The AI Blind Spot

Why the Smartest Technology Investment of the Decade Is Making Family Offices More Fragile

Dr. Ralph Welpe | Chairman & CEO, Mirai Capital Partners
April 2026 | Mirai Strategic Insights Series | Issue No. 3
ralph.welpe@miraicapitalpartners.com | www.miraicapitalpartners.com

A week ago, I had dinner in London with the CEO of a European single family office. Roughly $8 billion under management, spread across private equity, real estate, and a handful of operating businesses in logistics and healthcare. I have known him for years. Sharp, meticulous, exacting about detail.

Over the past six months he had thrown himself into what he called a full AI transformation. Portfolio analytics, automated report consolidation, due diligence review, an internal knowledge chatbot. He brought in a Co-CTO, hired a consultant, spent about $800,000 on the programme. I remember thinking at the time that he was ahead of most of his peers.

Then he told me what happened the Monday before our dinner. Someone on his accounting team received an email that appeared to come from the family patriarch. It referenced a real estate deal the family had been discussing internally. Right tone, right details. It told the accountant to wire EUR 2.4 million for a deposit on a property in Lisbon. She had worked there nine years. She sent the money.

It was a deepfake. A good one. The attacker had gathered publicly available material (LinkedIn, a conference speech the patriarch gave once, bits from the family foundation's annual report) and run it through a generative AI to produce something that read exactly like the old man. Four hours later the money was gone.

"Ralph," he said to me, "I spent $800,000 making us smarter. I did not spend ten minutes thinking about how the same technology could be used against us."

I have not been able to stop thinking about that dinner. Not because of the fraud, which is fixable. What bothers me is the assumption behind it. That buying a technology is the same thing as understanding it. I have spent two years watching family offices commit to AI. And the more I see, the more I think most of them have the sequence backwards. They are building capability without building comprehension. The gap between those two is where I expect the damage.

This memo is my attempt to work through that gap. What the numbers actually say, what the attacks tell us, what the research on human judgment is starting to show, and what a genuinely careful approach to AI adoption would look like in practice.

The Stampede and the Silence

I want to start with some numbers, because I think two stories are being told simultaneously, and almost nobody is reading them together.

The first story is exhilarating. According to a 2026 global study by Ocorian of 200 family office executives across 16 jurisdictions, 86% are now using AI technology in their operations.[1] That figure was 12% in Deloitte's 2024 survey of 354 single family offices, and 33% in BlackRock's 2025 survey of 175 offices overseeing more than $320 billion.[2][3] The acceleration is extraordinary. UBS reports that 69% expect to use AI for financial reporting and data visualisation within five years.[4]

I do not dispute the direction. AI is genuinely transformative and the efficiency gains in analytics, document review, reporting, and risk monitoring are already visible in the offices I work with.

But there is a second story, and it is told more quietly. In the same BlackRock survey, 61% of family offices said they "do not know where to start" when embedding AI in their workflows.[3] McKinsey's 2026 AI Trust Maturity Survey of approximately 500 organisations found that only about 30% have reached even a moderate level of maturity in AI governance and strategy, despite widespread deployment.[4]

And in February 2026, the US Treasury released a Financial Services AI Risk Management Framework containing 230 control objectives across governance, data, model lifecycle, third-party risk, and adversarial resilience, developed with more than 100 financial institutions. The framework exists because the gap between adoption and governance has become a systemic concern.[4]

Neither set of numbers surprises me on its own. What surprises me is that they are both true. 86% using AI. 61% who admit they do not know where to start. You cannot square those two figures. They describe an industry that has bought the tools but not the understanding.

I have been around long enough to recognise this pattern. It is the same dynamic I watched in the early 2000s with offshore structuring, in the 2010s with direct investing, and again with private credit in the early 2020s. The opportunity arrives first. The infrastructure to manage the risks arrives later, usually after the damage is done.

What No One Asked the Vendor

I have sat through a lot of AI vendor pitches over the past couple of years. The questions from the family office side are always the same. What does it do? What will it cost? Can we see it work? Not once, in any of these meetings, has anyone asked what I consider one of the key questions that matters: what new attack surface does this create?

I think about this question differently now than I did a year ago, and the reason is the data.

