Webinar Transcript

Unseen and unchecked - The real AI risks facing your boards

So welcome everybody. Welcome to our webinar today titled Unseen and The Real AI Risks Facing Your Board. Our webinar team today is Helen Van Orton and Alexei O'Brien. My name is Sean McDonald and I am your host and moderator in the background for the next forty five odd minutes.

Firstly, thank you so much for attending today. We always appreciate the efforts 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 as against chat. It just enables us to keep a track of the questions as they're coming through.

And we'll try and get through as many of the questions as we have time for. And finally, if you stay through till the end, which of course we hope you will do, we have a really short one minute survey at the end of the webinar for you to consider. Your feedback really helps us bring relevant content to you week after week and enables us position the wealth of expert presenters for you. So, please take a minute to complete the survey as you leave the webinar today.

Now, for those not too familiar with BoardPro, we are a board software provider, sometimes called the Board Portal. And we serve just over thirty five thousand users around the world, and we're represented in about thirty four odd countries these days.

And we enable organisations to prepare for and run their board meetings more efficiently and effectively with our clever software, with less time and deliver more impact and value for the organisation.

And 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, but especially those organisations with resource constraints.

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

And these webinars that we host every week are also a great way of accessing key governance knowledge without the time commitment and costs associated with in person events.

So for the next forty odd minutes, just relax, listen and add to the discussion by asking as many questions as you would like.

The full recording of the webinar along with the slide deck and other resources will be sent to you twenty four hours after our session today. So that will be tomorrow.

So let me have my wonderful panel introduce themselves starting with you, Helen.

Everyone. I'm Helen Van Oten. I'm an experienced professional director and chair. Sit on boards like the Cooperative Bank and Centrix, the Credit Reference Bureau. I'm also the CEO and founder of Directly, which is a company that specializes in AI governance training, exec coaching, and keynote speaking, targeted at boards, directors, and senior executives. And really that's about empowering boards and execs to lead with confidence in the AI era and really understand what AI means in your governance journey. And I am joined by the fabulous, O'Brien, who I will hand over to, to introduce herself.

Thanks so much, Helen. And thanks for having me, Sean. My name is Alexa O'Brien, and I'm the director of Leadership Academy dot ai, and I work with boards, executive teams and businesses across Australia to build AI fluency, moving from curiosity about AI to confident governed use. My background is commercial.

I spent the bulk of my career in retail and financial executive roles at Lululemon and Rip Curlton, I'm a few. So I come at this from a perspective of someone who sat in the rooms where these decisions are getting made. More recently I've worked with businesses like Pillow Talk, Entain, Harris Farm Markets and St Kilda Football Club on building AI enabled cultures and getting real work done. I'm a current board director as well and a graduate of the AICD.

Looking forward to having this conversation with you today, Helen.

Beautiful.

Thanks, Sean. So some of you would have hopefully been to or followed up and watched the webinar that Lexi and I did last year. We did a session on AI board risk.

This year, the landscape has really moved and so has the urgency around it. And so today we're gonna talk about some of the risks that are already sitting inside your organization. Some of which you're gonna be able to see, some of which you probably can't. So we'll just start on this slide with where things are standing right now. So that stat I think at the on the left hand side is a little bit concerning so eighty seven percent of workers are already using AI at least weekly.

More concerningly, nearly half of them are uploading sensitive company data into an AI tool without anyone's approval.

And two thirds of boards have got limited knowledge or oversight of what is going on. And those are global numbers, but that is not just a global picture. Like a recent study showed that sixty eight percent of Australian workers are already using AI. So that's two and three. And in New Zealand, eighty seven percent of large organizations have adopted it.

But looking at the Australian number again, only a third of Aussie workers have actually had any formal training from their employer. So there's a really big gap here. And that gap between what's going on in the grand and what's being governed at the top is really significant.

And I'm gonna really start off with an uncomfortable truce.

Sorry about the background noise. I'll go on to mute as soon as I've finished. If your board has not set a framework for AI use, your people will have already built their own. They are gonna use whatever tools they're finding useful, however they see fit.

So the question isn't whether your board is going to lead the conversation, it's about whether they will lead it or then just discover afterwards that something has gone wrong. I'll hand over to you, Alexey, and quickly put myself on mute.

