T R A N S C R I P T

INTEGRITY AND OVERSIGHT COMMITTEE

Inquiry into the Operation of the *Freedom of Information Act 1982*

Melbourne – Tuesday 12 March 2024

MEMBERS

Dr Tim Read – Chair	Eden Foster
Hon Kim Wells – Deputy Chair	Paul Mercurio
Ryan Batchelor	Rachel Payne
Jade Benham	Belinda Wilson

WITNESS *(via videoconference)*

Professor Lyria Bennett Moses, Director, UNSW Allens Hub for Technology, Law and Innovation, University of New South Wales.

The CHAIR: We resume our public hearing for the Integrity and Oversight Committee Inquiry into the Operation of the *Freedom of Information Act 1982*. Just before we start, Professor Bennett Moses, there are some formal things I have got to cover, so bear with me.

The evidence taken by the Committee is protected by parliamentary privilege, and you are protected against any action for what you say here today, but if you repeat the same things elsewhere, including on social media, those comments will not be protected. Any deliberately false evidence or misleading of the Committee may be considered a contempt of Parliament.

All evidence is recorded by Hansard. You will be provided with a proof version of the transcript for you to check when it is available, and verified transcripts will be placed on the website. Broadcasting or recording of this hearing by anyone other than Hansard is not permitted.

I welcome Professor Lyria Bennett Moses from the University of New South Wales, and we are expecting to be joined at some stage by Professor Toby Walsh, if the technology allows.

My name is Tim Read. I am the Chair, and from my right we have Eden Foster and Ryan Batchelor, and on my left, Paul Mercurio from the Committee. And Jade Benham is online.

What we might do is kick off then. Professor Bennett Moses, thank you very much for joining us. Did you have any brief opening remarks, or shall we go straight to questions?

Lyria BENNETT MOSES: I am happy to go straight to questions, unless you wanted me to summarise my written submission.

The CHAIR: We have read your submission, and thank you very much for making it, so why don't we kick off? I ask you: In your view has the public's expectation of government transparency and accountability changed in a digital age and with the advent of artificial intelligence [AI], and if so, how?

Lyria BENNETT MOSES: I suppose there are a number of ways to answer that question. One relates to the point I made in the submission. I mean, people want to know what information government has, both generally if it is in relation to a program they are interested in, but more particularly if it is in relation to information a government department or whatever has that relates specifically to them. One of the real challenges with artificial intelligence is that the way that data is held is quite different to the traditional way data is held. So, if you go back and you think about 'What files does the government have on me?'– kind of a question – they would be files in a drawer in a particular agency, for example, that would be clearly identifiable as documents, records. The terminology varies often by jurisdiction, but they are similar concepts and it is very easy. Then what has happened, not just with AI but with earlier technologies, is the complicating of that scenario to the point that sometimes it is not clear either which agency is responsible for which thing or even whether the thing needs to be provided. If you go back, even just to concepts around cloud computing, it is no longer very clear that one agency has that filing cabinet sitting in its department. It can now be a situation where a number of different agencies have access to particular databases and what the level of access is for that information to be disclosed.

But even more concerningly, if you think about artificial intelligence, large language models and the way that they work, the way that information is held is incredibly diffuse. It might be words – 'John' and 'Smith' are close to each other on a particular dimension in relation to a particular thing that means that 'John Smith' links to maybe a particular activity and so forth in the way that the words are arranged on a sort of complex multidimensional set of parameters, right. It is not as simple as the government saying, 'Here's what we know about you.' You could query the language model, the chatbot if you like, and get slightly different information about any particular individual at different times. But if someone actually wants to know what is known, its probabilistic relationships to concepts in a really complicated giant parameterised graph – that is sort of how I picture it anyway. Toby would be able to describe it much better because he is actually a professor of artificial

intelligence. But the point is it is very diffuse. You can have something like that and not be able to point to any documents or records or anything concrete that government has, and I think that creates a real problem.

