The Hype, The Hope, and The Hard Truth

AI, Financial Advice, and What Actually Matters

Pete Ridlington
CEO & Founder, Ningi
Chartered Financial Planner (FPFS)

April 2026

I Run an AI Company. I'm Telling You to Calm Down.

I feel like I need to start with a disclaimer. I am the CEO and founder of Ningi, a technology company that builds AI-powered platforms for financial advisers. I have a Mac mini on my desk running open-source AI models. I use Claude, ChatGPT, and Cursor daily. My livelihood depends on AI being impactful and transformative. I am, by any reasonable measure, all in on this technology.

I also see what you see. I open LinkedIn and more than half the long-form posts are AI-generated content about AI. Every conference has an AI panel. Every vendor pitch starts with AI. Professional Adviser just ranked 15 different AI tools competing for the attention of financial advisers — note-takers, report writers, compliance checkers — and most of them didn't exist eighteen months ago. UK AI startups raised £6 billion last year. Global AI spending hit $644 billion. The message coming at you from every direction is the same: this is urgent, this is existential, act now.

And I'm telling you to calm down.

Not because the technology isn't impressive. It is. Not because nothing is changing. Things are changing. But because the gap between what is being predicted and what is actually happening has become so wide that it's distorting how an entire industry thinks about its future. Financial advisers are being told, with increasing urgency, that their profession is on the brink of extinction. That AI will replace them. That the empathy they provide is a delusion. That if they don't radically transform their businesses in the next twelve months, they're finished.

I don't think that's true. And I think I can say that with more credibility than most, precisely because my financial future depends on AI working.

The most reliable prediction of the future is based on the present.

The fundamentals of how people interact with each other, how they make decisions, how they engage with money and risk and uncertainty — these things don't change at the speed that technology changes. They change at the speed that humans change, which is generational. The burden of proof required to suggest the world is changing at some fundamental level should be high. Saying so isn't burying your head in the sand. From a scientific perspective, it's the most sensible starting position.

The Prediction Machine

Before we look at the data, it's worth asking a question that doesn't get asked enough: who is making these predictions, and what are their incentives?

The most alarming predictions about AI disruption come from the CEOs of AI companies. Dario Amodei of Anthropic predicts 50% of entry-level white-collar jobs gone within one to five years. Sam Altman of OpenAI talks about artificial general intelligence arriving by 2027. Matt Shumer, founder of an AI startup, writes viral essays about how he is no longer needed for the technical work of his own job.

Now, I'm not saying these people are lying. They may fully believe what they're predicting, and they have data and conviction behind it. But consider their position. These are people who have raised billions of pounds from investors who need a return. Their companies are valued at extraordinary multiples based on the promise that this technology will be as transformative as they say it will be. Their personal wealth, their companies' survival, their investors' patience — all of it depends on these predictions coming true.

> *"If someone is selling you electric scooters, they're going to tell you electric scooters are the future of urban transport. If they've raised two billion dollars in venture capital to prove it, they're going to say it with even more conviction. That doesn't make them wrong. But it means their predictions are inseparable from their incentives, and you should weigh them accordingly."*
>
> — Pete Ridlington

The pattern repeats with remarkable consistency. Elon Musk has been saying that fully autonomous self-driving cars are six to twelve months away for roughly a decade. People have gone viral making montages of these announcements on YouTube — the same prediction, the same confidence, the same timeline, repeated year after year. The cars still aren't fully self-driving. It's not that the technology hasn't improved — it has, enormously. But the predictions have been consistently, spectacularly wrong on timing.

▶ WATCH

Elon Musk Promises Full Self-Driving "Next Year"

MG Investing · 7.3K views · YouTube

Top comment: "I come back to watch this every year"

youtube.com/watch?v=B4rdISpXigM

youtube.com/watch?v=B4rdISpXigM

The metaverse is another cautionary tale. Five years ago, Mark Zuckerberg told us we'd be living in a digital third world. Meta invested over $36 billion. Hundreds of companies pivoted. Entire business plans were rewritten. And then, quietly, it was wound down. Nothing fundamentally changed.

