

# Webinar Transcript

## Governing AI Risk:

### What every director needs to know

Hi everybody. Welcome to our webinar today titled Governing AI Risk What Every Director Needs to Know. This is with Helen Van Orton and Alexie O'Brien. My name is Sean MacDonald and I shall be your moderator in the background for the next forty odd minutes.

Firstly though, thanks for attending today. We always appreciate the effort you make to be here for our live webinar events. During the session, if you have any questions, please try and use the Q and A button on your toolbar. We'll be answering as many of these questions and try to get through as many as we have time for. And finally, if you stay through till the end, which of course we hope you will do, and as is very customary for our webinars now, we have a special treat for you. By answering a one minute survey at the end of the webinar, you'll go into our draw to win one of our beautiful gift hampers.

Okay. For those who don't know too much about BoardPro, we are a board software provider, sometimes called the board portal. We serve around, oh, about thirty five thousand users around the globe these days.

We enable organisations to prepare for and run their board meetings more efficiently and effectively with less time and deliver more impact and value for the organisation.

As much as we are a board software provider, part of our wider mission here at BoardPro is to make the fundamentals of governance free and easy to implement for all organisations especially those organisations with resource constraints.

And one of the many ways we do this is by providing free access to hundreds of business templates, guides and resources, which you'll find in the resources center of our website.

And these webinars of course that we host are a great way of accessing really great governance content without the time commitment and costs associated with in person events.

So for the next forty odd minutes just relax, sit back, listen and add to the discussion by asking as many questions as you would like. A full recording of the webinar along with the slide deck, the transcript will be sent to you on email tomorrow.

So let me hand over to the team starting with you Helen to introduce yourself to everybody.

Thanks, Sean. Kia ora, everyone. I'm Helen Van Oten. I am a professional director and board chair across a range of, different New Zealand organizations, but I also lead two companies where AI is deeply embedded in how we work.

I literally use it every single day, whether it's streamlining operations, supporting decision making, driving innovation, but in every sort of different area of operations that we run. I'm the managing director of a company called Directly, where we actually work with boards and executives to help them lead with confidence in the AI era. Very much thinking about AI governance, but including training, keynote speaking, and strategic board sessions focused on governance, risk, and opportunity in this really fast moving space. Over to you, Alexey.

Hi everybody, I am Alexey O'Brien and I am also a director and chair of a private ancillary foundation and a small private family office.

I am also the director of Leadership Academy and what I do is help businesses bring to life a culture of AI enablement.

It's important to look at it from the top down, which we'll be doing today, but also really once we've got those guardrails in place, embedding everybody with the tools and training to operate within those guardrails from experimenting to making sure they're abiding by their governance rules. So that's me. Back to you, Sean.

Back to you, Helen.

Hey, beautiful.

I just wanted to start off the session by framing it up because we talk about AI and governance and risk, actually it's probably really important to just take a little step back before we kick off. So just to reiterate, today is not it's not gonna be a technology presentation, hopefully, Alexei and I went through in any terms that are not familiar, and if we do, we'll explain them as we go.

This is really about governance fundamentals, and so I started off by saying, you know, why is AI at such a tipping point for boards? And part of that is because your board accountability is increasing. So the Institute of Directors in New Zealand did a survey and sort of said, where does what are the key issues for directors over the next year? And AI was right up there in the top five.

There's a lot of organizations talking about it, and Deloitte actually did a survey. Actually, they do a very regular survey, but the one that came out just a couple of months ago said that actually only thirty one percent oh, sorry. Only thirty one percent of boards don't actually have AI on the agenda at all yet. So if it's not on your agenda, you should be really, really worried.

And sixty six percent of boards were saying that they don't actually know enough about AI. So for those of you who are attending this, fantastic. That's about upskilling and sort of starting to learn about it. But what we want to do today is to sort of start talking about some of those risks in the AI landscape, which we'll go through in a minute.

The second one is about the regulatory scrutiny, and Alexei is going to go through and talk about some of the details of the legislation that's relevant in Australia and New Zealand. Sorry to those of you phoning in from further afield. We've got, like, a little worldwide bit, but we're not going into legislation of every comp country.

But that landscape is changing, so the EU's AI Act is just sort of coming in in force. The American market is complex and confusing, and I'll leave Alexei to talk to that one.