All good Helen. So I want to name a couple of things that have shifted materially since we spoke last year on AI risk. First, AI isn't something that people are now going out and searching and finding purposefully. It's actually already in tools that your organization is paying for.

So Copilot is in your Microsoft environment. Gemini is in your Google Workspace. Canva and Slack, just to name a few, all have AI capabilities. So these tools are actually starting to surface your organizational data through AI right now.

And in most cases, nobody's actually made a conscious or deliberate board level decision to actually turn that on. Number two, as Helen referenced in our first couple of slides, your people are actually ahead of your policies.

Almost half employers globally are using AI in ways that don't comply with a company policy if a policy actually even exists. And they're not trying to be malicious, they're actually really trying to be productive.

As an organization, we may not have set the guardrails for them yet.

And thirdly, the financial cost of getting this wrong is actually no longer hypothetical. So IBM found that ShadowAI, which I will touch on in a couple of slides time, is adding nearly seven hundred thousand dollars to breach costs And only seventeen percent of companies actually have the technical controls that are catching up to the capabilities that AI has for us, that will actually help prevent employees uploading that confidential data to public tools. So these three things together and now the new risk surface.

So the next slide please. So here's a full landscape. These are just nine and Helen and I were talking just before this call. Are definitely more than this, but we just wanted to surface these.

We've organized them in the order we'll cover them today. The first five we have touched on before, so we'll give it a bit of a light touch today. Please go back and check them out on the BoardPro site. And if you want to see, see a little bit more on hallucinations, bias, undocumented infrastructure and so on.

We will go deeper on the last four as we really believe that they're the most urgent, probably a little bit underestimated right now. Shadow AI, agentic, prompt injection and deep fake fraud. So we'll start with those first five and back over to you, Helen.

Thanks, Alexey. So as Alexey said, this is a bit of a super quick run through these on the basis that we have done them before. So if you haven't seen that previous webinar, please go back and watch it. We know that we're doing this super fast, but that is because that if this information is available elsewhere.

So these are five, risks, these first five that your board should probably already be tracking, and hopefully, you've heard about them. But each one, to be honest, deserves its own session. I sometimes go into boards and talk about each one of these for like probably fifteen, twenty minutes each. So the first one is a hallucination.

So this is where your AI is presenting complete false information with complete confidence.

I've literally the amount of times I've had AI telling me something with complete confidence, and and I know that it's wrong.

But your organization still carries the liability if your team is relying on this or indeed if one of your chat bots is relying on it. So Air Canada's chatbot a couple of years ago gave out incorrect bereavement fair information, and they tried to go to court and defend it saying, AI told me. Needless to say, the court found that that was not a legal defense.

So watch out for hallucinations. The second one is bias and discrimination. This is a really big one. So this is about your AI making unfair decisions about people at scale.

This has been really evidenced in recruitment, whether it's hiring, lending, etcetera. So Amazon's recruitment tool a few years ago actually downgraded any CVs, which had the word woman in it.

And it because it was based on a decade of, like, CVs that were predominantly males in the tech sector.

But if you are relying on a tool that is using bias, again, the board is liable whether the human or the algorithm has made the call.

The third one, and this is actually something that's probably feels really familiar to you. So this is around undocumented AI infrastructure. So it's probably a little bit underappreciated. But if you think back to the days when people were going into an Excel spreadsheet and they build this incredible macro in Excel and then they left, and nobody else knew how the macro worked, and actually the finance team were relying on the outputs of this folder, it's the same thing, but it's scaled up in in the AI space. So this is where your people are building workflows, automations, decisions, custom GPTs, where they're encoding all of the business logic, and there's no documentation. There's no IT asset registration.

And when that person leaves the organization, everyone's lost that information, and nobody can see it or maintain it. So this is not just a spreadsheet that one person understands. This is machine scale. So really, really important one to be thinking about.

The next one is your supply chain and vendor AI risk. Again, we kind of those organizations have we're hardly getting our heads and our grip around what's going on right now in our own organization, but actually AI is embedded in the products that you're buying as well, whether that's HR platforms or CRM or accounting software, not necessarily with disclosure. And going back to that recruitment piece in the US, there's now already been some case law. So a company called Workday was using their agent, and it was actually discriminating against people. And then all of the companies that had used Workday for their recruitment are actually being caught up in that class action as a part of that discrimination.

