, and it had taken initial steps and was developing an AI strategic road map to identify and address these gaps. TMR had assessed some ethical risks for QChat, but was yet to complete a full ethical risk assessment or establish governance arrangements for the system. While CDSB manages the system-level safeguards within QChat, which helped detect and reduce inappropriate use, TMR needs to improve its monitoring of the system and staff training to support its responsible use.

Similarly, for the mobile phone and seatbelt technology image recognition AI, we found TMR is identifying and managing some aspects of ethical risk. It had completed privacy impact assessments and implemented risk mitigation strategies to support system reliability and accuracy and to protect privacy, but it had not completed a full ethical risk assessment or established adequate governance arrangements over this system.

Both the mobile phone and seatbelt technology and QChat were implemented before the introduction of the AI governance policy. However, at the time of the audit, the AI governance policy had been in place for 12 months without the department conducting full risk assessments on these two systems.

I made seven recommendations in the report and also compiled a checklist for managing ethical risks of AI. This is the first report in my series of AI audits. Looking ahead, my forward work program includes a future audit that will assess how effectively the public sector is realising the benefits of artificial intelligence. This committee may also wish to hear from agencies themselves on their progress in implementing our recommendations, and I am happy to take any questions the committee has on this report. Thank you.

Ms PEASE: The report makes very interesting reading. Thank you. In the audit, you refer at length to the ethical risks of AI and in particular when it comes to the use of systems that support service delivery or decisions that affect the public. Do you consider that the use of AI or related systems to process applications for grants or funding as something that would require the highest level of risk assessment?

Ms Vagg: All systems require risk assessment, and particularly AI systems require a risk assessment that includes ethical risk considerations. Yes, I think any system that is used, in particular AI systems, should have a risk assessment and then appropriate responses put in place for those particular risks. That might include monitoring of activity, monitoring of outputs, training for staff—those types of activities.

Ms PEASE: Should the use of AI or AI related systems in hardship assistance grants for those affected by natural disasters require foundational artificial intelligence risk assessment?

Ms Vagg: I did not review that particular type of AI where it is being used there. AI is being used across government, and our recommendation, and that of CDSB, is that the foundational risk assessment is performed for each system in use and then appropriate responses put in place.

Ms PEASE: Is there any consideration that you would review those particular areas, particularly when it is hardship grants and things like that, but it is performing as expected?

Ms Vagg: It is something that I could do. Ultimately, it is the responsibility of each director-general to make sure that they appropriately assess risk in their organisation and then have appropriate responses to it, and then there is a role for CDSB to have a look across government to see that each director-general is appropriately doing those things and that they are collecting information about the risks that have been identified and then how they are being managed across government. I am the ultimate final checking point. It would be better to actually have that within the agencies, and any support mechanisms that a director-general might have in place for checking on activities within that department, that would be a better place to start. In terms of my forward work plan and what we look at when it comes to technology, systems controls and artificial intelligence, I am always more than happy to take suggestions on that as well, if you think this is a valuable topic for us to look at.

Mr BOOTHMAN: You made a comment about QChat and dealing with inappropriate risks. Can you elaborate on what you deem as inappropriate risks? I will have a follow-up question after that.

Ms Vagg: I will give you a bit of framing on QChat. It is a generative AI tool which was developed by CDSB and is available for public sector employees to use. CDSB has some monitoring tools and some guardrails and controls in place. Then each agency, as they are using that particular tool, should also monitor its usage, should train their staff, and have their own controls in place. Therefore, if they do not make their own assessments, they do not know about appropriate use or not of that particular system. It is more about each agency that is using that system making sure they have their own monitoring tools, training of staff and other appropriate internal controls.

Mr BOOTHMAN: With the actual system itself, obviously it has been used for a little while now. How has it impacted the customer service outcome for Queenslanders? What has it done? Is it creating efficiencies? Is it streamlining? Are there any concerns about potential incorrect information? Could you elaborate on that?