▶ WATCH

Why The Metaverse Was Doomed From The Start

The Art Of Storytelling · 318K views · YouTube

Despite investing over $73 billion... Horizon Worlds will transition to a mobile-first experience

youtube.com/watch?v=hLItp5SiRGw

youtube.com/watch?v=hLItp5SiRGw

You can point at the same pattern specifically with AI. In 2016, the CEO of a prominent fintech company predicted that half of all financial advisers would be replaced by AI within a decade. We're nearly there. It hasn't happened. In 2017, PwC published a widely cited report claiming AI would eliminate 30% of UK jobs by the early 2030s. Oracle ran studies predicting that AI-based robots would soon replace financial advisers entirely. Deloitte forecast that 114,000 legal jobs would be automated by 2036. None of these predictions have materialised in the way they were framed.

Is it different this time? Is the technology different? Of course. But it's been "different this time" every time this prediction has been made.

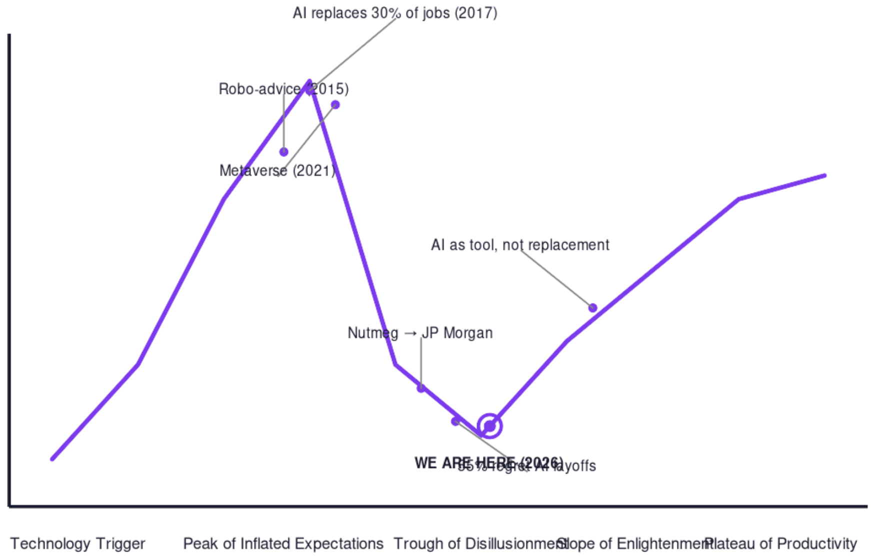

Figure 1: The Technology Hype Cycle — Where Are We Really?

The All-In Test

There's an interesting discussion from the All-In Podcast where David Sacks frames the strategic question for software companies clearly: the value capture is shifting from individual tools to whoever owns the workspace — the unified layer where AI has access to all context and can act across everything. That's a real insight about how technology businesses will need to be structured. But it's a structural observation about software, not a prediction about the end of human professions.

Even among the tech investors who live and breathe this stuff, there's significant scepticism. Scott Galloway responded to Amodei's predictions at Davos 2026 by pointing out that every previous technology created more jobs than it destroyed, and he saw no reason to think AI would be any different. The debate is far less settled than the headlines suggest.

What The Data Actually Says

We're far enough into the AI era now that we can stop speculating and start looking at evidence. And the evidence tells a more measured story than the hype would suggest.

Adoption Is Real But Slow

Global adoption of generative AI tools reached 16.3% of the world's working-age population in the second half of 2025, according to Microsoft's AI Economy Institute. In the UK, only one in six businesses (16%) are currently using any form of AI technology, according to the UK Government's own research published in February 2026. An academic study published in ScienceDirect found that the actual AI use rate among the general population rose from 3.7% in September 2023 to 5.4% in February 2024 — growing, but hardly the overnight transformation the headlines suggest.