Let me lay out the data. Deloitte's 2024 cybersecurity report found 43% of family offices had experienced a cyberattack in the previous two years. Among offices above $1 billion, 62%.[4] GCHQ's cyber arm said in 2025 that AI will "almost certainly" make intrusion operations more effective and more frequent, and warned about a "digital divide" forming between those who keep up and those who do not.[4] That year the NCSC dealt with 204 nationally significant incidents, 129% more than the year before.[4] The FBI, in its 2025 report, created a dedicated AI fraud section for the first time: 22,364 complaints and $893 million in losses. Bear in mind that is only what people actually reported.[4]

These are not hypothetical risks. At Arup, back in January 2024, someone in finance was tricked into making 15 wire transfers totalling $25.6 million. Every other face on that video call was a deepfake, including the supposed CFO.[4] A few months later a Ferrari executive picked up the phone and heard what he thought was his CEO, Benedetto Vigna, Southern accent and all, asking him to execute an urgent currency hedge. He only caught it because he asked a personal question the voice could not handle.[4] When Omega Systems surveyed family offices in 2025, 83% said they were worried about deepfakes. Only 60% thought their people could actually spot one. For context, the industry average in financial services is 69%.[4]

What worries me most, though, is how fast the offensive side is moving. In March, Anthropic leaked internal documents about a model they had not yet released, called Mythos. Their own people described it internally as an "unprecedented cybersecurity risk."[4] The stock market noticed. Cybersecurity shares dropped on the news.[4] Since then Anthropic has confirmed Mythos is real and has given early access to Amazon, Apple, Cisco, JPMorgan Chase, and Nvidia through something called Project Glasswing, essentially asking them to find and fix their weaknesses before a broader launch. What makes Mythos different is that it works autonomously. It finds vulnerabilities in operating systems and browsers, strings together minor bugs into serious exploits, and writes the attack code itself.[4] The UK AI Security Institute assessed it as more capable at offensive cyber operations than anything that came before.[4] Anthropic's own words: "The fallout - for economies, public safety, and national security - could be severe."[4]

Sit with that for a moment if you run a family office. The AI behind your portfolio system can now be turned around to find the holes in the systems protecting you. Your analytics touch financial data. Your document tools process confidential deals. Your chatbot has years of institutional knowledge sitting inside it. Every one of those is a way in. Dave DeWalt, who ran FireEye and then McAfee, said something in April that I keep coming back to: "The gap between what the offense can do and the defense can do widened dramatically. We are entering an actually dark period."[4]

The man I had dinner with is not a careless person. He understands technology better than almost anyone in this space. But he had only ever thought about AI in one direction: how do we use it to improve? The other question, how does it change the ways we can be hurt, had simply never come up. I do not think his office is unusual in that regard.

The Problem of Vanishing Skills

I want to describe something I have been noticing in the offices I advise, and I am not sure most people have a name for it yet.

The medical profession does. The Lancet published a study last year of 19 experienced endoscopists at four centres.[4] After just 3 months using an AI polyp detection system, their ability to find polyps without the AI dropped. Adenoma detection rates went from 28.4% to 22.4%. Three months, and doctors who had spent years mastering this procedure were measurably worse at it. The researchers called it "de-skilling."

I am seeing echoes of this in finance. A 2024 study had 157 people make financial forecasts with AI assistance. The AI was wrong, over and over, and people kept following it anyway. Here is the part that got my attention: when the AI made mistakes, people's trust dropped less than it would have if a human adviser had made the same errors.[4] A separate piece of work from MIT and Stanford looked at accountants using AI tools. Junior staff followed the AI no matter what. The experienced ones were the only ones who pushed back when the system flagged low confidence.[4]

Why should a family office principal care about this? Because the thing that actually makes a good office good has never been the technology. It is the people. The CIO who reads a deal memo and can feel what is not there. The operations head who catches a pattern in the administrator's reports weeks before the data confirms it. The CEO who picks up the phone when the patriarch goes quiet because he knows, from experience, that silence means trouble. None of that comes from software. It comes from doing the job, year after year, in situations where nobody hands you the answer.

When you automate the reporting and the due diligence and the portfolio analysis, yes, you get faster. But you also take away the repetitions. The daily practice through which people build and maintain judgment. And there is something even worse, which researchers have started calling "never-skilling": people who train in AI-augmented environments and never build the underlying skill at all.[4] The British Journal of Psychiatry drew the aviation comparison earlier this year and I thought it was apt. Autopilot made flying dramatically safer. But when it fails at 35,000 feet in unusual conditions, pilots who have not maintained their manual skills have lost control of the aircraft.[4]

I do not think you need me to spell out the parallel. When AI fails in the middle of a market dislocation, or a fraud comes through a vector nobody planned for, or the vendor's model starts producing wrong answers in conditions it was not trained on, you reach for the human backstop. If that backstop has been allowed to atrophy, you will find out in exactly the worst possible moment.