But regulators are increasingly scrutinizing what's happening in AI, and you don't want it to become your cyber where you just didn't think about it, and you're trying to retrospectively fit things in. You need to be addressing this as AI starts to be deployed in your organization.

And finally, globally, seventy six percent of companies are now saying that they have got AI in use, and sixty nine percent have got generative AI. So k p m g that's from a KPMG University of Melbourne survey. And the more local figures on that is sixty one percent of New Zealand companies are saying they're using AI and sixty five percent of Australian companies. So we've got really high penetration that obviously differs between the size of the organization, but there's really high penetration of AI. And so as a board, thinking about the governance frameworks for that and what we need to be thinking about is really, really important, hence this webinar.

So, Sean, if we can pop on to the next slide, that would be fantastic.

I just might add, I think, you know, it's not just regulators that will be increasingly pushing for visibility around this. It'll be investors as well. Totally. Superfunds starting to ask questions around, you know, managing that algorithmic algorithmic biases.

Yep. Yep.

Helen, I've got a question that's come in here, which we'll put up right now. Lots says lots of organizations have gone down the ISO twenty seven thousand information cybersecurity standard. Should boards promote the ISO forty five thousand AI standard for their business? What are your thoughts on that?

I think, honestly, it depends on the size of your organization.

The bigger the organization, the the more you need to be taking it seriously. I think if you're a small SME, then that's probably the complexity of implementing any of those ISO standards is really it's a lot of work.

So make sure that it's fit for purpose. And we I don't know if you cover that at all in yours, Alexey, but I would do.

Forty two thousand and one.

Yeah. Let's come back to that a little bit later. I'm gonna start a little bit more of a helicopter view. And what I want to do is just actually start unpacking what some of the risks are.

So it's there's a lot of jargon in the AI world and what I wanted to do is just actually bring some of these things to life so people can get a bit of a feel for what they are. So I'm just gonna run through some of the big risks that are in the AI space that boards need to just start thinking about. So let's start at the top there with ethical risks. So that could be in your bias, your discrimination.

It's the unfair outcome. So if I think of an example of that, in Australia, there's been a big study a big case recently about an organization that their recruitment was actually really discriminating against migrant women, against refugees, against First Nations people.

Other AI recruitment tools have disagreed sorry, discriminated against people who've taken career gaps. There's a whole load of examples in health care, in lending, where there's racism. There's sexism. There's ageism.

Now this isn't because AI is being manipulative or difficult. It is training on the data that is out there. So if the data that is going into the model is informing certain things, then you might find that that that bias just comes through automatically. And so as a board, understanding what is going into my model, and we'll talk about that a little bit in explainability, making sure that you don't have unfair outcomes.

You know, Amazon had a recruitment system that they brought in, and it had massive discrimination against women, so they actually had to get rid of it. Hence, that example on the top left there. So thinking about what is going into your algorithms and into your AI models is really, really important to make sure from an ethical basis what your company is putting out. And that's not just your company, by the way, that's your suppliers, you know.

They might be using AI to provide you with something. So, this goes all the way through the supply chain. It's not just within your organization.

Explainability is AI is very often a black box sitting in the middle of the table and your exec team will come in and they'll talk to you about it and they'll be like, oh yeah, it's doing this. As a board, you need to understand what is sitting underneath that black box. So, is there cognitive bias? Is there historic bias? Algorithmic bias? Is your data complete? You know, where are your datasets actually coming from?

And actually, is it copyrighted data that you've imported from elsewhere that is not actually meant to be part of yours? And even when you've got your model absolutely beautifully set up, then you start moving into things like model drift. So does it continue to be? So having explainability on your agenda of constantly going, do I understand what is going into this box and how are management keeping this transparent and open and valid? Like one of the organizations I'm on, we work with the University of Auckland to make sure that our AI models are always explainable, just to make sure that we've got that that depth of assurance in place for us as a as a board.

A little bit more about hallucinations. This is probably a term that you're a little bit more familiar with. So hallucinations is where AI, which always wants to please you, and if it doesn't know the answer, it will happily make it up.

So this is where AI is making it up. So just a little example of that, Chicago Sun Times, about five or six weeks ago, published this article for their, you know, their American readers, their summer reading list and what they were going to read whilst they were on their summer holidays. It had all of these book recommendations.