Even where people are using AI, the depth of usage is shallow. The most common use cases remain search and basic research — essentially using it as a slightly better Google. My mum uses ChatGPT. That's surprising in some ways, but it took years to get there, and she's mostly asking it questions she'd previously have typed into a search engine.

16.3% of the world's working-age population used generative AI in the second half of 2025. We are years into this groundbreaking technology, and most people aren't using it.

The Spending-to-Results Gap

Here's the number that puts all of this in context. MIT's Project NANDA reported in 2025 that 95% of enterprise AI initiatives fail to deliver measurable return on investment. Only 15% of AI decision-makers reported any improvement to their bottom line in the past twelve months. Global spending on generative AI hit $644 billion — up 76% year-on-year — but Gartner has officially moved generative AI past peak hype and into the "Trough of Disillusionment." Forrester's headline for 2026: "AI Moves From

Hype To Hard Hat Work." The money is flowing in. The results are not flowing out at the same rate.

The Rehiring Wave

Perhaps the most telling data point of all: the companies that loudly replaced workers with AI are now quietly hiring them back.

A report covered by People Matters Global in March 2026 found that 55% of companies that laid off workers and attributed it to AI now regret the decision. Among those that conducted AI-driven layoffs, 32.7% have already rehired between 25% and 50% of the positions they cut. Gartner forecasts that by 2027, half of all companies that attributed customer service headcount reductions to AI will rehire for those roles. Inc Magazine ran a headline in February 2026 that should be pinned to the wall of every conference room in Silicon Valley: AI layoffs are backfiring, and half of companies will start rehiring.

I have underlying scepticism about the big reports where companies announce they're replacing staff with AI. Is that really what's happening, or are they using AI as a positive market statement about downsizing teams they would have cut anyway because they overhired during the pandemic? The data increasingly suggests the latter. We saw the same dynamic with offshoring in the 2000s: companies loudly announced they were moving jobs overseas, then quietly brought many of them back within a few years when the reality didn't match the press release.

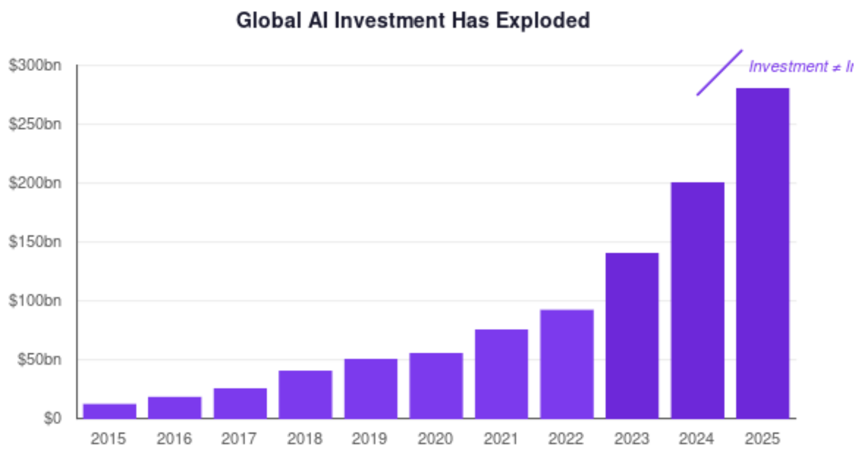

Figure 3: Global AI investment has reached staggering levels — but investment is not the same as impact.

The Cost Equation

There's a financial reality to AI that rarely features in the predictions. If you replace a software engineer with AI and force one of the remaining engineers to move faster using AI tools, the token costs can approach what you'd have been paying in salary. The cost of AI could come down — and probably will — but these companies also need to start making a profit. They have investors to repay. A lot of mid-term analysis, especially when you factor in the infrastructure required to power this technology, suggests that costs will actually go up before they come down.