> *The question is not whether AI makes your team faster. It almost certainly does. The question is whether it is also making them weaker, and whether you will only discover the answer when you need your people most.*

When the Model Changes and You Do Not Know

Let me describe something that illustrates a different dimension of this problem.

In mid-2025, Upstart Holdings, a publicly traded AI lending platform, deployed an update to its proprietary credit model, which management called "Model 22." For several months, executives cited the model's performance in raising revenue guidance. Then in November, they disclosed that Model 22 had "overreacted" to macroeconomic signals, reducing borrower approvals and conversion rates in ways that management had not anticipated. Revenue guidance was cut. The stock fell nearly 10% in a single day. By April 2026, investors had filed a class-action lawsuit alleging that the company had made materially misleading statements about the model's performance.[4]

The specifics of Upstart's situation are not my concern here. What interests me is the mechanism. A model was updated. Its behaviour changed. The people relying on it did not understand the change until the consequences were already unfolding. And these were sophisticated financial professionals at a company whose entire business is built on AI models.

Now consider a family office using AI-powered portfolio analytics, risk categorisation, or due diligence tools. The vendor updates the underlying model, and they are updated constantly. The risk categorisations shift. The anomaly detection thresholds move. The document summaries emphasise different things. Your office is making decisions based on a system whose underlying logic changes without your knowledge or consent. The Financial Stability Board flagged precisely this risk in November 2024, warning of "complexity and limited explainability of some AI methods" and the systemic danger of "third-party dependencies and service provider concentration" across the financial system.[4]

Most family offices are small. The median has fewer than 20 employees. They do not build AI tools in-house. They buy them from vendors, and each vendor relationship creates a dependency. Your data flows through their systems. Your workflows are built around their interfaces. Your institutional processes become encoded in their logic. When a vendor changes pricing, deprecates a feature, or suffers an outage, there is no fallback. Switching requires rebuilding workflows from scratch.

In a family office managing multi-generational wealth, where consistency and institutional memory are essential, this silent drift is corrosive. You are not just outsourcing tasks. You are outsourcing the logic by which you evaluate the world. And you may not realise it until the logic changes and no one in the room can explain why the conclusions are different.

A Question Worth Sitting With

I keep coming back to the dinner in London. Not to the fraud, but to something the CEO said afterwards, almost as an aside. "The strange thing is, I feel like I understand less about how my office works than I did a year ago. We are faster, we are better, but I cannot explain to you exactly how the decisions are being made. The tools sit between me and the information, and I am not sure I could remove them if I wanted to."

I think he put into words something a lot of principals are feeling without quite being able to articulate it. Their own operations are becoming less legible to them. Not because the technology has failed. It has not. But they moved faster than their own comprehension could keep up.

I want to be clear about something: I am not against AI. I use it. The offices that get this right will have a real edge, and that edge will compound. But getting it right means something specific. It means putting your security budget on the same footing as your efficiency budget. It means keeping your people sharp enough to work without the tools if they have to. A CIO I know well makes his team do a fully manual portfolio review every quarter. No AI involved. Three days of painstaking work. He told me he considers it the most valuable thing they do all year. I believe him.

Charlie Munger once said, "People are trying to be smart. All I am trying to do is not be idiotic, but it is harder than most people think." I have always liked that line. It applies with uncomfortable precision to the current moment.

The offices that get hurt by AI over the next five years will not be the ones that ignored it. They will be the ones that adopted it without thinking hard enough about what they were trading away. Fast reporting, sharp analytics, lean operations, and underneath it all, a team that cannot function without the tools. A security model built on the assumption that threats only come from outside. A set of vendor dependencies that nobody fully mapped until it was too late.

The blind spot was never the technology. It is the quiet assumption, almost never examined, that faster and smarter mean the same thing.