And then the final one, which is insurance coverage gaps. So as directors, we're all quite obviously concerned about our D and O insurance, and literally one of the boards that I'm on, we've just been back to our insurance companies to really understand this. So Barclay Insurance, which is one of the big US insurers has just introduced just before Christmas an absolute AI exclusion. So they're basically excluding all AI related claims from professional liability policies.

Other US companies are starting to follow. Obviously, that is gonna cascade down to Australia and New Zealand and up to Canada, because people on the call from Canada. So again, question for your board is, does your insurance actually cover an AI related incident?

You will probably assume that it does, but actually that answer is increasingly uncertain. So just make sure that you're having that conversation with your broker. We know that was quick. Don't worry. The next four, we will go into more deeply and more slowly as these are probably the most urgent ones right now. So I'm going to hand back to Alexey.

Yeah, just a couple of other things on that slide. The hallucinations to bring it a little bit closer to home in Australia, Deloitte last year, big four firm had to actually refund the federal government because they had a paper that was on a government website for actually three months before an academic flag that a lot of the quotes were actually AI fabricated court quotes. And that dependency question also, undocumented AI infrastructure, there's a dependency question buried in there. So I am asking organizations that are going deep on using AI as an operating model.

What if it got switched off tomorrow? What are we going to do? What breaks? So it's not just the infrastructure is undocumented.

It's that nobody's mapped out what we're actually now dependent on. So that's a continuity risk hiding inside an innovation story as well.

So next slide. Awesome. So I'm going to do a deeper dive on ShadowAI. You probably have heard of ShadowAI.

So this is where people are using tools often free publicly available tools without the organisation's approval, visibility or controls. Again, the data is stark. So more than three quarters of AI users are pasting data into these tools. Eighty two percent of those paste them from personal accounts, not enterprise managed platforms.

In Australia, only thirty one percent of businesses are actively focused on implementing AI governance policies, and even enterprise tools in the shadow AI space aren't immune. So Microsoft earlier this year confirmed that Copilot had actually been reading emails marked confidential. So bypassing the exact data loss prevention policies that organisations had put in place to stop that. Governments are also starting to take notice banning some of these uses.

What makes it hard to govern? ShadowAI actually doesn't look like a breach. It looks like your best people trying to do their best work. The financial analyst building a reconciliation dashboard, pasting ledger data into ChatGPT or your executive using a browser extension, summarizing every meeting, including maybe the ones that they shouldn't be recording.

They're not intentionally going rogue. They're trying to add value, but none of it is actually sitting inside our governance frameworks. Nobody's vetted the tools. Nobody also knows necessarily where that data went.

So if we asked our boards tomorrow, what AI is our organization relying on? Would we actually be able to answer that? So imagine this happening with our board pack, our M and A strategy or our legal advice. So that's the kind of risk that we're actually talking about with ShadowAI.

Helen, over to you.

Yeah. Just one a couple of other points on that ShadowAI as well is one, like even looking at the notes today, there's a lot of you have jumped on using Fireflies or Otter and they're great tools. But two thoughts around that. One is if you are recording meetings that are sitting outside the the covenants frameworks, there's a there's a whole piece around confidentiality and smoking guns and that data being available if anything happens later on in court, which is particularly relevant when you're recording a board meeting, if you haven't thought that through. And there is a webinar that we did last year on that.

But the other one as well is what's happening to that data? Where is that data sitting? If it's getting, if that company is breached has curious your data and particularly depending on the sector that you're in there's a whole load around privacy concerns if you're actually talking about client data in there. Have you thought through the full protections that you've got in place.

So one I had a conversation with a client the other day, and they were saying, oh, well, we've got all of our machines locked down, so nobody can actually use their laptops. And I'm like, but AI is multimodal. So there's nothing to stop somebody taking a photo of the screen and then putting it into their own GPT or whatever it is. So it's a it's a real people led conversation, this as well as a policy led one. Sorry. Onto the next one.

So I just also want to talk a little bit about Agentic AI. Now that's probably a phrase that you've started to hear in the press. It's, and people most boards are probably not really consciously registering yet, but it is not a future problem. It is definitely here.