Ms Vagg: I do not know the answer to that yet. It is still a fairly new system. It is being increasingly used by the public sector. The recommendations in this report are about making sure that the controls are in place in terms of appropriate use. In terms of gaining efficiencies and outcomes for the public sector, that would be more in my next report which is actually 'Realising benefits of artificial intelligence' and I have left that this year—it will be in later years in the program—to give the public sector an opportunity to more fulsomely use AI, have it implemented, have suitable controls and do their own assessment about how effective the use of AI has been for them.

Mr BOOTHMAN: When it comes to the appropriate use of the actual systems, do you have concerns that it might create too much reliance for public servants on these systems? Do you feel that, as a run of the mill, people will adjust, learn and realise that it may be a great tool, but you still need to use your own initiative so to speak?

Ms Vagg: The use of any systems—AI or not, so technology—is a way for a government to improve processes, efficiencies and the effectiveness of what they do. This would be another tool to improve the operations of government. Making sure there are suitable controls or actions to manage risks is really important and that is what this report is all about. Looking at those eight principles of ethical use and making sure that, if one of those is at risk, there are appropriate controls within the agencies to do that. One of those things is keeping humans in the loop, so public servants in the loop, about what the system is doing and the checks and challenges about that particular system. I guess it is adapting the skill set of public servants as well. Sometimes it is about doing it themselves to then reliance on a system and checking the suitability of the outputs of that particular system. That is not just with AI; that is with any implementation of technology.

Mrs NIGHTINGALE: A big nod to whoever came up with FAIRA. That is quite clever. Firstly, we have talked about systems that are exclusively AI. My question is around systems that largely integrate AI capabilities but are not exclusively AI. Should they be included in the ethical risk assessment?

Ms Vagg: I will turn to Tony.

Mr Guerrero: If they contain a component of AI, they would perform the ethical risk assessment. They might consider the decisions of systems that are downstream that are using AI for decisions and actions, and the impacts and the controls that they may need to apply there as well.

Mrs NIGHTINGALE: That component of it will still have that ethical lens over it?

Mr Guerrero: That is correct.

CHAIR: Can I get your view on how AI could possibly be adopted by osmosis in that it may be taken up over time imperceptibly because it accompanies other applications or is just an accepted way of doing business and is, therefore, not necessarily recognised as having risks that need to be managed or that could escape the scrutiny that we accept needs to be employed?

Ms Vagg: That aligns with the recommendations in the report. From a whole-of-government perspective, CDSB should have awareness of where AI is being used. That is not just systems that are AI specific but also other pieces of technology that have AI components in them. That is definitely part of the recommendation to CDSB. Then at the agency level, there should be an inventory or an awareness within that particular agency or each agency about how AI is being used and then appropriate assessments put over that usage. Awareness is a very important part of the recommendations in this report.

Mr HEALY: The real concern is that, when you talk to people in the real world about AI, there is a significant amount of suspicion about AI. When we look at technology and its application in everyday life, I think there is a suspicion that we lose our gut feeling. I understand the intricacies of government policy and I understand the enormity of the job and measuring. Are we losing our gut feeling in this space with increased technology? Are we losing our connection? A lot of people out there are not aware or familiar. I accept that we live in a world where we are swamped with information and starved of intelligence. How do we address those sorts of concerns? You would be aware that people have those issues.

Ms Vagg: While it was not a specific topic in this report—

Mr HEALY: I noted that.

Ms Vagg:—as this was about a very particular part of it, there are recommendations in there about appropriate training and guidance of staff. Really, that is a part of it. There is education about ethical use of AI, but there is also overall capability assessment of staff in terms of what they need to do in their organisation. We need appropriate learning and development and growth activities for staff for the role that they need to undertake in the world where technology is used. I think that is for any organisation to assess and then work through changed requirements for the development of staff.

Mr HEALY: Auditor-General, if someone is sitting at a workstation, can they can make their own assessment or will they be directed or encouraged? At what stage does the implementation become an essential component in any due process of analysis or does it become the decision of the level of government that that person might be working at? Does that make sense?

Ms Vagg: That would be up to each director-general or agency head about which systems they use, which technology they use and then appropriate staff and skills to respond and manage that use of technology. An important control in the use of systems and technologies is keeping humans in the loop, particularly looking at the inputs into the technology and then the outputs from the technology. Having appropriate checks on either side of that is really important and training staff to be able to do that is important. That is for each agency but it is also as we bring new staff members in to train them coming out of the university systems.

