
Beyond AI:
Human insight as the advantage.
Grant Show
FUTURIST, SPECULATIVE DESIGNER,
INDEPENDENT CONSULTANT

BENNELONG POINT / DUBBAGULLEE VIEWED FROM DAWES POINT IN c1804 (MITCHELL LIBRARY, STATE LIBRARY OF NSW - V1/1810/1)

I’d like to begin by acknowledging the Gadigal people of the Eora Nation, the Traditional Custodians of the land we meet on today. This area, now known as Martin Place, sits within a network of deeply significant places for Aboriginal people, who have cared for and connected with Country for tens of thousands of years.

BENNELONG POINT / DUBBAGULLEE VIEWED FROM DAWES POINT IN c1804 (MITCHELL LIBRARY, STATE LIBRARY OF NSW - V1/1810/1)

Just a short walk from here is Bennelong Point. Before being called that it was called Cattle Point, and it’s original name was Tubowgule, which means “where the knowledge waters meet.” It was a place of ceremony, gathering, and learning. People would travel from across neighbouring Nations — the Darug from the west, the Guringai from the north, and the Tharawal from the south — to meet here with the Gadigal.

BENNELONG POINT / DUBBAGULLEE VIEWED FROM DAWES POINT IN c1804 (MITCHELL LIBRARY, STATE LIBRARY OF NSW - V1/1810/1)

In a way, our gathering today mirrors that long-standing tradition. We've come together from many places to connect, learn, and build something shared. It is a reminder that the act of coming together on this land has deep roots, ones that extend far beyond us.

I pay my respects to Elders past and present, and I extend that respect to all Aboriginal and Torres Strait Islander peoples here today.

By introduction, I'm a strategic leader and consultant. I've worked directly with many of the world's leading brands, and alongside many more as clients. Many in the room here today.

My career has taken me across industries and continents, spanning design, full-stack development, entrepreneurship, and running a national platform for innovation, creativity, and leadership.

At the core of my approach is building the transformation and innovation capabilities organisations need to thrive.

Deloitte. Allianz salesforce AUTODESK

My career has taken me across industries and continents, spanning design, full-stack development, entrepreneurship, and running a national platform for innovation, creativity, and leadership.

At the core of my approach is building the transformation and innovation capabilities organisations need to thrive.

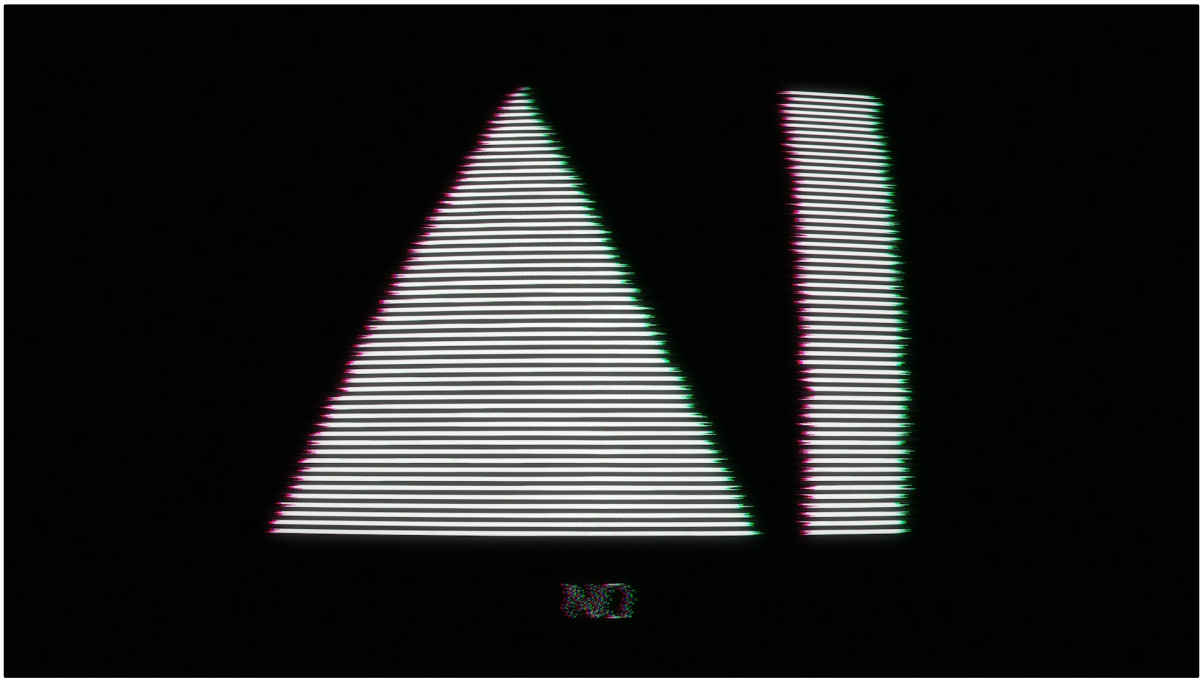

As we move through this conversation today, I might come across as anti-AI. I am not. I think AI is extraordinary, and when it is used to enhance human capability it will take us to heights we have never reached before. And indeed in many ways already has.