For a financial advice firm, the same logic applies. If you're now paying £200 to £400 a month across different AI licences and subscriptions, has your revenue gone up by that amount? Has your capacity genuinely increased? For most firms, the honest answer is no. The fact that you're reading this white paper, or attending a conference about AI, or experimenting with yet another tool, is itself evidence that AI is currently taking your time rather than giving it to you.

Why People Hire Experts: A Model

To understand why AI won't simply replace financial advisers, you need to understand why people hire them in the first place. It's more nuanced than the "empathy" argument suggests.

I've carried a model in my head for a long time about how people engage with their finances. Imagine two dimensions. On one axis, you have cognitive capability — not intelligence in general, but the specific capacity and inclination to understand financial concepts. Someone can be a brilliant doctor or lawyer but have a brain that simply doesn't engage with financial planning. On the other axis, you have motivation — how much someone actually wants to spend their time on this stuff.

The Four Quadrants

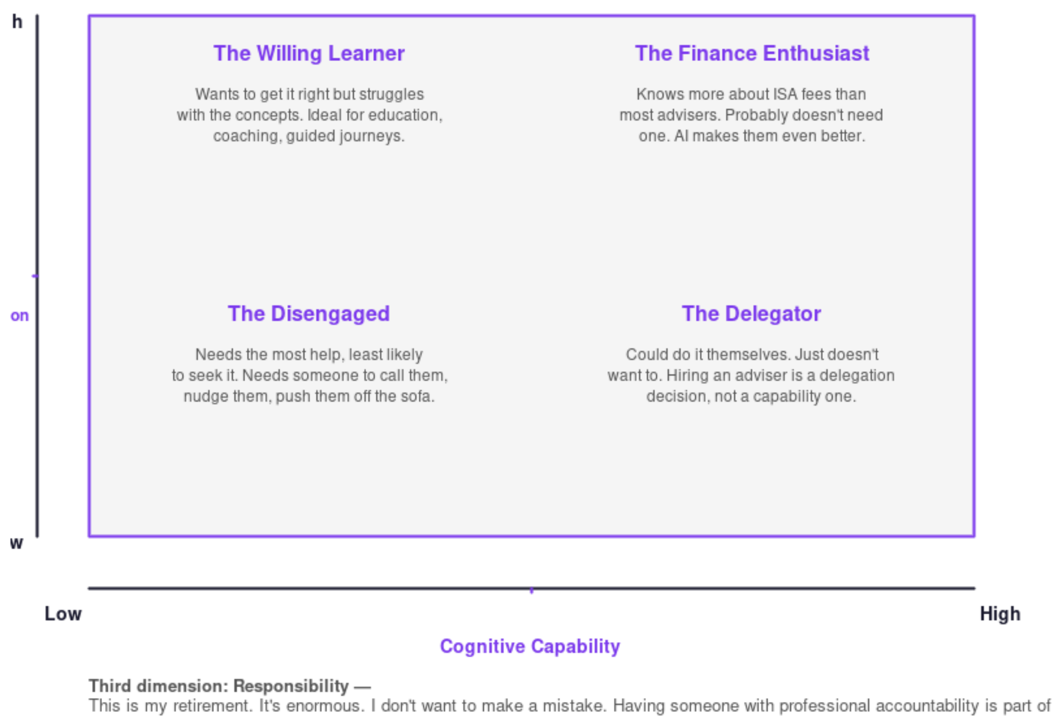

Figure 2: The Motivation × Capability Quadrant

High capability, high motivation: These are the personal finance enthusiasts. They know more about exchange rates, ISA fees, and investment platforms than most financial advisers. I have friends in this category. They love it, they're good at it, and they practise what they preach. They probably don't need a financial adviser, and AI tools will make them even more capable.

High capability, low motivation: These people could do it themselves. They have the intellectual capacity. They just don't want to spend their time on it. They'd rather outsource the cognitive load and get on with the things they care about. For them, hiring a financial adviser is a delegation decision, not a capability decision.