Mr HEALY: In your audit you say that the CDSB has limited visibility on AI use and the ethical risks associated. Their response to recommendation 7 outlined that they would address it for CDSB projects. Should this recommendation be applied more broadly to capture all projects?

Ms Vagg: The intention of the report is that there is an awareness across government of AI use so it should be at the individual agency level. They should be aware of where AI is being used in their organisation and make risk assessments at that level. Stepping back from a whole-of-government perspective, there should be awareness of which systems are being used, what those ethical risk assessments look like and then if a whole-of-government response is further required. Beyond what CDSB is currently doing, if they need to do more then that information will help guide that additional work that they may need to do.

Mr BOOTHMAN: When it came to choosing the department you chose TMR. Obviously, transport legislation is very black and white. There is very little grey area so to speak, so it would make it easier for an AI system to interpret a request and to then put that forward. I know you would have done only that department. Are there any thoughts, as an audit office, that you would do other areas where AI had to be more creative to find information? Is that a concern? I know you have put forward your seven recommendations. I am trying to get my head around this: you have something that is very black and white but potentially, in other departments, you would have a lot more grey areas. I am curious to hear your thoughts about that.

Ms Vagg: There is complexity across government and there is complexity within TMR as well. I will get Darren to speak in a minute about why we have chosen TMR. When we do a series of audits, typically we select different agencies each time and we choose an agency that best supports that particular subject matter. In our future reports on AI, you are likely to see other agencies being selected. We are also grappling with the use of AI in our financial audits. We audit financial statements. That is another area where we are likely to see where AI is being used from a financial processing perspective and a financial reporting perspective and will, therefore, audit it at that level across every entity. That will be on that part of our audit. Darren, in terms of sample selection, did you want to add a bit more to that?

Mr Brown: In terms of this particular audit, we chose TMR on the basis that the AI policy had been out for about 12 months and, bearing in mind that there is not a complete registry of what AI is in use across all of the public sector, we looked across the public sector at where we could see AI was in use. Across the public sector it is at various stages of either development, testing or implementation and being used. We found that TMR was probably ideal in the sense that it had been using QChat for a period of time and a sufficient period had been progressed since the AI policy was implemented so we would expect that that AI policy would have been applied to those particular systems. Because MPST had been out for a period prior to the AI policy as well as QChat, we identified that that was probably the best agency for this particular audit. Also, given that TMR was using QChat, which is being more broadly used across the public sector, the value that we could provide to the public sector would be greater.

Mr BOOTHMAN: For your audit process, do you use AI systems?

Ms Vagg: No, not yet. We use technology greatly. We use data analysis and automated data analysis within our audit processes. We are also very conscious of the protected environment that I have so we are carefully stepping forward in terms of the development and an appropriate strategy for protected information in the audit process. I am working with other auditors-general across Australia for us to share what we are doing and some joint strategies and risk management from that side.

Where we are progressing as a business is more in the unprotected space so more generally available information and how we can best use AI to analyse that information for us. An example would be, as we are assessing activities of the public sector to develop the forward work plan of QAO, AI can be very helpful in terms of synthesising and analysing huge pieces of information for us to help inform the forward work plan. As a business, we are going forward with unprotected information and using it now. With the protected information part of our audit activities, we are stepping forward in a more careful manner to make sure we protect information appropriately.

Mrs NIGHTINGALE: On page 12 of the report, you refer to entity feedback on FAIRA, which is one of the supporting materials provided by the CDSB. Which entities provided feedback on the supporting materials and could you please elaborate on the feedback that was provided?

Ms Vagg: I will turn to Tony for that.

Mr Guerrero: Departments were involved in that feedback process. I do not have the specific names of those at hand, but all departments were included in it. The feedback pointed out that the tool was a good communication tool for stakeholders. There were some concerns around the tool being a little bit complex for more simple AI systems, so they did request a solution that would be easier for them to navigate while still managing those ethical risks.

CHAIR: Member for Theodore, you seem to have had some insights here and some questions. Would you like to ask another one? He must be younger than me; he understands this better.