The only problem is is the person who'd written that had used ChatGPT to do it, hadn't gone back and checked it, and actually half the books didn't even exist. They were made up storylines. Some of the authors didn't even exist.

But the person had just taken this list and gone, oh, yep. I need AI to do this. They hadn't validated. They hadn't checked their work, and they had created this thing that was a hallucination. It's quite a funny story, right, because it's a newspaper and it doesn't really matter. It's not something important. But a more important one of those is if you look at that Air Canada example where the Air Canada chatbot made up an answer and said, yeah, you can get a flight at a special reduced rate because it's bereavement.

Well, actually they couldn't, and the AI had made that up, and suddenly Air Canada is in court, and the board are going, oh, well, wasn't us, or who are defending it, it wasn't us, it was the AI. Needless to say, they were found liable and had to refund the customer. So, hallucinations can be sort of really quickly embraced in your organization. And again, back to that KPMG research that I was just talking about earlier, that the global numbers for that is saying people who are using AI at work, fifty percent, that is half of the people, are not actually checking their work.

And, they are just letting it go out as theirs, they're claiming ownership for it, but they're not actually going back and checking their work. So, hallucinations is a really yeah, false citations, Duane. Absolutely. The amount of legal cases in the US where made up case law has come out, and there's been all sorts of issues in the courts where people are going there.

And, yes, Marco agree ethical risks are border broader, and I did talk about violating IP and stuff. We can only cover so much, so I'm just trying to do a really high level jump into the all of the different topics that are a very, very high level.

But thank you for sharing that. People, please read the chat because there's some really fantastic little insights coming in on the chat there.

ShadowAI. This is where people in the organisation are using AI when they're not actually authorized to use it or using it outside of the guidelines that are in place.

Globally, again, that KPMG survey, and this is literally only a couple months old, this data, sixty percent of people are using AI at work, and almost half of them are using it in a way that they know contravenes the policies at work. So in New Zealand, that's thirty four percent and forty four percent. But, you know, this is that's more than a third to half of your workforce who are

using AI outside the guidelines and guardrails, if you've even got any. So there's a whole load of risk in there where people are, know, what of your company information are they putting into AI that is just going up and training the model and putting your information out there.

Absolutely.

Helen, it creates exposure to data leaks and copyright violation, privacy laws as well. And it's happening without visibility sometimes too.

Totally, totally. And then just super quickly, the legal and regulatory, Alexei is going to cover this in a lot more detail later, but I couldn't have the list go up without just covering it off, which is whether you're breaching privacy laws, whether you're breaching director duties, discrimination laws, etcetera. Then there's just a final little lens that I want to put onto the AI risk landscape because this is very internally focused. So, this is looking at you and your organisation and what are the risks within your organization.

But as a director, you've got that duty to make sure you're looking more broadly. What are the external implications of AI for you as an organization? So one of the ways I like to look at it is, you know, if you use the Porter's five forces model and you go, well, okay, from a supplier perspective or a buyer perspective or a new entrant or a substitute, all of those different things, where does AI actually pose a threat to your organization of somebody being able to do something differently in a way that massively impacts you. So it's very easy to get caught up and, you know, we've only got limited time, so these are a sort of a super quick run through.

It's very easy to get caught up on the what do I need to think about on the inside, but if you're also not looking externally and not thinking about that wider focus, you are also exposing yourself to a whole load of probably less obvious risk from AI.

Helen, I've got one more I'll just throw in the mix as well. Very common as well, misuse risk. So the team's using it, but they don't understand what it's doing and whether, you know, what, how they're using it for the task and whether it's appropriate or not, which also can be considered or countered with education and training.

Yeah. Yeah. If anyone wants to get into this in a lot of depth, the the Paris AI Summit earlier this year did a really good piece about AI risk, and you can go and get lost in that document. It's very long.

But there's a there's and that's where there's a hollow about misuse and systemic risk and manipulation and all sorts. But we're trying to keep it high level and digestible today, not go too much into the depth. Okay, Sean, can we pop onto the next one, please?

I just thought we'd tackle a few questions, if you don't mind.

I've got three that Can you see them there?

I was looking on the chat as opposed to question. I'll just do the one on the chat. At a meeting level, minutes taken by AI.