So up until recently, AI has been quite reactive. So you ask a question and you get an answer. A Genetic AI is really different, and this is the one that should if you have not got any thinking about it in your board space, it's the one that should be really scaring you.

So a JentiKi will act. It will execute a task. It will trigger a workflow.

Whatever parameters that someone in your team has set, it is going to be making a decision without necessarily a human reviewing the action. So that is the machine at three o'clock in the morning running something and creating decisions, creating actions off the back of it.

So that might be your CRM that's responding to a customer inquiry or your accounting software or categorizing invoices, scheduling, making diary decisions. It's not necessarily dramatic deployments. It can be like quite quiet and little incremental ones that are just switched on by default.

But it's it's starting to spread through organizations. So Gartner did a study at the end of last year, and they're projecting that by twenty twenty nine, which is only three years away, half of all knowledge workers are gonna be deploying AI agents on demand. And that's not IT teams. This is frontline people in finance, operations, marketing, customer experience, etcetera.

Interestingly, you sort of look at the big companies. So Shopify CEO last year said that teams had to prove that a job can't be done by AI before they're allowed to hire a human. So there's a real shift in that future of work piece.

It's not experimental. It's becoming that operational expectation.

But from a governance perspective, that's actually surfacing some really new questions. So who authorized the AI to take the action? What are the boundaries? What happens when it acts on bad data? Because if it's hallucinated earlier in its process, it's, you know, it's gonna be changing. So it's not experimental.

Where's your audit trail? What's happening?

And, you know, we just popped an example up on this on the slide last sorry, on the slide, which is an an Amazon experience from last year, their own coding agent actually decided it was going to delete and recreate a live production environment, which actually led to a thirteen hour outage for AWS.

There was no human checkpoint. It had been given permissions.

And the systems are doing exactly what we tell them to do. We just might not have meant exactly what we wanted to do. And from a governance perspective, how are we putting in the

frameworks and the guardrails and the guidelines around AgenTik AI, as well as AI more broadly. Alexey, I'm sure you've got a couple more comments to add on this one.

Yes. I mean, I'm seeing it live in action inside of organizations, you know, playing with it, you know, having people actually take action. And, you know, I think it's really important that the comment that you made around people actually understanding that you're still prompting the agent. And so it's still acting on your instructions and what you want the objective to be.

And unless you're super clear on where to stop and where you want to have that, human in the loop conversation, it will follow your instructions, and what you've inadvertently asked it to actually create and do. So we've got to really make sure that we're locking down, you know, where we want it to stop, what permissions, what access that we're giving it as well all the way through, because there is so much capability. I keep joking with my teams, you know, with great power comes great responsibility. With this great capability comes a real opportunity and risk that we need to keep a hold of.

Right.

We have a question that's coming from Simon here, which I'll read out.

Simon says AI is moving so fast that enterprise subscriptions are not keeping up, meaning new apps need to keep pace with competitors and faster pace of work. What's the answer to this?

And Simon, you're really right. I think a lot of people in who work in the AI world find we don't tend to use Copilot, which is what a lot of organizations have because it it is not necessarily keeping pace. They are releasing a lot of updates to that.

The challenge is to make sure that you understand what you're signing up for, and you make sure that your people in your organization are only using the tools that you have tested and checked and sanctioned and authorized and that that, like the most important thing is like even if someone is using an AI tool so for example like before the call we were just talking about gamma as a presentation creator, you know, could use gamma in your organization if it I'm not not suggesting that you did, but if someone was using that in a way that would it was just putting in generic data and information and creating the outline and then bringing it back into PowerPoint in your own environment and then putting the company data in at that point, that's actually a reasonably safe way to use AI.

But going and putting your company information into Gamma is not safe. So it's again, it's around the guidelines that you're giving to your team to say, don't use the unauthorized tool, but knowing that somebody will, if you were going to, these are the things you need to think about. Never put the company data into it. Make sure that you're not just quietly doing this on your laptop, that you've spoken to someone in IT so you know that those pieces are in place. Alexa, you've probably got another couple of examples from your organizations, but it's about visibility and how people are using it are probably two of the most critical elements there.