Low capability, high motivation: People who really want to get their finances right but struggle with the concepts. They're the ideal clients for education-led propositions — courses, coaching, guided digital journeys. AI can serve this group well, but they often want a human to check their thinking.

Low capability, low motivation: The largest group, and the hardest to serve. These people need the most help but are the least likely to seek it. No amount of AI sophistication will change the fact that they're not going to proactively engage with a financial planning tool. This is the group that needs someone to call them, to nudge them, to push them off the sofa.

The Third Dimension: Responsibility

But there's a third vector that complicates the picture further: the desire to outsource responsibility. "This is my retirement. It's enormous. I don't want to make a mistake with it. Having someone I can rely on — someone with professional accountability — is part of the point."

People frame this as "I want someone to sue," and there's a regulatory dimension to it. But the deeper truth is more about hierarchy and expertise delegation. When something is big, scary, and consequential, humans instinctively want to hand it to someone they perceive as an expert. Not just someone who has information — someone who has standing.

Malcolm Gladwell tells a powerful story in Outliers about plane crashes. In the late 1990s, Korean Air had a catastrophic safety record. When investigators analysed the black box recordings, they found a consistent pattern: co-pilots who knew the plane was in danger — running low on fuel, approaching terrain — would express their concerns in soft, deferential language rather than direct warnings. The cultural norm of deference to the captain was so strong that co-pilots would essentially hint at impending disaster rather than state it plainly. The captain, unbothered, would continue on course.

The fix wasn't better data or better instruments. It was restructuring the hierarchy in the cockpit — giving co-pilots explicit authority and training them to communicate directly.

Financial advice has a similar dynamic, but in reverse. People don't just want information about their finances. They want to delegate the decision to someone with hierarchical authority on the subject. When I'm threatened with legal action, I often do some initial research with ChatGPT. But I want a lawyer there in the court with me. Not because the lawyer has access to different information — in many cases, AI could surface the same case law. But because the lawyer has standing, experience, and accountability that I don't.

> *"If I ever sell Ningi for serious money — and let's be clear, Jym is nowhere near it when that happens — I already know which financial advisers I'd go to. Not because they have access to information I couldn't find myself, but because they have a kind of authority I'd defer to. I'd want someone who'd stop me doing something stupid with it. I'm very good at convincing myself I need a new car, and I want someone who won't let me get away with that."*
>
> — Pete Ridlington

The Personal Trainer Analogy

Apply this model to personal trainers. All the information exists online for me to learn how to run a good diet and how to exercise. I can build a plan. Some people have the capability and motivation to do that. But I go to a personal trainer and say: keep an eye on my form, make sure I'm doing these deadlifts correctly, come up with a plan, and make sure I follow through with it. I'm outsourcing my motivation, my cognitive load, and the responsibility to learn and understand this stuff.

And here's the thing: if my personal trainer lets me off the hook — if I don't show up and they don't give me grief about it — they're not a very good personal trainer. AI is desperate to please. It's a yes-bot. That's fine for answering questions. It's not fine for the part of financial advice that actually changes outcomes: holding people accountable when they'd rather not be held accountable.

The industry has never been about something that other people fundamentally can't do for themselves. Hargreaves Lansdown has been showing you stocks in the back of a paper for decades. You could have called the death of the adviser with robo-advice, with online investing, with Reddit personal finance. You don't technically need a financial adviser. You probably never did. AI is just another iteration of the DIY versus expert argument — one that's been playing out since long before large language models existed.

The 1980s Time Traveller

Here's a thought experiment. If you dropped someone from 1985 into today's world, their mind would be blown. Smartphones, same-day delivery, video calls to Australia, instant access to all human knowledge.

But would it really be that alien? You can buy things from digital online shops and they get delivered tomorrow — but is that so completely different from going to a physical shop? You can access all the information in the world — but is that fundamentally different from what was starting with TV, radio, and books? You can call someone on the other side of the planet and see them on a screen — but you could also pick up a telephone in 1985 and call someone in Australia.