1. Ocorian, Global Family Office Study 2026 (March 2026) - https://www.ocorian.com/knowledge-hub/insights/family-offices-turn-ai-avoid-investing-sector-now
2. Deloitte Private, Family Office Insights Series: Digital Transformation (2024) - https://www.deloitte.com/global/en/services/deloitte-private/research/family-office-insights-series.html
3. BlackRock, 2025 Global Family Office Survey (June 2025) - https://www.blackrock.com/corporate/newsroom/press-releases/article/corporate-one/press-releases/blackrock-family-office-survey-2025
4. UBS, Global Family Office Report 2025 - https://www.ubs.com/content/dam/assets/wma/static/documents/ubs-gfo-report.pdf
5. McKinsey, "State of AI Trust in 2026: Shifting to the Agentic Era" (March 2026) - https://www.mckinsey.com/capabilities/tech-and-ai/our-insights/tech-forward/state-of-ai-trust-in-2026-shifting-to-the-agentic-era
6. US Treasury, Financial Services AI Risk Management Framework (February 2026) - https://home.treasury.gov/news/press-releases/sb0395
7. Deloitte Private, Family Office Cybersecurity Report (2024) - https://www.deloitte.com/uk/en/services/deloitte-private/about/family-office-cybersecurity-report.html
8. UK NCSC, "Impact of AI on the Cyber Threat: Now to 2027" (2025) - https://www.ncsc.gov.uk/report/impact-of-ai-on-cyber-threat
9. NCSC Annual Review 2025 (October 2025) - https://www.ncsc.gov.uk/collection/ncsc-annual-review-2025
10. FBI Internet Crime Complaint Center, 2025 Internet Crime Report - https://www.ic3.gov/AnnualReport/Reports/2025_IC3Report.pdf
11. CNN, "Finance worker pays out $25 million after video call with deepfake CFO" (Feb/May 2024) - https://www.cnn.com/2024/05/16/tech/arup-deepfake-scam-loss-hong-kong-intl-hnk
12. Bloomberg (July 2024); MIT Sloan Management Review analysis - https://sloanreview.mit.edu/article/how-ferrari-hit-the-brakes-on-a-deepfake-ceo/
13. Omega Systems, Financial Services Cyber Resilience Report (October 2025) - https://omegasystemscorp.com/insights/blog/omega-systems-releases-2025-financial-services-cyber-resilience-report/
14. Fortune (March 2026), "Anthropic accidentally exposed internal documents about unreleased Mythos model"; IRONSCALES analysis - https://ironscales.com/blog/anthropics-mythos-leak-is-a-wake-up-call-phishing-3.0-is-already-here
15. CNBC (March 2026), cybersecurity stock sell-off following Mythos disclosure, cited in IRONSCALES - https://ironscales.com/blog/anthropics-mythos-leak-is-a-wake-up-call-phishing-3.0-is-already-here
16. CBS News (April 2026), "Anthropic's Mythos AI can spot weaknesses in almost every system" - https://www.cbsnews.com/news/mythos-anthropic-ai-project-glasswing-hacker-threat/
17. ASPI Cyber and Tech Digest (April 2026), UK AI Security Institute assessment of Mythos - https://aspicts.substack.com/p/anthropic-unveils-autonomous-cyber
18. Dave DeWalt quoted in InfoRiskToday (April 2026) - https://www.inforisktoday.com/ai-based-threats-usher-in-dark-period-for-cyber-defenders-a-31184
19. Budzyn et al., "Endoscopist deskilling risk after exposure to artificial intelligence," The Lancet Gastroenterology & Hepatology, Vol. 10, Issue 10 (August 2025) - https://pubmed.ncbi.nlm.nih.gov/40816301/
20. PMC/Behavioral Sciences, "Trust Dynamics in Financial Decision Making" (October 2024) - https://pmc.ncbi.nlm.nih.gov/articles/PMC11505338/
21. Choi and Xie (MIT/Stanford), AI in Accounting study (2025), cited in CFO Dive - https://www.cfodive.com/news/ai-cuts-monthly-financial-close-time-75-days-mit-stanford-study-accounting-accountants/757610/
22. Oettl et al., "Never-skilling" framework, Journal of Experimental Orthopaedics (March 2026) - https://pmc.ncbi.nlm.nih.gov/articles/PMC12955832/
23. British Journal of Psychiatry, "Artificial Intelligence and Deskilling in Medicine" (January 2026) - https://www.cambridge.org/core/journals/the-british-journal-of-psychiatry/article/artificial-intelligence-and-deskilling-in-medicine/641F0B39DC9494C7B9561F63E718EAC6
24. Banking Dive, "Upstart investors sue over 'overreactive' AI" (April 2026) - https://www.bankingdive.com/news/upstart-investors-sue-overreactive-ai-revenue-adjustment-damages/817226/
25. Financial Stability Board, "The Financial Stability Implications of Artificial Intelligence" (November 2024) - https://www.fsb.org/uploads/P14112024.pdf

About the Author

Dr. Ralph Welpe is Chairman and CEO of Mirai Capital Partners, a London-based boutique advisory firm working with some of the world's most influential families and entrepreneurs. With over 25 years of experience spanning investment banking, private banking, and alternative investments across Europe, Asia, and the Middle East, he advises family principals on strategy, governance, and cross-border investment structuring.

Contact: ralph.welpe@miraicapitalpartners.com | www.miraicapitalpartners.com | London, United Kingdom