Yeah, I think also being in the conversation with your teams as well, you know, there's an arms race back to Simon's question going on, you know, like last week, Claude released so many new products, you know, they're predicting that open AI and Codex is going to, do, you know, so many more things identically because you're going be able to run it on your computer at the same time. Whereas Claude code, you've got to, you've got to give it control and stand back. I think having the conversations with your team, you know, having, AI champions in each department that they can actually start to surface, you know, where are they using it?

What tools are they discovering? Making sure that we're looping back into, you know, having it part of that AI, champion team so that we're aware of what teams are using it, why they want to be testing these new tools, and bringing them into the, to the fold to make sure that we're, covering off on those things because people, your teams, they're out there, they're watching YouTube, they're getting the subscriptions on notifications, on capability. And it is like, it's hard for Helen and I to keep up, you know, it's changing so quickly and there's so much of it. But it's just, you know, making sure that we're staying in the conversation, continuing to, you know, have the team surface up what they're finding useful.

Why is that part of our ecosystem and have we got eyes on what tools and what data people are using?

All right, so we'll go to the next week's risk. We're going to talk about prompt injection. So what happens when AI is actually following instructions that you didn't even give it? So that's our next risk.

It's a newer territory one. It's called prompt injection. Sounds a little technical, but it's really simple. So most of us understand social engineering.

So someone might manipulate a person into giving up information or access that they shouldn't prompt injection is the same kind of idea, but the target is the AI, not the person. And here's how it works. So someone might hide an instruction inside a document, an email, a data source, and it might be invisible to the human reader, but when your AI is actually processing the content, summarizing it, and generally when you're uploading documents or using extensions, it will read the entire document. And it might pick up that hidden instruction that you don't know is there and follow it.

And it could mean leaking data, ignoring a security policy, actually deleting data, or producing a misleading summary. So there's a case on the screen. The Ecolease case came out in twenty twenty five. It was a normal email and an attacker sent through someone's Outlook and Copilot allowed the email to have read that because it was AI reading that and it silently exported some sensitive corporate data through trusted URLs from Microsoft.

There wasn't a click, there wasn't a phishing link, there was no download, the data just left. So the defenses that Microsoft had in place at that time, their prompt injection classifiers were actually bypassed by something almost embarrassingly simple. The attacker wasn't in code. It just wrote the email as if it was talking to a human.

And that's how thin the line is. So this isn't just an email risk inside a Copilot either. There's a named vulnerability called CellShock where researchers are now demonstrating that Claude for Excel can be tricked into writing formulas that will, exfiltrate your financial model. So, really making sure your teams are also trained in what these risks are because they're the ones using the data.

So if they have called for Excel, spreadsheets, which can be hundreds of thousands of lines long, could be hiding white on white, instructions that will actually take action. So the point is anywhere AI can read content that it didn't write, spreadsheets, slides, PDFs, the attack surface is there. And the sobering part is, OpenAI has publicly acknowledged that prompt injection is unlikely to ever fully be solved. We have more and more, access more and more, patches, but even, the UK cyber security agency said that it's not necessarily a bug, that can get patched.

It's a structural characteristic. So we've got to be aware of what, what files that we're using, what data we're giving our AI access to, and making sure that we've got eyes on what we're actually using to make sure that, and the teams know don't use spreadsheets and files with these extensions in play. If you don't trust them and you don't know those sources as well.

How about you, Helen?

Yeah. And I think another thing that increasingly, if you've got people in your team who are quite sophisticated AI users, you can go and download skills. So for example, if you're a Claude Copilot user, there's, like, libraries of skills that you can go and just download and import into Claude, which are, you know, I'm your marketing SEO expert or I'm your customer experience complaint handling expert. And, you know, they're written as if they're the most amazing thing to transform your workforce.

Now Claude has, in partnership with GitHub, has a whole load of checked ones, but you can get these like they're all over TikTok or, Instagram or anywhere you want to look. If your team just goes and gets a skill from somewhere that is not an authorized place, there could well be a prompt injection in the skill, which they've just brought unknowingly into your organization. So again, prompt injection is it's quite a scary risk surface, but one that you definitely need to be aware of. Okay.

Jumping on to the, the fourth one.

Sean?