Whilst things are radically different, they're also incremental steps forward. The same thing is going to happen with AI: exactly the same, only differently.

Generational Change, Not Overnight Change

There's a commonly observed psychological bias about how it feels like things are changing faster now than they used to. We have rose-tinted glasses for how things used to be. As you age, time genuinely does seem to speed up. But the actual pace of deep behavioural change is measured in generations, not product cycles.

If my mum is going to buy a dress for a night out, she goes to the shops. She could go online. But for the first thirty or forty years of her life, she went to the shops, and that's her experience, and that's how she likes to do it. The same thing will be true of AI. The way that forty-year-olds of today do things probably won't fundamentally change. You almost need a generation to grow old and be replaced by AI natives for these shifts to take hold. You've seen it with computers — everyone uses them, but the people who grew up before them didn't fundamentally restructure their lives around them.

We overestimate the change that will happen in the short term and underestimate the change in the long term.

That well-worn observation — often attributed to Bill Gates — is exactly what's playing out. In 1997, the US Department of Labor predicted that the internet would make insurance agents, travel agents, and financial advisers largely redundant within fifteen years. Travel agents were significantly disrupted. Insurance agents adapted. Financial advisers grew in number. The same technology produced three different outcomes in three adjacent professions, because the human dynamics of each were different.

These changes are coming. You need to be aware of them and in front of them. But this is a generational shift, not a twelve-month deadline. The world of financial planning in the early 2030s will look quite similar to the world of financial planning in the early 2020s. Different tools, different efficiencies, the same fundamental human dynamics.

The Industry Has Heard This Before

The financial advice industry has been told it's on the brink of technological extinction roughly once every five years for the past two decades. It's worth taking a brief tour of the predictions that didn't come true, because the pattern is instructive.

A Brief History of the End of Financial Advice

In the mid-2000s, the rise of online brokerages and execution-only platforms was supposed to disintermediate advisers entirely. Why pay for advice when you could trade for free? The platforms grew, but so did the advice profession. Different markets, different needs.

In 2012–2015, robo-advice was the existential threat. Betterment, Wealthfront, and in the UK, Nutmeg, were going to democratise portfolio management and make human advisers obsolete. Billions were invested. Conferences were held. Predictions were made with great confidence. Nutmeg grew to £8.5 billion in assets under management and 265,000 clients and still couldn't make the standalone model work. It folded into JP Morgan. Wealthfront was acquired by UBS. The robo-advisers didn't replace human advisers — the best ones became tools that human advisers used.

In 2016, the CEO of a well-known fintech company predicted that half of all financial advisers would be replaced by AI within ten years. In 2017, a widely shared Oxford study ranked financial advisers among the professions most likely to be automated. PwC predicted 30% of UK jobs eliminated by the early 2030s. Reports from consultancies including Deloitte and Accenture forecast mass displacement across professional services.

We are now in 2026. The number of regulated financial advisers in the UK has remained remarkably stable. The profession didn't shrink; if anything, the shortage of qualified advisers is a bigger problem than any technological displacement.