Funnily enough, I think it was Alexei and I did a webinar on that just a few weeks ago. And certainly for a board meeting at the moment, my recommendation is very strongly be cognizant of when you're using AI and be very intentional about when you do and don't use it because there's a whole load of issues around reaching confidentiality, about your board minutes being available and to be subpoenaed in court. If it's a conversation that's under legal privilege, does the AI break the legal privilege that's not being proved in court yet? So, there's a whole load of things to think about, having AI in your board meeting.

So, I would we'll send out the resource for the Paris Agreement. I would be very careful about using the AI for board minutes, and I've watched that webinar back. It's on the BoardPro portal. But yeah, and sorry, Paris resource we will send out as notes with the meeting.

I'm just gonna jump into I can take that first.

How do you ensure the algorithm is going into the toolbox as well rounded and unbiased? We can't control what's going into it because it is trained on other data, anything that's available to it. However, what we can do in our outputs is we can check and add I've got in my global settings inside of my tool which is ChatGPT.

Give me where this output might have an inbuilt bias in it. So it's just like flagging it to me that there might be bias in the output. So I recommend that everybody in like, you can do that even in a prompt when you're doing some thought work or add it into your global setting. So we'll remind you every single time.

The KPMG University of Melbourne, I can't remember the name of the study, but again, we can get that sent out with the resources afterwards. Just to note, there is a global study, but then there is also a study for each of the individual countries. So, if you're in Australia, you can download the Australian one. If in New Zealand, you can download the New Zealand one. And I literally have the two summary pages of those studies printed out and on my desk for both countries the whole time, because there's just so much insight into there.

How do you stop hallucinations that are given direct to the public with no organisational input? Well, you make sure you've got the right guardrails and guidelines in place for your team to say, don't send stuff out with filtering. And if you have got a customer facing large language model, you need to make sure that it is really well tested.

That's probably something Alexei will answer better than I will because she's more on the operational side, I'm more on the governance side.

In terms of that stop, I don't think you can stop the hallucinations. You can't stop. If someone was wanting to feed false information about your business outside of your company, you couldn't stop it. But I think that's where, you know, this human in the loop principle is really important and that we have everybody trained to understand you like and working appropriately

in their sphere of expertise so that they know that the output is correct or not. It is a risk and that's where I think that we've got to, you know, have those guardrails, understand that a hallucination actually isn't a bug necessarily. It's just the way the generative AI model works.

And this is how we can make sure our teams using guardrails and training know how to interact with the AI.

Yeah. And I just want to build on that comment that Alexei made there about human in the loop. It is a really, really critical part of how you are looking out after your array, making sure that there is that human interaction, there is that checking. And we're not going to cover it today.

But for those of you who are sort of more into AI and are starting to explore agentic AI, then you start moving human in the inner loop or human in the outer loop. And if you are moving into agentic AI, having that human oversight and thinking about the risk of your systems doing more without human oversight is really, really critical. So we're not gonna go there. But if you are in an organization starting to use AgenTik AI, there's a whole load of other risks that you need to start thinking about.

I'm really conscious of time, so we're gonna just jump onto this one.

And this is really just a checklist of things to think about. And it's, again, it's not absolutely everything. We'd be here all day.

But it's just a sort of a high level one. Interestingly, the New Zealand government in the last couple of weeks has just published its strategy for AI. And in that, it's sort of it barely references governance. But one of the things it does say about governance is making sure that you are building that AI, oversight into your existing systems and setups.

So I think that's actually really quite powerful. So if you look at the second one of those, have you got a robust AI governance process? Have you got policies and frameworks that you need to manage it? But also within that, where is it sitting?

So within your organization, is it sitting like some of the boards that I'm on, it sits at board level, but it also sits in the technology committee and it sits in the people and culture committee because we're looking at it in people and culture, from an ethics, from a change management, from technology, we're looking at sort of the practical and logistical operations.

So, there's different areas, but as a board as a whole, we're looking at it, and it's also sitting in our risk committee because we're making sure that we're looking at it from a So, it's sitting across multiple committees.

And if on your board it hasn't even made it onto the agenda yet, be mindful of that. You know, there are organisations who are using AI not only within the organisation, but if you are governing it well, it's giving you a massive competitive advantage.



That first one about where is it used, having that inventory of where it is in your organisation, and that's not just the actual AI, that's the shadow AI that we were just talking about.