Mr BOOTHMAN: A little bit more savvy, I suppose, in that respect. With regard to the answer you gave me before about how audit offices around Australia are collaborating with each other, is there a set process or a set governance model, so to speak, to ensure that your organisation uses the AI in what is deemed a manner that ensures that the information is properly filtered, correct and also—obviously you are dealing with very secure information here—properly secured? Can you talk about how this works and how we can ensure that there is a high degree of trust in what you are doing?

Ms Vagg: I can. All of these principles that apply to public sector entities also apply to me and are based on national principles and international standards on governance of AI, so those same principles from a governance perspective apply. Do I have an AI governance policy which is based on the international standards and then do I have an AI strategy—that is, identification about what I am going to use and when and what expected benefit I will have from that AI and then what controls I have in place? We are progressing through at the same pace as other organisations and other audit offices are a good benchmark for me to see whether we are progressing at the same pace. My earlier comments about the information that I hold and then how I am using technology to assess that information indicates that we are progressing through that. I will do so carefully.

In the non-protected space—where there is publicly available information—I can use AI to assist with our assessment of that information and our summarising of that information, but we will always have our humans and our teams checking that information. I think at the local level that happens. I also check it via my extensive stakeholder relationship management and I talk to individuals out in the public sector which also helps with that checking function. When it comes to protected information—so that is all of the information which is within government that I have access to and it should not be disclosed because it is confidential—it requires more careful consideration about what those controls look like to make sure that I still protect that information but then also how I assess that using technology, so we are at the early stages of that. I apply a series of auditing standards and controls which make sure I still comply with all of the standards as required by the law as well, so they are all still in place.

Where we have found the most benefit so far with the use of technology is in the collection of data across government and in our processing of that data and applying some automated procedures in there, and that is automating some audit procedures that a human otherwise would have done. We check the outcomes of that. We undertake testing based on that, so there are definitely checks and balances in place, and that is not the use of AI; that is just the use of technology within our audits. I have a technology savvy organisation. We have definitely been advanced in terms of the data-processing side of QAO. We have well-established controls in place with executive oversight and controls as well. Is that helpful? I am happy to take further questions.

Mr BOOTHMAN: Sort of. I will allow my colleagues to ask some more questions and then I might get back to you.

Ms PEASE: It is a very big undertaking—huge. With FAIRA I understand it was an internally created platform by CDSB; is that correct?

Ms Vagg: CDSB but I think it is based on some standards. Tony, I might just get you to expand there.

Mr Guerrero: Yes, so it is available on a public website. It is a tool that helps step through understanding what AI does, making sure you understand the outputs and decisions that flow on.

Ms PEASE: So FAIRA is a publicly available platform?

Mr Guerrero: Yes.

Ms PEASE: Okay. You have had some criticism in your report with regard to the oversight for the department. Do you know if the department is adequately resourced given that it is using a public platform process that is overseeing all of government policies around AI use? Is the department adequately resourced to undertake that?

Ms Vagg: I have not made a comment on that in the report. It is probably a question which is better asked of the director-general. They define what their responsibilities are and what they should be doing and we make recommendations for improvement there. In terms of the resourcing of the department, it is probably a question for the director-general.

Ms PEASE: Just further to that, do you know what they are doing and using in other jurisdictions in this space? Can you provide some advice on that?

Ms Vagg: We do have some jurisdictional analysis in the report.

Mr Guerrero: Yes, that is correct. It is in the appendix, so there are different approaches that different jurisdictions are taking. Some are mandating a central oversight body and others are mandating principles, but we have a bit of a jurisdictional comparison in the appendix of the report on that.

Ms PEASE: Thank you.

Mr DILLON: In listening this morning there has been consistent reference to both the protected dataset and the non-protected dataset, so obviously confidential information vis-a-vis something that is relatively potentially open source. We see reliance on AI becoming everything from ChatGPT to various forms of subtle AI that we are not even perhaps aware of as consumers of technology that is occurring and there is becoming an in-built reliance on the accuracy or the validity of that data without necessarily knowing the source of what AI is drawing the dataset from. Was there any observation that you made or detected when you were doing this audit about the AI tools and their combination of datasets, both protected and not, and any affectation that that may have had on either the validity or perhaps the department's reliance upon the data? Maybe that feeds to the ethical question that you have raised; I am not sure. My question has become very convoluted, but have you from an audit perspective been able to ascertain any concerns around reliance on dual datasets here given that you raised the protected dataset?