The second part of the question is: What are the expectations of transparency, in particular around AI? I think that people do want to know more about systems making decisions than perhaps they want to or need to know about individuals making decisions. If you have an officer in a government department who is making a decision about me, I do not really have a right to know what their hobbies are, what they do in their spare time or anything else. I am entitled to reasons for decision, but that is really all. Once you are talking about systems, people want to know a lot more about the system than what they might be entitled to know about as an individual, and there are really good reasons for that. There is a lot happening at the moment, and one of the things I work on is artificial intelligence standards. Standards Australia has a committee on AI. I sit on that committee, and we work through ISO/IEC Joint Technical Committee 1 in international meetings to draft international standards for artificial intelligence. Transparency is critical in terms of the ethical considerations, which it is sometimes described as, and what people's expectations are around the use of AI systems. But it is important to understand what people want to know about those systems, because it is not necessarily an easy answer either. If someone in government made a decision about me, I do not want a functional MRI of the person's brain while they are making the decision. I do not need to know which neurons fired, right; I want to understand why the decision came out the way that it did. But computers are very different to people, so getting a computer to give you its 'why' can be a lot more complicated – and a lot more things can go wrong. One of the things is simply fidelity. A human could lie, I suppose. A human might have a real reason for the decision but not tell you, and hopefully that is dealt with through supervisory arrangements and guidance to officers within agencies and so forth. But a computer has to generate the reason through some process. It is not necessarily the case that the computer will lie about its reasons; it is simply that the reasons that are the output of that system do not correspond to the actual logic underlying the original decision. But, of course, the logic that underlies the original decision is not going to look always like human logic, depending on the nature of the system. You cannot ask it the same way you ask a human. A lot of work has gone into the question of what kinds of information you want in what circumstances, and I can get into some examples of that if that is useful.

Perhaps the final form of transparency that really matters – and this is something New South Wales has actually recently done, which was a sort of audit of all of the automated decision-making and AI systems in government that was run by the Ombudsman and then providing that information publicly.

The CHAIR: I just want to go back to the first part of your answer, about how the hypothetical John Smith could be associated with various other bits of information and AI could join that together. In your submission you talk about whether that is information or not. I wonder if you could just tell us a little bit about how AI is being used currently or might feasibly be used in the not-too-distant future, leaving aside its potential for mistakes – but how it could be productively used and how we should think of it being used.

Lyria BENNETT MOSES: AI is a very large category, but maybe I will go through large language models, which is where a lot of the current interest is in these kinds of tools. Large language models can be used in government to help write text, for example, so if you prepare a letter to somebody or something else. It can be used to help provide reasons for a particular decision. It can be used in a query-and-response-type format, even to provide a decision that could then simply be copied. Now, the levels of that, in terms of what would be authorised and what would not be, quite a lot of that is actually still to be determined. I do not know Victoria as well, but in terms of guidance given to public servants around the appropriate and inappropriate uses of that technology, a number of organisations are looking at taking the general large language models, so ChatGPT or whatever it is, and having them trained on local data, so letters to constituents or letters to citizens and members of the public, or whatever it might be. They take the underlying large language model but are sort of fine-tuned on the datasets so that they produce stuff like what is already in the dataset. Now, government will also have a lot of data at its disposal, and it could even use that data through a large language model interface to construct the various things: the correspondence, the decisions, the reports, whatever it might be.

I think there are a lot of issues there, but I will stick to the one that we raised in our submission. Anyone who has ever typed their name into ChatGPT, or a friend's name to see what it says about that person – and I think a lot of people have played with it – will find out that actually it makes mistakes, but it does have a concept of that person. If I type in 'writer's CV for Lyria Bennett Moses', it will put in things that are not true. It is not a very good CV, but nevertheless it is connected to things that are true. It will pick up that I am an academic, for example. It will come up with fake publication names, but they are the sort of thing maybe I could have

published. The words that are in the title are similar to the words that might be in the real titles and so forth. So there is a concept of me in there – not directly, not that there is a Lyria Bennett Moses parameter as such in the actual language model – but maybe my names are there and close to each other in terms of where they sit on the giant multidimensional graphs, so that once it sees 'Lyria' it knows that the next word is more likely to be 'Bennett' than it is to be 'Smith', for example, and that would also be near other words associated with me. It might even be broken down into syllables. I do not know because I cannot actually see inside the model – and no-one can; these are generally closed models – but nevertheless there is some kind of concept.