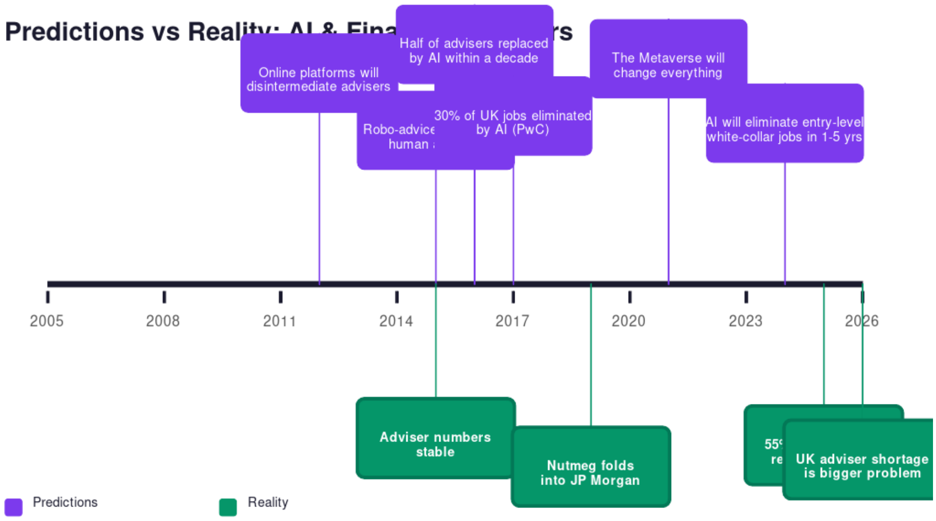

Figure 4: Two decades of predictions vs what actually happened.

> *"It hasn't happened more times than it has happened. At some point, you have to weigh the evidence of past predictions against the confidence of current ones. And the evidence is overwhelmingly on the side of "the adviser will still be here.""*
>
> — Pete Ridlington

What AI Has Actually Delivered So Far

The general transformer architecture was invented in 2017. We're a few years past ChatGPT's launch. And what has the financial advice industry actually got from AI so far? Note-takers and report writers. Which, as a technological solution, isn't that different from the dictation and report generation software that financial planners were using on DOS in the early 1990s. You put a few basic details in, and you got a report generated. We've added a conversational interface to something that already existed.

When I got into AI in 2017 at Wealth Wizards, we built genuinely impressive technology. I thought it was going to be the end of days. There were people making predictions that financial advisers were the most at-risk category from AI, that they wouldn't exist in a few years. That never happened. Not because the technology was bad — it was good. But because technology adoption in a regulated, trust-dependent, relationship-driven profession doesn't move at the speed of technology development.

There's also something I find genuinely interesting that doesn't get discussed enough. I read something where someone asked: "Why are LLMs writing poetry and music and doing research and thinking, when meanwhile I'm still washing up the dishes and putting the bins out?" Get the AI to do the stuff we don't want to do, not the stuff we do want to do. The technology is solving for intellectual problems while most people's daily frustrations remain resolutely physical and mundane.

Pete Ridlington, CEO & Founder, Ningi

Hold Both Hands

Here is the most important thing I can tell you, and it's the thing I tell myself every day: you have to hold both perspectives in your head at the same time.

Perspective one: AGI is coming. AI is going to eat your lunch. It's the end of days. What does that mean for you? Keep your team small, lean, and agile. Be regenerative and profitable. Learn these tools. Be front-footed. Don't put your head in the sand.

Perspective two: This is hype. It's overblown. There will be some change and impact and opportunity, but it's not earth-shattering. What does that mean for you? Keep your team lean. Adopt this technology. Keep cracking on.

The response is the same. The practical actions are identical regardless of which future materialises.

Alan Smith at Capital Asset Management frames this well as a kind of Pascal's Wager: the downside of preparing for transformation that doesn't come is that you become more efficient, more tech-savvy, and deliver a better client experience. The downside of not preparing for transformation that does come is existential. Three of the four possible outcomes are positive or neutral. Only one is catastrophic, and it's the one that comes from doing nothing.

I agree with that framing completely. But I'd add this: don't let the fear of the catastrophic outcome drive you into panicked, reactive decisions. Don't buy an AI note-taker because a LinkedIn article scared you. Don't bolt another disconnected tool onto a tech stack that's already held together with spreadsheets and prayer. The calm, deliberate version of preparation is better than the anxious one.

So What Should You Actually Do?

The business case for transforming how financial advice is delivered hasn't changed. I first encountered it in 2010, working on a project where I thought: we can automate this, we can remove risk with technology, we can improve the customer experience with a client portal. Sixteen years later, the same problems exist. Advisers still deal with fragmented technology stacks — twenty different systems that don't talk to each other. They still spend 60% of their time on administration. They still serve around 115 families and dream about serving more.