Ms Vagg: To answer that, when I was talking about those types of classifications I was talking about QAO in our activities that we are undertaking. Government has its structure which we comply with about how they classify information and the requirement to protect or not with information, so that is established within government. In terms of ethical risk assessments, one thing that is suggested is that reliability of systems is one of the ethical risks to be assessed and then appropriate controls put in place. That would be something we are saying each agency should do and in fact is something that CDSB also requires with their ethical risk assessment through the FAIRA framework. In terms of what we have identified in our assessment, we were looking at the application of those ethical risk questions and whether the department has appropriately responded to them. We did not see anything that we have not identified in this report of concern, so it was more that those assessments should occur and should occur in a fulsome way within the departments. Did you want to add anything to that, Tony?

Mr Guerrero: Yes. QChat did have some controls to help with preventing the uploading of protected information, although there are still ways in which employees might inadvertently provide protected information into it through prompts, and that is why it is particularly important that there is training for staff around what information should go in and which should not, so clear policies on that really help to prevent those risks occurring.

Mrs NIGHTINGALE: Given that this is all so complex and so new and quite challenging, is there a state or territory that you have identified as offering a best practice approach to the governance of AI in the public sector?

Ms Vagg: Queensland has some examples of best practice in terms of application of international standards about the governance of AI. Each jurisdiction has different approaches. Each has benefits and issues associated with theirs, so it is really about identifying those elements which are most appropriate for the state of play, so there are pros and cons against each jurisdiction.

CHAIR: I wanted to ask whether the framework, controls or risk assessments associated with AI should trigger a deeper look at the systems and processes and policy that they work on. I recall when the seatbelt camera technology was first rolled out, I had quite a few constituents—and I would not be surprised if there were others on the committee here who had the same experience—who came to me saying, 'I got this infringement notice saying I wasn't wearing a seatbelt, but I keep a letter in the glove box from my doctor which for the last 38 years has been satisfactory for a police officer who's pulled me over to say, "You're not wearing your seatbelt" saying that I have a medical reason not to do so.' It appears that contemporaneously that recognition was no longer acceptable and a new form had been created that had to be filled out, certified and filed neatly in quadruplicate somewhere in order for that to be recognised, so the individual was not only impersonally picked out by AI for not wearing a seatbelt but also did not have the opportunity to repace it in the way they had already been able to in the first place. I suppose that is a bit convoluted, but do you see my point about whether the advent of AI could be cause to trigger a deeper look at the operations and the policies and procedures of the agency that is employing it and should that be therefore considered in their risk assessments when they introduce AI?

Ms Vagg: When processes change, those types of assessments should occur for users of the processes and people who are affected by them, and that is regardless of technology and AI, so that should occur. In terms of the ethical risk considerations, having the ability to be accountable and to challenge the outcomes of a particular process is part of that risk assessment process, so that is at a higher level response. In terms of those details about how it works with this particular system, I think you would need to ask the department about that.

CHAIR: Speaking globally.

Ms Vagg: Yes. That ability to check and challenge a system and to be transparent and accountable is an important part of the ethical risk assessment.

CHAIR: Okay. I am sure all of us have a few little ideas popping up in our heads about examples that we have encountered along the way. The more I think about it, the more I think of. Are there any questions from the opposition?

Mr HEALY: It is probably more of a statement than anything. Because we are seeing such rapid growth in the technology in this space, I am just wondering if you have identified any areas that need to be assessed a little more frequently than others because of the implementation of new technology and the impact that has on AI.

Ms Vagg: One of the recommendations in the report is that CDSB give more guidance about how regularly ethical risk assessments should occur and at what stages of an AI system's lifecycle—sorry, life cycle it should occur, so not just at the beginning but as things change.

Mr HEALY: You see, we are already talking about lifestyles.

Ms Vagg: Yes, that is right—life cycle.