I guess my concern is this, if I put it in a nutshell: if we move into a future where these tools get better to the point that they are being used more and more in government – and we are seeing that over time that is true; there are less documents and records as such and more databases or, going in to large language models, much looser associations between entities in a fairly diffuse concept – then what we risk is losing some of what Freedom-of-Information laws give citizens, which is the ability to make enquiries, like I said. You know: 'This government department, what files do they have on me?' That becomes a much more diffuse concept. So, we need to go back and actually think in terms of the language used in legislation, to go beyond words like 'document' and 'record' and so forth and think about data and what data citizens are entitled to. That is easier if it is a database that you can save it in, and they might be entitled to have a particular query put through the database to extract information that is, if you like, about them inside that database. It gets a lot more complicated with large language models, which have a probabilistic component to them, so they will never give you the same output in response to the same query. It is a hard question, but I think it is a really critical one, because if we do not do that, either we say, 'Don't use the tools,' and we say, 'We're not doing that because that creates a problem in terms of Freedom of Information and public transparency,' or you have a world in which somehow loose, diffuse information that relates to me probabilistically is used in some way that I can never get my hands on in government decision-making and, going to the third point, if we are not even telling the public that such systems are being used in the first place, without even me knowing that is happening.

The CHAIR: Thank you. I feel like you have actually answered the next couple of questions we were going to ask. So, I wonder, Paul, if you want to ask maybe four.

Paul MERCURIO: Yes. I was just thinking, I have never put my name in ChatGPT, and I am a bit scared to in case what they say about me I like better than I actually like myself.

The CHAIR: We are all going to do it now.

Paul MERCURIO: It is scary, isn't it, because some of what I take away from what you have said is that AI does not actually have to be based in truth with its answers because it is actually sort of trained in a sense. Just that sort of concept is a bit frightening. So how does technology, including AI, impact on what is required of agencies regarding information management and record keeping for Freedom of Information, and how can a government facilitate this?

Lyria BENNETT MOSES: Well, going to the point about truth, just to sort of stop back there for a second because I think that is a really important point, the really good example is if you type into ChatGPT 'The cat sat on the', it will answer 'mat'. It has never seen a cat. It has no idea whether cats actually sit on mats.

Paul MERCURIO: Or what a cat is.

Lyria BENNETT MOSES: What it knows is that that is a common phrase in its training set, right, so there is no truth filter at all. It has no concept of truth. If you like, it sits in a box; it has never observed the universe. It has only been trained. But going to your question about the responsibility of agencies – I mean, rather than sort of interpreting current law, which I think is vastly inadequate – I think it is largely a question of what government wants to do to protect democracy, rule of law and public transparency going forward. At the very least I think the kind of action New South Wales has taken of knowing and allowing the public to know what systems are used in decision-making is really critical. I think the public have a right to that. I think it is important to preserve the right to reasons for decision in the sense of administrative decisions taken that affect an individual, but it is important to preserve that with integrity. In other words, the reasons for decision have to be as good as if not better than – and I mean that because systems are different to people, so what we can expect for them can actually be more in some contexts – and certainly not less useful than the ones we are getting out now. And that includes bearing in mind that reasons that the system gives can be misleading. You can see that,

you know, if you do play around with ChatGPT – and I am not picking on them – or any of these kinds of tools, it will give explanations for things that are not the actual explanation.

There was that example in New York of a lawyer. I do not know if you saw that one in the paper, but there was a lawyer in a hearing who had made submissions on the basis of cases that did not exist. They were fake cases. But, of course, he said, 'Well, I thought it was true, because I asked ChatGPT if these cases were authentic, and ChatGPT said yes.' That example, if you like, highlights that when it says 'yes' that does not make it more real than it was the first time it said it. It is not really supporting anything. So we need to ensure that public decision-making can maintain the same quality of reasons for decisions, recognising the flaws in it – but, you know, this is a constantly developing field, and it is certainly plausible that reasons will get better, or the constraints we put around those reasons, or the reasons could be checked by humans or some other mechanism, or some combination of those, that means that we can do that. But I think that is the second kind of important thing.

Then the third one is that agencies understand their responsibility for data that is in more diffuse formats. It is very easy to read legislation for a bit of paper in a filing cabinet which is a record and a document and all of those things; it is much harder when even just taking documents from the cloud to know which agency has the responsibility. It is much harder when looking at databases, where there is already case law; I think this might be New South Wales, so it is not Victoria. But it is potentially problematic, if not clarified in legislation, that agencies do not have an obligation to query databases. In other words, if information is in a database and an agency is asked for that information and the record does not already exist, because no-one has actually already queried the database and printed out the response, then they do not have to search the database. Now, to me that is already problematic, because that is just how information is held by government now. It is always in databases. There is always some process to get things out, and no individual item or combination of items in that database is necessarily going to satisfy the kinds of language that we used in a paper-based world with freedom-of-information laws.