Making sure that your risk is always an offset of risk and opportunity at a governance level. So, is it in your enterprise risk management framework? What's your risk appetite around AI? Because it may be that it's, you're going, actually, we're really keen to do this, but it's as long as it's within the guardrails. So, you know, don't when AI first came out, was a whole load of organizations saying, hey, you just can't use it. We're not letting you use AI because, well, actually you want to be able to take the competitive advantage of it, but do it in a way that is safe. So then the compliance one, you can read that faster than I can talk, but just making sure that we have that ethical piece referencing all the ethical things we were talking about on the previous slide.

Have you got AI in your strategy? And if not, I suggest you have a strategy day reasonably soon to make sure that that gets captured within it because that is a massive missed opportunity for you, but it's also a gap that your competitors can seize if you haven't actually thought about it.

And then that expertise gap is not just as within your organisation, but as a board, where are your knowledge gaps? Do you actually understand AI? What the opportunity is?

You know, just the basics of it. And again, that's, Alexei and I can both help out with that, her and Ozzy and me in New Zealand, because we both do a lot of work in that space. Alexei, over to you.

Thank you, Helen. So this is the core leadership challenge for directors, not just managing AI, but stewarding it with integrity, which means setting that tone. So directors are the tone setters, ethical boundaries for AI, defining our organisation's AI risk appetite and ensuring there's a cultural expectation of transparency and explainability, not just performance. So asking questions like do our boardroom conversation reflect our values around fairness, transparency and responsible innovation?

If not, we risk AI potentially operating in a vacuum of ethics. So we want some oversight structures. Helen touched on this before as well and stewardship isn't passive. So do we have those technology and ethics subcommittees or AI ethics dashboards visibility, some risk and audit reviews, I'll touch on that in the next slide when we talk about governance.

Thirdly, directors duties. We've got a duty to anticipate and mitigate our foreseeable risk. So AI harms like bias, privacy breaches, faulty decision making are now considered foreseeable. So we're not expected to prevent all harm as directors, but we are expected to show that we've made reasonable inquiries and understood the risks and ensured that those governance frameworks are in place. So failing to ask the hard questions around AI is no longer a mutual act and it could actually be a breach of our fiduciary duties.

And then capability building. So we've got to resist the temptation to outsource judgment and outsource the tech. AI governance isn't a judgment issue, it is a judgment issue. It's not just a technical one.

So we've got to build that sort of fluency, not just to be able to code or not to code, but to confidently be able to ask, is the AI system explainable? Who's actually accountable for the outcomes? It's great that we've got it on multiple subcommittees and in multiple parts of our organisation, but ultimately who owns the AI strategy, the AI risk register and have the ethical implications been assessed. So capability building, it also isn't a one off training.

It's got to be an ongoing commitment to that board level literacy.

Sean, would you mind passing over? So we are going to now touch on our regulatory landscape. So across Australia, New Zealand and globally oversight is tightening and the expectations of board accountability is certainly rising. That have taken longer in Australia and New Zealand than some of our global counterparts.

We are now starting to see some shape to how we're going to be seeing the regulatory landscape or at least the guidance from an Australian and New Zealand perspective. When we think of regulation, we're often thinking of restriction but Australia and New Zealand are taking quite a principle based approach to AI governance as of today. So Australia is laying foundations for what this is going to look like but directors are as of today already expected to govern within the impacts of AI through existing laws. So like privacy, there's just been an update to the privacy act to really tighten rules on AI and data as well as anti discrimination acts and our directors fiduciary duties.

In October, twenty twenty four, the Australian government actually concluded a consultation on ten mandatory guardrails as well for AI systems that operate in high risk settings like healthcare and law enforcement.

And these guide rails are really about emphasising transparency, that human oversight, the human in the loop as well as accountability.

And the government is still looking for how do we actually legally enforce these mandatory guard rails?

Is it through existing laws, a standalone AI act or blended framework? So it's a bit of a watch this space. However, there is advocacy certainly from the Business Council in Australia to urge the government to use a light touch regulatory approach. So we're not seeing additional legislation come into play. So how do we create more oversight use existing legislation to govern that AI risk?

APRA is lifting its standards as is ASIC and ACMA. So if you're on boards that are regulated by those bodies, make sure that you're checking out those areas as well. In New Zealand, they are also taking a principle first approach with strong signals from the government around transparency, fairness and algorithmic decision making. So as boards, we've got to assume that the AI impact will be continued to be scrutinised through the lens of this ethical practise and social licences.