Okay. So this one tends to get people's attention. So this is a case from, like, eighteen months ago. So it's a little bit old, but it's actually really relevant because of what's happened in the last week.

So twenty twenty four, UK engineering company, a Hong Kong employee jumps onto a call. On the screen, you've got a numb the CFO, you've got a number of the exec team sitting on the call, and the CFO asks Skye on the call to make a transfer. And two hundred million Hong Kong

dollars, which is twenty million pounds, is transferred by this person. The only thing is is everybody on that call was an AI generated deepfake.

There was not a single real person on the call other than the person who was actually doing the transfer. It sounds a bit like a thing out of a Hollywood movie, but this is actually a real case confirmed by Hong Kong police, etcetera.

So the numbers around deepfake fraud are really staggering. So attacks on businesses are surging. Voice cloning frauds up like nearly seven hundred percent. Attacks on businesses have surged like three thousand percent over the last couple of years.

It, as a director, this is the thing that should scare you. Because if you're anywhere on video or your voice is recorded anywhere, it takes about three seconds of audio to create an eighty five percent voice match, whether that's a conference presentation or webinar recording. So Alexei and I have got no chance because do too much content. Or even a voice mail.

But this is where this gets really surreal. So in the last week, Mark Zuckerberg, who's obviously the CEO of Meta, has announced that Meta is actually building a photorealistic AI clone of himself, which is trained on not only his voice, but his mannerisms and all of the strategy stuff and all the things that are really important so that he can hold one on one meetings with seventy five thousand employees on his path. Now he's doing that really deliberately and transparently, but this is the same technology that's being used for deepfake. Right?

So if the CEO can of Meta can build this convincing replica of themselves, so can a bad actor.

And the question for you as a director is, has your organization updated its controls so that in a world where authority can be convincingly imitated, you can tell the difference or your team can tell the difference between sanctioned leader and a malicious one. And have your financial controls been updated to account for the fact that voice, video, etcetera, can no longer be treated as reliable authentication. You know, it used to be we'd get an email into our inbox saying it was from the CEO going, please, can you transfer some money or whatever it was. Now it's a lot stronger. It's on a video call. It could be the entire exec team trained in their voice, their mannerisms, having a real conversation. And if one of your team has leaked data about what's going on in the organization through ShadowAI, they've actually got really current information about a project that's going on with people who sound, look, feel authentic.

You need to be thinking about what controls that you've got in place. K. I'm really conscious of time, so we are going to jump on to the next one. So as boards, we kind of we'll discover that staff are using the unapproved tools, and our instinct is go, it's a compliance failure.

They're not pushing the they're not following the rules. Generally, as we said earlier, it's not about your people being reckless. They are trying their best to be productive. They found something that works and the they feel that there's a gap between what they can do is their best work and what they're able to do.

So making sure that you've got a Gen AI policy in place. And as we said earlier about talking to that and making sure that's not just a policy that sits from the boardroom, it's actually all through the organization and that managers are having their conversations with their people around how to think about shadow AI and why it's important. Like Australia right now has got one of the lowest rates of Gen AI policy adoption, although I think that might have changed slightly. Think it's updated just in the last week, hasn't it?

Alexei KPMG has done a new study on that. I take that back.

So don't worry about that. It's not the lowest in the world. But the challenge is is that the the governance frameworks are running way behind what people are wanting, and it goes back to that question earlier. So as a board, it's making sure that you've set the guide rails. It's made sure that you are thinking about whether you know, it used to be the BYOD. It used to be cloud adoption.

We're at that point with AI right now. So we're gonna rather than just going, hey. This is a problem. We're gonna super briefly talk about what we need to think about and do it, but I'm conscious that we've only got a few minutes left. So, Alexei, back to you.

So what does good actually look like?

It's five things. None of them require technical expertise for your board. So number one, get AI onto your risk matrix. If it isn't, not as a single line item, but across the dimensions that matter.

So data security, governance, strategic alignment as well. AI services, your strategy, not the other way around. If AI isn't on your risk register yet, Number three, have AI as a standing board agenda, not a deep dive one off or a crisis response, but a structured regular conversation. Number four, adding AI literacy to your board skills matrix.