Then I do not even know what the answer is going to be for large language models. I think that is more complicated. At least with a database you can ask for a particular query, but with large language models, as I said, it is hard. So, I think there are some really big questions there about how citizens can query what associations government is making about them where they use those kinds of tools.

Paul MERCURIO: Thank you.

The CHAIR: Alright. Ryan Batchelor.

Ryan BATCHELOR: Professor, lawyers and drafters like to put things in black and white and define everything within an inch of its life. Our FOI [Freedom of Information] Act is 40-something years old, and technology likes to change very, very fast. How do we get our legislation to be responsive enough to the pace of technological change to ensure some of these fundamental protections, like knowing what information government holds about an individual, which is one of the bases of FOI law, so that citizens can know what the government knows about them? How do we go about the task of doing the drafting in such a way that can enable that, and is anyone doing it well now?

Lyria BENNETT MOSES: I am happy to share an article I wrote a few years ago on this question. I did a giant audit of all jurisdictions – so territories, States and federal – of all of the laws, and everyone uses different words. Everyone uses different words for two things. One is the relationship between data and an entity, so 'holds', 'possesses', 'controls', 'custody' – you know, there is a whole variety of language that is used. Then also of course there is language for the thing – is it 'data', 'documents', 'records', 'information'? Then what is the definition of those things? And it gets even more crazy – speaking of lawyers defining everything – because a lot of the time those definitions loop. You know, there are examples of 'hold' being defined in terms of possession or control and then 'possession' being defined in terms of control and so forth, so you try to actually read the sentence with all the substituted words and it is a minefield. The problem is that none of the words that are used most commonly are actually helpful as data becomes more and more diffuse. I am using that word, but I think I have explained what I mean by that now. Take 'hold' as an example, which is often defined in terms of other words, but just thinking about it as an English word it suggests a physical grasping: 'I am holding in my hand the file or holding in my office the file'. It gets much more complicated when you start to look at even cloud computing, which is what I was talking about then. If you go through the words, most of them have those kinds of problems.

What I suggested in the paper was that we really need a few core concepts that I think are relatively impervious to technological change. One is that an agency might be able to access data. They have not yet necessarily accessed it, right? They might never have used it, but they have access to, for example, a database portal that would give them access if they made a query for some information. That is scenario number one. The second one is that they have accessed it, so there has been some kind of data processing by that agency. They have actually processed data about me. They have done something with data about me, not just have the potential to be able to do so but in some context they have. The third is that the agency has control over data. In other words, there is a database, and they decide what is in that database and if you wanted to make an amendment to that database that is the agency you go to do that. The fourth is that they have physical possession – and I do use that word here, 'physical possession' – over the media on which the data is stored, so the server sits in the office of a particular agency.

The rights and obligations with respect to those things might be different, in particular what people have a right to request. There could be a difference, for example, with me being able to request what an agency has processed in terms of data about me versus being able to request an agency to start processing data about me simply because it has access to the possibility of getting that data. You can define these things, but I think those categories are relatively clear. Are they perfect? I do not know, because no-one knows what the future will bring, but I think they are a more helpful way of thinking what obligations different agencies have in different circumstances.

The second thing I think is very hard for any single Parliament, Victorian or otherwise, to do, and that is the challenge that everyone is using different words. That creates a huge problem because of course a lot of data crosses State lines, and for that matter jurisdictional lines, so State agencies getting data from the private sector, sharing it with an agency in different state for some purpose and so forth. Having different words is unhelpful. I say that there are better words, but I also acknowledge that against that challenge is the need for uniformity or some level of common language to use with the different laws in the different jurisdictions.

Overall, I have written quite a lot in my research about how you design legislation in the context of constantly changing technology. You can never do it perfectly, and you would not want to do it perfectly. The example I have used in my own scholarship is: imagine you are working in road safety. You can be really technology-specific: all cars must have ABS brakes. That is very technology-specific. Maybe there are other kinds of braking technology that you want cars to be able to use, so you might set a performance test: cars must be able to stop on a slippery road within 10 metres of pressing the brake or whatever. You say, 'Well, that is better because you can do that in different ways.' But then I tell you, 'Ha, what about if cars don't stop more quickly but they harmlessly bounce off each other due to some special force field I have invented?' I am being very sci-fi now, but you get the point. The old rule is now obsolete because we have a different way of protecting passenger safety.

