



AndersonRisk The AI Risk Dialogue

A thematic, board-level review of AI uncertainty, dependency and control

Boards are approving AI tools, partnerships and investments faster than most of them can answer two simple questions: who really understands the risk, and what happens if a key supplier walks away? The issue isn't whether to embrace AI: you have to. It's whether you're taking on exposures faster than you and your people can understand them, challenge them and keep control of them.

The problem with most AI governance

Most AI risk is being managed with tools built for other things. A policy. A supplier inventory. A risk register. All of them have their uses, and all of them can exist without ever answering those two questions. None of them tells you where authority really sits, or which dependencies are becoming hard to undo.

The harder issue to grapple with is that AI risk isn't one thing. Some of it you already feel directly: errors, bad customer interactions, fraud. Some of it you can settle with evidence: model performance, bias, security, data quality. And some of it is genuinely unresolved: nobody yet knows, for instance, what a shift to smaller networks of specialist AI models would do to cost, concentration or who would end up holding real power. Force all three risk types into one risk score and you get the appearance of control, not the reality of it.

Three lenses, not one framework

What do we actually know? For each significant use of AI, I separate what you've experienced directly, what evidence can support, and what's still genuinely uncertain. Then I ask the only question that matters: what evidence would change your decision?

Does the right dialogue happen? Governance lives in dialogue, not in policies. Does your risk team mean the same thing as your technology team when they say "controlled"? Does bad news reach the people who can respond to it, or does it get smoothed out on the way up? Does challenge survive when there's a commercial imperative to just say "yes"?

What could you not undo? AI capability increasingly sits outside your organisation, in models, cloud platforms, data providers, specialist suppliers. I map where the dependency sits and test a simple question: if this supplier changed its terms tomorrow, could you cope? The aim isn't zero dependency; that's not realistic. It's dependency you could still walk away from if you had to, or at least a route towards that goal.

How the review runs

Three to four weeks, built around your board's actual concerns rather than a generic checklist.

Four phases of work:

1. **Scoping:** I start with a conversation with the Chair, CEO, Risk Committee Chair or CRO to identify the decisions that matter.
2. **Documentation review:** I then review what's already there: strategy papers, policies, supplier contracts, risk assessments, board minutes.

3. **Interviews:** Eight to twelve confidential conversations with directors, executives and specialists. This is usually where the real picture emerges, often a different one from the papers.
4. **Feedback:** It closes with a facilitated board or executive workshop to test the findings and agree what happens next.

What you're left with

What you get coming out of the review:

Output	What it gives you
AI uncertainty map	What's known, what's assumed, what's still unresolved
Risk-dialogue map	Where information and challenge move, and where they get stuck
Dependency assessment	What you rely on, and what you could and couldn't replace
Board decision framework	What needs board sign-off, what's delegated, and when to escalate
Governed learning agenda	What to test before going further, and what would stop you
90-day action plan	Practical next steps, prioritised

Who this is for

Any board moving from AI experiments to AI that actually matters, especially if you're concerned about supplier concentration, agentic AI, or losing capability you can't easily rebuild.

This isn't model testing, cybersecurity assurance or legal advice; those need their own specialists, and where they're needed I'll say so. This sits above that: whether those specialists' answers reach the board, and whether the board knows what to do with them.

Why me

I spent over twenty years, including eight as a partner, at PwC, and later chaired the Institute of Risk Management. Since then I've chaired board risk committees and sat as a non-executive director across banking, payments and critical financial infrastructure, including through periods where the theory of good governance met the reality of a live board under real pressure.

My work has always been about risk appetite, risk culture and the conversations that decide whether difficult information reaches the people who can act on it. This review brings that experience to AI: treating it as a governance and behavioural question first, not just a technology one.

I do this work myself. There is no team of junior consultants: the person you speak to first is the person who conducts the interviews and stands in front of your board.

Getting started

If this sounds like a live issue for your board, the best next step is a short conversation, no commitment on either side, just a chance to work out whether it's the right fit for where you are. Get in touch to arrange a time for a preliminary discussion:

Email rca@andersonrisk.com

Phone 07807 780284

Richard Anderson

Founder, AndersonRisk

More about AndersonRisk: andersonrisk.com

AI Risk Reading, my fortnightly newsletter: [LinkedIn](#) · [Substack](#)