Doesn't have to be a technical deep dive, but enough to be asking the right questions and setting direction. I think I would add in here, actually making sure that your PD and your AI understanding your knowledge and education on AI is sufficient that you can ask those questions.

Number five, honest visibility on what AI tools are actually in use across the organisation. Can't govern what we can't see.

So we do these five things and we're ahead of the vast majority of boards out there. It requires that governance leadership.

Back to you, Helen.

Beautiful. Okay. So we've said that getting onto the risk matrix is great. So when you're sitting in your risk committee, think about it.

Yeah. This is very much how your risk committee would normally be thinking. Okay. So risk appetites.

What is your risk appetite for AI?

Have you actually had a conversation around that?

Because if it isn't reflected in your risk appetite, then you actually haven't made a conscious decision about it.

So it's not just AI is a risk because some parts of AI, there's fantastic opportunity. And for some organizations, it it's good to have a sort of a more optimistic and open, you know, rather than closing down the opportunities, but you need to understand as a board where you sit on that. That second piece is visibility. So that's so important around how you're tracking that AI use within your organization. And as we said earlier about shadow AI, it's not just about the tools that your IT team have approved, but it's actually ones that people are using on their own devices, their personal accounts.

The third one is strategy. And just super quickly, like, AI is a risk conversation. Absolutely. But it's not only a risk conversation.

It affects your business model, your competitive position, your strategic plan. If you're not talking about it as part of your strategy and the opportunity that it brings you, you really need to be. You're really not governing well if you're not thinking about the opportunity around AI. And then the final piece is that assessment and that's kind of how we answered that previous question, which is what are the tools and processes that you're using to differentiate between high and low risk AI applications?

Because it doesn't all carry the same risk. And again, your governance framework needs to reflect that. So those are some of the questions that boards need to be thinking about to move from awareness to action. And you really need to be working on those and mapping those back to the conversation that we've had today.

So back to Alexey.

Thanks, Helen. So before I just jump into this slide, Sonia has asked a really great question in the chat. You know, what happened? Our management pushed back on, questions about AI saying it's too operational.

So how do we respond to that? So, AI, you know, the core thing is that AI isn't just an operational question. It's as we've been discussing today, it's strategic risk, it's strategy. And it's a governance question, which makes it squarely a board matter.

So when management's saying that they're too operational, we've got to make sure that we've got a clear strategy and ask the questions. We're not asking how you're configuring the tools or necessarily using them. We're asking, have we got appropriate guardrails around the

technology that touches on data privacy, on cyber risk, IP protection, regulatory governance, as well as workforce strategy. Helen touched on that before as well.

That's really governance. That's the role of the board as well, and it's not operational.

So there's a couple of thoughts on that, for you, Sonia, hopefully that helps.

So before we go to further questions, Helen and I are both offering AI board workshops, thirty to sixty minutes. They fit around our busy board agendas. They're purpose built for directors who really need to govern AI, not be in the weeds, with confidence and not just curiosity. You can't govern it unless we understand it, have a working use case of understanding it.

So these, are targeted to the risks that really matter most to your sector designed for also non technical directors. So that's something your board would benefit from. We'd love to have a conversation. So as you can probably tell from the accents, I'm based in Australia, Helen's in New Zealand.

Please get in touch. Our details are there on the screen. Back to you, Helen.

Okay. I think we've got time for a couple of questions. So thank you everyone for jumping on the call today and we will go to the questions. I can only see one in the q and a at the moment.

Realistic status of regulation, particularly in view of the current USA regulatory stance.

Yeah. Great question.

Both the New Zealand and Australian government, sorry, Canada, I don't know what the status is over there, have both put in guidance around AI rather than any legislation. We did actually have a slide in this pack on regulation originally.

But it really at the moment, it's things like the Companies Act and the Privacy Act. There's not any more robust legislation than that. I think a really important thing to remember though is depending on the size of your company and where it's operating, if you fall into, like as in you are operating in the European Union, you are still obligated to comply with the EU AI act. You have to say that really slowly.

And the fines for the EU AI act, if you get it wrong, are up to seven percent of your company's global revenue or thirty five million euros. So they are really, really hefty fines if you are not compliant with that act. And that's particularly around fairness and bias. Those larger fines come into play. So, yeah.