You can constantly strive to be better and more technology-neutral, but there is no end point beyond saying 'cars must be safe' or something similarly meaningless that will not change action. But what you can do in things like Freedom of Information is unpack the concepts and look at the extent to which I think generally speaking most legislation assumes something like a document. We went from physical documents. We now recognise electronic documents, but the imagination is still a Word file or a PDF or something like that. I think the concept of 'data' or 'information' – and those words, I query whether they should be synonymous or not; some people use them synonymously and others distinguish. But I think those concepts are relatively impervious to technological change, unlike concepts like documents and record, at least as I have seen those definitions play out. If you think about what the thing is that people are entitled to know – the data, the information – and you think about how does an agency become responsible for doing something with respect to that, and you can think about the four different kinds of relationships an agency might have, I cannot promise that there will not be some future technology that will raise questions and will need to be addressed, but I think it is a long way from where we are now.

That still does not necessarily deal with the full large language model challenge, because what you have got there is much more complicated. It is not a data entry for Lyria Bennett Moses even in a database or a data entry that relates to some characteristic of mine, like my suburb or anything else. It is, like I said, essentially the closeness of concepts to other concepts in a multidimensional space. How do you provide any measure of transparency over that? I think it is very difficult. I think the most you can do is kind of evaluate its behaviour, so you can say that when this person's name is entered and particular questions are asked about that person,

these are the different outputs we get of a few iterations or whatever. You could do that. Whether that is feasible for government agencies to do in response to citizen complaints would be another question. Whether you set up an API so that citizens can access the same tool directly themselves and query it themselves to see what kind of thing government is seeing when their name is inputted and whether that is feasible, I do not know. At the moment it would be very helpful if Toby had not had technical problems, because he is the AI person; I am only the lawyer. But I suppose I can say they are the questions that I think are important.

The CHAIR: Thank you. Let us go to Eden Foster.

Eden FOSTER: Thank you, Chair. Thank you, Professor Bennett Moses. My question really is more asking you to elaborate a little bit on the impact that AI has had on the section 33 exemption relating to personal privacy in Victoria's FOI Act and what you think can be done to address this.

Lyria BENNETT MOSES: I think the main difficulty with privacy law at the moment is federally it is in a state of flux, and I do not know how the States are going to respond to fundamental changes if they happen at the federal level in terms of privacy law. It is hard to know which issues are already going to be resolved initially through that privacy law review federally when we see a sort of amendment Act come out of that and the extent to which the States are going to look to that for their own reform processes. To some extent, in other words, a lot of the questions I have might be being solved. But I think there are a few, and I mention them in the submission.

One is inferred information, and that, in particular, as I said, is a real challenge I see with large language models in particular, because everything they do is inferred information. They do not actually have facts stored as such. Everything is, if you like, linking concepts. I mean, you can look at it through the lens of government use of these kinds of models in its own work and what is permitted and what is not, but you can look at it even more broadly if you go back to just machine learning type artificial intelligence. There will be inputs. Say you are just looking at a machine learning example – so forget large language models, because they are a little bit more complicated – but government is trying to learn who is most in need of a particular service or benefit so that it can reach out to them and ensure that they get that provision, as an example. It will use what it knows about people from its own datasets, which it will have. It will know, potentially, certain things about me that government has collected and it might use those data points, but along the way in terms of actually drawing an inference for the question that you are asking the system to predict 'Does Lyria need this service?' it is going to make other inferences: both the end inference, 'Lyria does or doesn't need it', but also inferences along the way.

The one way to think of a neural net, which is just a complicated series of steps, if you like, is that each of those is potentially an inference but no-one even knows what the inference is, potentially. So, a particular layer of that neural network might actually be drawing some intermediary conclusion about everybody, but in fact that was not decided by the people who built the neural network, it is just chance that that has happened in layer N. But nevertheless, there actually is information about me – because there will be something that happens as I am fed through that neural network – and it is actually leading to a government decision about me at the end of the day that I am or am not given access to whatever the program is. So that is a complicated scenario in the sense of Freedom of Information, but it is also a complicated scenario in the sense of privacy, because I had certainly never consented to that inference being drawn and so forth. As I said, in some contexts, government will not even know about intermediate inferences that are drawn. So, I think there are important questions about how that kind of information is treated under both privacy law and Freedom-of-Information law.