CHAIR: We will be talking about Cairns in a minute.

Mr HEALY: Yes, that is true. That is a great lifestyle, Chair. Thank you for that segue.

Ms Vagg: So it is a recommendation that CDSB give greater guidance to the sector about regularity of assessments.

Mr HEALY: When there is so much money being invested in the private sector and it is such a fast-moving area, it would have to be a regular occurrence. The updates would need to be fairly frequent.

Ms Vagg: The assessments and any appropriate adjustments to the risk responses.

Mr BOOTHMAN: My question goes to machine code. With QChat there would be a shell code which the department would then use. They would put in their own instructions, as per the requirements. With your investigations, did you have concerns that AI software is designed to replicate tasks to make it easier for itself to look for subject matter, to obtain more easily and more efficiently? Your recommendations are wonderful and they would pull it in a lot but, when it comes to how AI will continue to develop itself, do you think that the departments themselves will become too reliant on this information? Whilst there are recommendations about how we put this software together, do there need to be further investigations as to how public servants actually use the software to ensure that there is transparency and that the correct information is put out there to the public because, like humans, AI can learn from bad habits, so to speak?

Ms Vagg: That is definitely part of the risk assessments. It is also the part of any judgement a director-general would apply when deciding to use a piece of technology. They should assess capability of their staff to use that technology and to use the outputs of that technology. In terms of the code itself within the AI systems, I did not look at that. That was not part of this audit. Our audit was actually about the risk assessments within the state and understanding of AI usage. It is a period of great change when it comes to use of technology and artificial intelligence. Therefore, adapting our workforce—that is, with any organisation to be able to appropriately use that technology and then to appropriately deliver the services that they need to deliver for that department—is an important part of the management of those agencies right now. It is also part of our training and development of graduates who are leaving university, that they are appropriately skilled and trained and that then we adapt that knowledge and skill into our agencies. It is a period of great change but also great benefits.

Mr BOOTHMAN: There are, yes.

Ms PEASE: You mentioned on page 38 of your report that Queensland 'stands out' because it mandates the use of ISO 38507, which relates to governance implications of the use of artificial intelligence by organisations. Can you please elaborate on that and what that means?

Ms Vagg: I can. I will get Tony to give more detail. There are international standards, ISO standards, that apply to many practices within organisations. They are good in terms of standardising approaches and then reliability for an external person looking into that organisation if they know that international standards have been applied. They are also usually a reflection of better practice if an international standard, an ISO standard, is applied. That is why Queensland has been called out in terms of the use of that particular standard. Tony, I might get you to add to that.

Mr Guerrero: That standard is quite useful for agencies to reflect on their current governance arrangements and establish new policies. It helps guide them on establishing accountability, responsibility and also risk management practices within it. It is quite a useful tool.

Ms PEASE: I would like to go back to our earlier discussions with regard to the ethical use of artificial intelligence. How does that impact the assessment for things that are not black and white? We have talked about TMR where you are either wearing a seatbelt or you are not. In terms of hardship or the impact of losing a house, how is that being overseen? There is no black-and-white answer to that; it is more of a feeling potentially or a hands-on approach. How are the ethics of the programs that are being used being determined, and is there oversight currently?

Ms Vagg: I could not give the oversight of that particular process that you are talking about. That would be a question back to that particular agency. In terms of the ethical principles—human, societal and environmental wellbeing; human centred value; fairness; privacy—all of those aspects of that ethical risk assessment should go to those judgement pieces that you are talking about. It is understanding how the system as well as the processes of the department around that system actually assess whether those ethical principles are being applied appropriately. That is over time as well as the system might change and develop in itself. It is about continued assessment to make sure that those ethical principles are being applied. That would be the same whether a system is making that decision or whether people are making that decision. You would like those same principles to be applied and that the department or the agency is making assessment of the appropriate application of those principles. It applies regardless of whether it is a system or people.

Ms PEASE: I am intrigued as to how that can be assessed. Does the AI system assess itself or is there someone out there assessing the ethical behaviour of AI?