So last week, as Helen mentioned, New Zealand actually released their AI strategy and the government is deliberately removing any unintended legal barriers to support AI adoption. They're not going after an innovation or invention, there are strong messages around adoption, but the legal environment or legislation environment is all about rewarding responsible innovation, not stifling it. So from a board perspective, our question becomes how confident are we that our organisation is ready to innovate within these guardrails without needing to wait until our AI laws are fully codified.

And what's really clear from the New Zealand strategy is that the government is wanting to continue to build confidence in AI by making sure the legal environment is clear, stable and pro innovation. So we don't need to wait for those hard laws but we've got to act responsibly and the risk isn't necessarily overstepping, it could be falling behind.

There's also a strong message across from the New Zealand government around adopting those international alignment principles from the OECD, forming a really ethical foundation. So even if the legislation evolves slowly across Australia and New Zealand, government expectations are really looking to what's happening globally, which I'll touch on as well. So we've got two distinct models emerging globally. So in Europe, there's this really hard line with regulation that is now in place with the EU AI Act.

It's come into force and sets a really high bar for risk and transparency. So, if you're on boards or in businesses that are operating in the EU, we've got to understand and comply with those hard legislation acts. In the US though, we've got a more decentralized approach. So last month, the Senate rejected a proposal to prevent states from making their own AI laws, meaning we could see that patchwork of state level regulations emerge again, much like we saw with data privacy.

And to answer the questions before around the ISO IEC forty two thousand and one, absolutely we should consider it. But like Helen mentioned, it really depends on the size of the organisation, the industry that you're in and AI maturity of the organisation. So it's the first global standard for AI management systems and it really provides a governance framework to ensure AI is used responsibly, ethically and transparency.

And for large enterprises, regulated industries, especially those deploying or building complex AI systems, it's an important signal for trust or risk management as well as future readiness. However, for smaller businesses or those early in the AI journey, full certification might be premature and cost prohibitive.

So we might want to, don't necessarily need to adopt it, but you know, using the principles behind that standard might be useful.

So it's not about memorising every policy, but it is about understanding that regulation is coming from multiple directions. Board accountability has got to be assumed, it's not optional and we need visibility and governance in place before enforcement catches up.

Sean, over to our last slide.

Just a couple of questions. Oh yeah, sure. There's one on chat. Are there any specific risks for nonprofit organisations that won't be covered here today?

I actually did a presentation on this earlier in the week and I think the risk, it's very similar. It's like what are our, how are we training our people to know what data should we be using on the models? It's, you know, the donor information privacy. So I don't know that there's any particular specifics. We've got to comply with our regulators to make sure any industry related specifics are complied with. But I think they're very similar risks across the board with not for profits.

Another one on Q and A. Do you have any advice for school boards, particularly around student assessment and cheating?

Helen, I'll throw this one to you.

Thanks.

I'm gonna have a two sided part of this answer. One is, obviously the cheating aspect we need to address, but as a school board, what are you training your students for? Because their future is going to be a future where AI is part of that future. So there's a there's an approach you can go like, there's a whole load of emerging tools starting to come out like Turnitin AI detection or GPT-zero.

So there's tools out there that you can use to stop students doing it. But actually the bigger question is, is what skills do you want your students to come out? So in the US and in China, they are teaching kindergarten kids how to use AI, because they think it is so important that it's the future. So, Trump signed that off about two months ago, teaching kindergarten kids AI.

So, actually, do you want to embrace it as part of your school curriculum or do you just want to stop people from using it? But making sure that you've got a really clear policy in the school, starting to think about it, I think it, is it Microsoft? Microsoft has got a whole education portal which has got an incredible range of tools and stuff to support teaching.

And it's probably worth, I mean, I'm not in-depth in school sector, but it'd probably be worth going and having a look at that. But I think the big question for your board is, so what?

Is it that I want to stop people from cheating, or do I want people to learn to use AI really, really well? Because I can promise you when they get into the workplace, they're going to be using it. And to the other question there about ChatGPT and can you tell if people are using it, If there is a double dash in the sentence, then I can guarantee you they've used ChatGPT or another form of AI. But even if you say you can't use it, people have got like their phone, right?