The second issue that I mentioned that is I think a real challenge in privacy law at the moment is the use of de-identification. Often people talk about the data life cycle, and the data life cycle tends to end with one of one of two steps: there is either de-identification, or there is deletion. Now, deletion in an ideal world is like literally taking a hammer and smashing the drive, because with pretty much anything else you do it is not really deleted, but leave that aside – demagnetising, et cetera, et cetera. But if you just look at the de-identification challenge, that is actually harder and harder to do. De-identification: ultimately, even though often in government policy it is described as an end state – 'This data has been de-identified' – in practice no data is ever de-identified. All that has done is reduce the probability that anyone could associate that data with a particular individual. That is the most you can do. You can make it really, really unlikely that anyone would, but any data, no matter how generalised, could potentially tell you something about an individual.

For example, and I will give you – this is not realistic, but it illustrates the problem: suppose you had data on how many men, women and non-binary people lived in Victoria and I had a list, being a different government department or whatever, of every person in Victoria except for one and their gender. I would then be able to infer the gender of the final person, who I did not know the gender of beforehand, because I have got the totals and I already had access to this data set. What I am using that to illustrate is not 'Don't publish gender data'; that is not the point. What I am using that to illustrate is, that any information that has ever related to a person, like gender, even if de-identified, even if amalgamated, can contain information about a single individual to somebody who has got enough other data. Now, what that means in a world in which there is so much data about everybody out there and lots of people have different parts of it or can buy different parts of it from other people or get access to different parts of it from other people is that de-identification can only ever be a process. Government can obviously decide, 'Yes, we are going to publish gender data, because the probability that anyone would be able to determine a single Victorian's gender from that is very, very small'. That makes sense. But it is only ever a process.

That creates a problem, because, like a lot of things in AI, there is a constant competition there. De-identification techniques can get better and more robust, but re-identification techniques also get better and more data is constantly being put into the system. And what that means is that it is not an easy question, and it is not a question of set and forget, because something that is a dataset that is de-identified today might, given a new technique tomorrow or the availability of a completely different dataset tomorrow, no longer be de-identified. And I think that is a challenge that privacy has not yet dealt with. It assumes that there is something at the end that is now a free-for-all.

Eden FOSTER: Thank you.

The CHAIR: Thanks. Jade Benham, do you want to do one last quick question? We have just gone a little bit over time, but –

Jade BENHAM: I think I know the answer to this, but I would not want to make an inference. Statutory protections – do you think that in Victoria we are going to need additional statutory protections moving forward?

Lyria BENNETT MOSES: Statutory protections for what?

Jade BENHAM: If we were to ensure that the FOI scheme is more responsive in the way that information is created, used and stored in the digital age – I suppose that is the key bit that I left out. Are we going to need increased protections there?

The CHAIR: And protections for FOI officers or departments against the consequences of releasing information I think might be the intention there.

Jade BENHAM: Yes.

Lyria BENNETT MOSES: Sorry, just so I can understand the question, then, it is about the safe release of government data. Yes, there is a need for increasing vigilance around that, and there is a need to, I think, recognise that any release has an element of risk, in particular any release that relates to personal information. I mean, go for your life on the weather charts or something, right, but for any release that is related to individuals there will be an element of risk. Sometimes the risk is so low that the risk of not releasing it and withholding information from the public, who have a right also and an interest in the workings of government – the greater risk is there. I think saying because of the historic release anything would be a similar mistake, but I do think that there is a need to think through what the process is for release. To my mind the best process is, if you like, 'red team'. In other words, if you are releasing something that is in a particular risk bracket, then you 'red team' it. You get a team to say, 'Can you identify people in this data?' Before you release it you actually get the data scientists on it, and they see what they can extract. Then you have tested it. That is one approach.

Another approach is to look at some of the – and this is not my area, but I know that there has been a lot of research with Data61 and so forth around PIF, I think they call it, which is a personal information factor which allows a calculation on data that can be useful. I am sorry, I am describing that really badly, because it is too technical for me, but there are certainly tests that can be performed that can give information. I do not know if it

is extra protections. I am not sure if that is how I would frame it. I certainly think it is extra processes that should be put in place to protect individual privacy in the context of release under freedom of information.

Jade BENHAM: I guess it is hard to answer before we know what the systems being used are. Okay. Thank you.

The CHAIR: Professor, thank you very much for joining us, and thanks very much for broadening our minds on this. It is possible that we may have some more questions for you as we continue in the inquiry, but thanks again. I declare today's hearing closed.

Committee adjourned.