Ms Vagg: It would be people in the agencies at the moment. They would look at inputs into the AI system and then they look at outputs of the AI system. They look at the judgments and activities that have been undertaken by that system and whether the outputs are appropriate with those ethical risk considerations, so it should be a structured process. It should be something that is articulated within the department and important within the department as they are implementing new technologies.

Ms PEASE: In any of your audits, given that you have done some audits in different departments, have you seen that that is actually happening, that they have those processes in place?

Ms Vagg: I might get Tony to speak to the two systems that we looked at.

Mr Guerrero: Yes, certainly with QChat there has been testing on inputs and outputs. When they configure it, sometimes AI can behave unpredictably, and it will get updated as the models will change. So when those processes happen, they do need to retest those to make sure that it is acting the way that it is intended to. With MPST, where it is identifying the potential infringements, there are weekly audits that are happening by the department to crosscheck whether the AI is making those decisions accurately. Then that information is then fed back into the model to help train it to improve its accuracy going forward.

Ms PEASE: With regard to the artificial intelligence that these departments are using, is it the same company that is providing it or are they just drawing it from everywhere? Where are they getting their artificial intelligence from?

Ms Vagg: It is from various places. CDSB has developed QChat themselves. Sometimes, as mentioned earlier, existing software may have some artificial intelligence added to it which the department should be aware of and make a conscious decision to incorporate that technology, or they may choose to purchase technology or develop technology with an external vendor.

Ms PEASE: Is there any oversight by the department as to which software developer they are purchasing from or engaging with in terms of the efficacy associated with that organisation?

Ms Vagg: It is similar to the purchase of any technology, product or system. You would want to do assessments as to the appropriateness of that particular product, the controls of that particular product—and there are whole-of-government suggested controls about where data might be stored, for example, and how data is used. As part of that, any AI assessment should also have the consideration of whether it is going to be appropriate for the process that it is going to be used for.

The recommendation in the report is that there is a greater understanding within each agency of AI and where it is being used in that agency and then from a whole-of-government perspective about where AI is being used and where the risks might be, which is why it is an important recommendation within the report to get a good understanding across government.

Ms PEASE: I have a million questions obviously, but what I am interested in is, ethically, with the data that TMR, for example, currently have and they have available, are there any controls to ensure that there is no unintended use of that information when accessed by other departments? Currently information sharing on various things requires some fairly significant background work to happen. With AI, given that it is so powerful and stores so much information, is there any risk associated with it and what protections are in place to make sure that there is not an unintended consequence of sharing that information? I am not talking about the dark web or someone hacking in; it is already here and available. Can it automatically be available to everyone, whereas once upon a time there had to be this big process about sharing that information?

Ms Vagg: There are different risks depending on the type of technology and AI that is used, but data protection is an important part of any technology used within government. As AI systems are developed or new technology is implemented, understanding the data that will be within that technology and then the protection of that information is an important part of the implementation of any system including AI. In terms of the mobile phone and seatbelt technology, MPST, the department has done a privacy impact assessment. That is about protecting that personal information that is contained within that particular system. An assessment has been done. Therefore, the department should have identified any additional controls and mechanisms that need to be in place for that particular system to protect that personal information. We know that that assessment has occurred. Any further information about where the information is held and how it is protected is probably a question back to the department on that one.

Mr BOOTHMAN: Just going back to data protection for an AI system, as the member for Lytton was talking about, AI systems would relate back to a MySQL table system, which would be then governed by the protocols and the requirements for the department to protect their data; is that correct?

Ms Vagg: Yes

Mr BOOTHMAN: When it comes to seatbelt software, it would be an interpretation from an AI modelling system which would actually have predetermined what would be classed as wearing a seatbelt or using a mobile phone. Therefore, the AI software would be very limited with what it can do, apart from communicate with the database to say that this is a licensed numberplate; is that correct?

Mr Guerrero: That is correct.

CHAIR: On behalf of the committee, I would like to thank you very much for your appearance here today, Auditor-General and executive staff. I take it that the gallery comprises staff from the Office of the Auditor-General as well. Thank you for coming today. I thank our Hansard reporter for being here to cover proceedings. A transcript of these proceedings will be available on the committee's webpage in due course. No questions were taken on notice. I declare this briefing closed.

The committee adjourned at 9.59 am.