What AI does is make the solution more achievable, not different. The opportunity was always: one adviser, with the right technology, should be able to serve 3,000 people across a continuum — from digital engagement and education at one end through to face-to-face regulated advice for complex needs at the other. AI makes that continuum more viable. But only if you approach it as a fundamental rethink of your operating model, not as a tool you bolt onto the side.

Don't Add Another Tool

If your response to AI is buying a note-taker, subscribing to an AI compliance checker, and adding another disconnected system to your stack, that is definitively not the answer. Most advice firms already have ten to twenty disconnected tools. Adding AI point solutions on top of fragmented infrastructure is like putting a turbocharger on a car with no wheels. The engine might be impressive, but you're not going anywhere.

AI's real power comes from context. An AI that can see your client's complete picture — their conversations, their documents, their financial plan, their risk profile, their meeting history — and reason across all of it simultaneously is transformational. An AI that can only see the transcript from last Tuesday's meeting is a slightly better dictation machine.

Get Your Data Ready

The single most valuable thing you can do right now, regardless of what you believe about AI's trajectory, is consolidate your data. Get everything into one place. Client records, communications, documents, compliance files, meeting notes. Not because you need AI today, but because when AI capabilities advance further — and they will

— the firms with unified, structured data will be able to deploy those capabilities immediately. The firms with data scattered across twenty systems will spend years catching up.

Find a Partner, Not a Product

The build versus buy debate in financial advice technology has always been the wrong question. Building your own technology is prohibitively expensive and distracting. Buying off-the-shelf locks you into someone else's opinion of how your business should work. The right answer is partnership: working with a specialist technology company that builds and maintains the platform while you shape what it does.

This matters more now than ever. AI capabilities are evolving so quickly that any product you buy today will be outdated in twelve months. What you need is a partner who is evolving with the technology, who is building AI into the foundation of the platform rather than bolting it on as an afterthought, and who will take you on the journey rather than sell you a destination.

Build the Muscles

Use these tools. Experiment. Learn. Spend time with AI every week — not reading articles about it, but actually using it in your work. Give it a client scenario. Ask it to draft a suitability letter. See where it's good and where it falls over. Build an intuition for what this technology can and can't do, so that when new capabilities arrive, you can evaluate them from a position of knowledge rather than fear or hype.

The advisers who will thrive aren't the ones who predict the future correctly. They're the ones who are adaptable enough to respond to whatever future arrives.

Ningi's Position

Ningi exists because we believe the financial advice industry is at an inflection point — not because of AI specifically, but because the gap between what technology makes possible and what firms are actually doing has never been wider. AI is one part of that, but the operating model, the data infrastructure, the client experience, and the business economics all need to change together.

We operate as a technology partner to financial advice firms. That means we build and maintain the platform, handling the engineering, security, and continuous improvement. Firms shape the product roadmap, bringing their expertise and client understanding. Both parties invest in the relationship because both benefit from its success.

Our platform is AI-native — not AI-bolted-on. The same team that built the back office built the AI, sharing the same data layer. When you transcribe a meeting, generate a compliance document, and send a client communication, all of it flows through a single system. The AI has context across the complete client relationship, which is why it can do things that disconnected tools simply cannot.

We don't know exactly what the next five years look like. Nobody does, despite what they might tell you on LinkedIn. But we're committed to being on the journey with the firms we work with — moving quickly, adapting constantly, and building something that serves advisers and their clients regardless of which version of the future arrives.

> *"We're not selling you a solution. We're building one with you."*
> — Ningi

About the Author

Pete Ridlington is the CEO and founder of Ningi, a technology platform for financial advisers. He is a Chartered Financial Planner (FPFS) who has spent over fifteen years at the intersection of financial planning and technology, including roles at Wealth Wizards and founding Ningi in 2020. He lives in Peterborough with his family, a Mac mini running open-source AI models, and an unwavering conviction that the future of advice is human.

Contact

pete@ningi.co.uk | ningi.ai | linkedin.com/in/peteridlington