

AndersonRisk

The Straight Line and the Boardroom

An executive briefing for board chairs and directors

Richard Anderson · October 2026

Two clocks

The money behind artificial intelligence runs on one clock and the technology runs on another. Data centres, chips and power contracts are financed over fifteen to twenty years. The competitive frontier they serve moves in months.

Most forecasts settle the difference by drawing a straight line: namely that tomorrow's AI will be today's, only larger, and bought from the same few suppliers. That may prove right. It's still an assumption, and few boards have discussed it as one.

What two prices assume

I've run a shareholder value model in reverse on the two largest AI listings of the year. I didn't ask whether either price is right. I asked what must be true for it to be right.¹


Anthropic is reported to be aiming at \$2 trillion. That needs operating margins of about 45%, revenue of about \$1.4 trillion by 2038, and no change in the kind of AI the world buys until the 2040s. Margins like that belong to a supplier its customers can't replace. The company also carries \$518bn of future cloud and compute obligations, more than a hundred times last year's revenue.

SpaceX, which now owns xAI, is valued at about \$2.3 trillion. The businesses I can model account for about \$250bn. The rest is a payment for businesses that don't yet exist. Its fastest growth comes from renting compute to four customers, the largest of which is Anthropic. Either side can end that contract on 90 days' notice.

Neither price is impossible. Each needs today's kind of AI to keep its shape for fifteen to twenty years. And the two depend on each other: one company's rent is the other's revenue.

You may never own either share. If your organisation builds on suppliers like these, the valuations are still about you. Margins of that kind, held for that long, come from customers who keep paying and find it hard to leave.

Two paths, one question



¹ Anthropic figures come from press reports of a leaked draft prospectus and may change. SpaceX figures come from its SEC filings and market data at 6 October 2026. Nothing here is investment advice.

I don't know which path we're on, and nor does anyone else. A board doesn't need to know. Both paths lead to the same question: could we change course if we had to?

There's an upside in that question. If AI turns towards smaller, cheaper, more specialised models, the organisations able to switch will get better tools for less. I call this sovereignty through substitutability.

Three questions for your next board meeting

1. **What are we assuming?** Which of our AI plans rely on today's suppliers, prices and architecture lasting as long as the investment? Which of those assumptions can we test, and which can nobody know?
2. **Do we mean the same thing?** When we say our appetite for AI risk is low, is each director picturing the same risk? Would a doubt about a supplier reach someone with authority to act?
3. **Could we change course?** Could we move provider and keep our data, our accumulated know-how and our capability? How long would it take, and what would it cost?

None of these is a technical question. Each belongs to the board.

Why me

I spent over twenty years, including eight as a partner, at PwC. After the financial crisis I wrote for the OECD on risk management and corporate governance. I then chaired the Institute of Risk Management and was principal author of its guidance on risk appetite and tolerance. Since then I've chaired board risk committees and sat as a non-executive director across banking, payments and critical financial infrastructure.

My work has always been about risk appetite, risk culture and whether difficult information reaches the people who can act on it. I do this work myself: the person you speak to first is the person who stands in front of your board.

Where this leads

These three questions open The AI Risk Dialogue, a thematic, board-level review of AI uncertainty, dependency and control. It begins with a single board session.

The reasoning sits in my essays at [Infinite Curiosity](#): *Anthropic: What does \$2 trillion assume?*, *SpaceX: what does \$2.3 trillion assume?*, the Straight Line Fallacy series, *Here Be Dragons* and *Here Be Aliens*. All of these can be found on my [website](#).

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](#)