The other thing I'd add to that Helen is, you know, we still have data privacy that we need to comply with too. So we've got to make sure that we are compliant with and there are some updates to the Australian Data Privacy Act that come into play in December, twenty twenty six. So check those out too. We need to make sure that we're understanding what's coming from a data privacy perspective, what we're using in our large language models and AI tools. It's

generally probably an AI risk, data privacy risk first, if we're using our, our company's data or customers or employees data, in AI.

Yep. And there's there's only one piece of case law that even refers to AI across Australia and New Zealand at the moment, which was I think about three weeks ago when it was basically saying that boards and directors can't rely on an AI summary of a board pack for you know, they need to understand the company better. So, yep, it it will evolve slowly. We've just started using Copilot.

Can you elaborate on your concerns? Look, I don't have any concerns about Copilot itself. It is it's a robust tool. It's got all of the enterprise security.

Like, it's a it's a good tool.

It's just it's not necessarily as advanced a model as some of the other models that are out there. So depending on what you're used to, it's it it can't necessarily act in an environment in the same way that some of the others were are in terms of its connectors, etcetera.

So it's not a concern at all. Like, it's you know, if you go back six months and said, wow. If we had something like this available, it would be amazing.

If you're working at the cutting edge of AI, it doesn't feel quite so amazing. Is that probably fair, Alexey?

Yes, absolutely. I think, I mean, is based on Claude, but just the responses just not as robust, not as detailed, but it's also better than not having access to it all. I know that a number of organisations don't allow any other type of large language tool. So they've shut that down. So, is better than not using anything at all.

And to that final question from Joanna, it's rather than what what to use, what not to use. What not to use is, like, make sure you're using a comp a tool that has an enterprise level of security around it. So any of the like Anthropic Claude, OpenAI, Google Gemini, all of those have got all of those have got like an enterprise environment. The worst thing you can do is use a free tool.

Because if you are using a free tool, you are the product, not them. And they're taking all of your data and they're using it to train the model. So make sure that you are paying for whatever you use and that you are using it within a plan that has got the model not trained on your data, that is got that enterprise security sitting around it. That the most important guidance we can give you around tools to use.

Basic Claude, do you mean a free Claude reasonable?

Think Helen, you would back me up in saying, that enterprise level or at least certainly a paid version of any of the large language models provides that additional level of security.

Enterprise version would be even better because it's all about the data. If we're putting data in, it's where is that being sent to if it's being offshored we're certainly at risk from a data privacy perspective. If we have an Excel spreadsheet built by Claude is it likely to have prompt injection? I wouldn't have thought so. I don't know the answer to question specifically.

Again, I would really just if you wanted to be really safe, is what spreadsheets are you actually building and using and make sure they're internal.

Yep. One other thought on the what not to use is just to be aware of if you are using a a Chinese AI like DeepSeek, just check their terms and conditions because currently their terms and conditions say they can share any data that you upload to wherever they want to.

So, just be aware of the terms and conditions that you're signing up to. Don't just go, Yay, that's great. Sean, back to you.

Thank you. Thank you. So, please feel free to connect with our presenters today on LinkedIn. I'm sure they'll look forward to your connection, Helen or Alexei. And also on the survey. Hopefully on the survey, I've got your names inserted in there. My apologies if I haven't.

Our webinar program is published on our website, which you will see listed here. We've got a great series coming up for you over the next couple of months. Tomorrow, we have a session on to pay or not pay your directors, which is proving to be really popular. We've got over nine hundred registrations for that particular webinar. So, that's a really hot topic. So, take a look on our webpage for the upcoming webinars.

You'll receive an email from me tomorrow, which will include a full recording of today's webinar, along with the transcript and the presentation slides that Helen and Alexei have put together.

They'll also be hosted on the webinar library of our website over the next forty eight hours. And of course, if you're considering board management software for your organization, we would love to hear from you, of course. Better still, why not try our free thirty day trial? It's simple and really straightforward with no credit card required, and it's really easy to get started.

So thanks again, everybody, for your attendance. I hope you enjoyed the session with Alexei and Helen. I know I did. I look forward to seeing you all at our next webinar. Everybody, don't forget our short one minute survey as you leave. It really helps us craft this great program for you.

So, a great day everybody.