So they're just going to go on their phone and shortcut it. If they've worked out how to shortcut it, and there's a whole thing around secret cyborgs and people effectively stealing time from the company because they're doing their job with AI, so there isn't a, on their own devices, you can't see whether they're using it or not. So, again, it's about having guidelines, it's about training

them, it's about embracing the technology and saying, how can I get my staff to champion this and use it really well to be more effective in their roles? Sorry, Alexey.

I was just going to add to the school board's question. Know, Part of our duties on school boards is to make sure that we're providing the education to prepare the students for the future. Exactly. Part of what our obligation for them is, how is AI going to be impacting their future, their university level, how are they going to need to use, how do you learn how to use it?

Absolutely agree that we've got to tackle the cheating question, but how are we teaching them how to think critically and not just rely on the output for the large language models if they're going to be using it anyway? We've got that shadow risk. So how do we put the guardrails, teach them how to use it responsibly and also our job is to shoot and prepare them for the future. Exactly.

Yeah.

Okay, our last slide. So we've covered the risks, expectations and governance challenges. So what matters most now is what do we actually do next? So here are five really quick next steps that you could look at for your board.

So map that AI usage, gain that visibility, take it as a standing agenda item for your boards, whether it's with your CIO or your chief risk officer, maybe do an external audit if it's appropriate for your board size, that forty two thousand and one gap analysis again, if appropriate. Number two, building that AI risk register. AI risks shouldn't just live in the ether. We've got to be documenting them, prioritising them and assigning them.

So treat AI like you would any other critical enterprise risk and who owns it.

Thirdly, develop that AI fluency. We don't need to become technologists, but we've got to have that fluidity. So start that capability building programme and include it in our governance off sites.

Do we need a board director on our board with credible AI expertise and have we had that AI specific board education sessions on a regular basis, part of our CPD hours now and policy implementation. Do we have clear principles and guidelines for how AI is deployed, approved and also monitored in clear English. It's got to be really simple for our teams to be able to understand and interact with. And lastly, governance.

We've actually got to be able to test the system, whether we can actually detect, manage and escalate AI risks as well. Can we establish clear thresholds for what warrants as well that visibility at a board level to make sure that we're empowered to act if something goes wrong with AI. So it's not just about tech, AI risk, it's all about stewardship and it starts now. So over to you, Sean, for more questions.

Sorry, I was just typing away furiously.

Do you want to just stick onto the final slide, Sean?

Sure.

So we didn't specifically think actually, Alexey talked about the resources.

So both the Institute of Directors in Australia and New Zealand have given guides, and I think, Sean, you'll be sending those out as part of the webinar notes. Yes. So, you are interested in sort of reading a bit more about governance frameworks, those two are both quite good.

Honestly, I shouldn't say this out loud, but I think I prefer the AICD one, even though I am in New Zealand, it's a little bit more structured.

And also, if you do want to understand a little bit more about AI, BoardPro has an on demand masterclass which you can download and I think it's about four and a half hours, and I'll just take you through everything you need to know about AI from prompt engineering, governance, risk management, the future, some of the new technologies that are coming out. This is a really broad scope of everything you need to know about AI, so a really good way of up skilling and getting in your CPDs, as Alexey said, around AI and governance.

And then, yeah, Sean, over to you for questions.

We're out of questions.

Oh, beautiful.

One more from Angela. I see that there's Are there any platforms that allow you to work in a confidential manner? Actually ChattyPT has a temporary turn off. It's basically like using your incognito session inside of chat. You can also, if like we always recommend, Helen and I have spoken about this a number of times, make sure you have a paid version and turn off the training to the algorithm which will allow the information to reside in your chat. However, we don't recommend putting confidential information out there, but there's a new temporary button inside of ChatGPT which will allow you to work basically incognito.

Interesting.

So I highly recommend you all reaching out to Helen and Alexei if you have any questions around AI for your organization.

Please feel free to even connect with either them on LinkedIn. I'm sure they'll look forward to your connection there.

So you'll receive an email from me. It'll be tomorrow afternoon now, early afternoon, which will include a copy of the recording, the transcript and the presentation slides plus those articles that Helen talked about. Just as you leave the webinar, don't forget to complete our one minute survey. Go in the draw for our hamper.

I'll announce the winner for that tomorrow as well. So thanks again everybody for your attendance today. I hope you enjoyed the session. Thank you Helen and Alexei for your great conversation.

Look forward to seeing you all everybody at our next webinar. Have a great day.