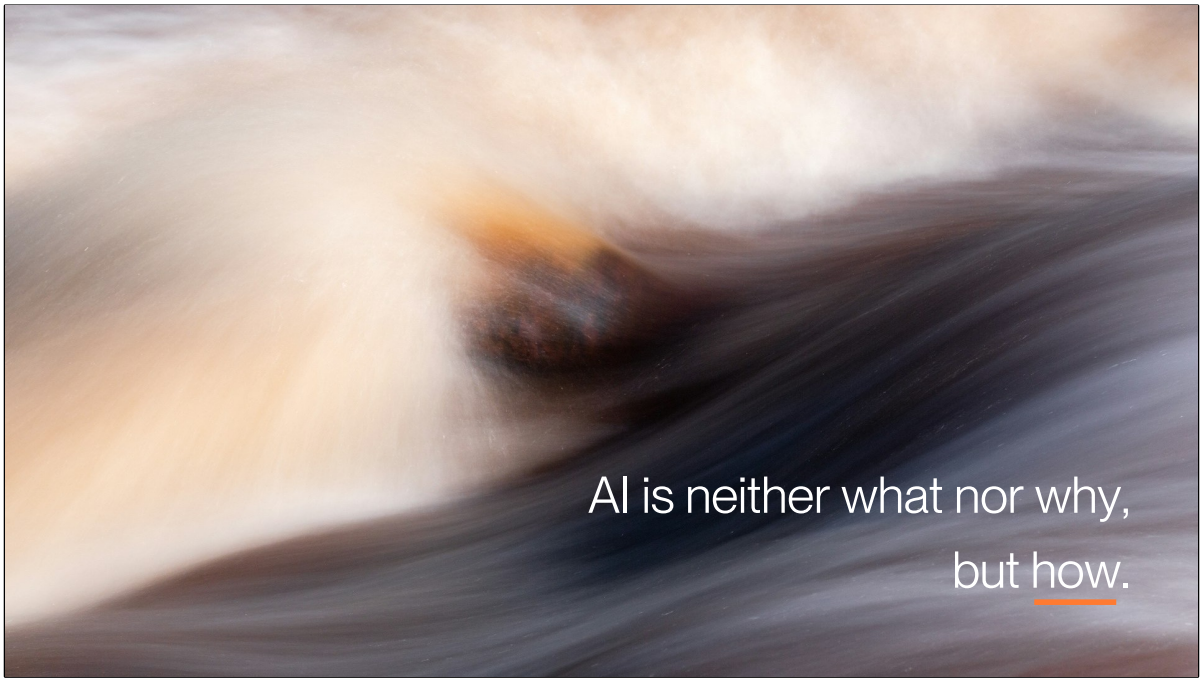

But I am also deeply pragmatic. As a rational optimist I believe progress is real, but AI must never be mistaken for the what or the why. It is quite simply a how.

And so, I want to begin here, with this quote:

The single biggest problem with communication is the illusion that it has taken place. George Bernard Shaw.

Shaw's quote captures something I've been thinking about as I work with organisations and AI. We're in the middle of the fastest technology adoption rush in history. AI has spread faster than mobile-first, faster than apps, faster than online banking itself.

It's a dizzying rate of adoption unlike anything we've seen before.

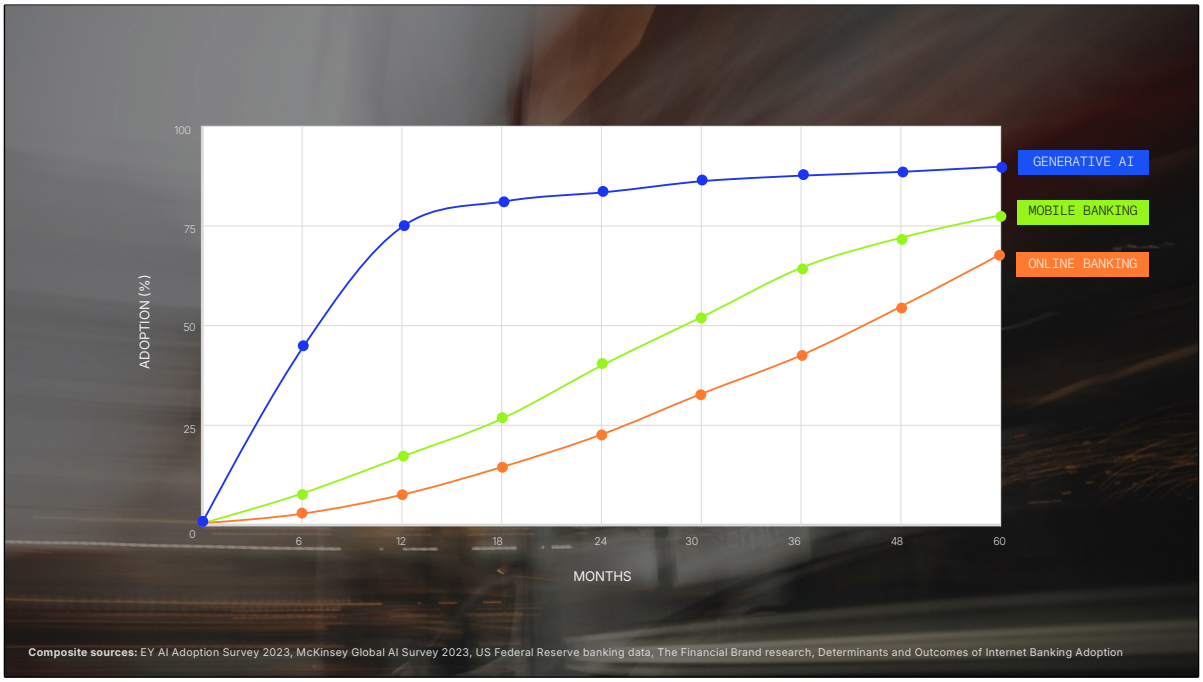

To illustrate this, I want to share a diagram of adoption over time. While internet and mobile are now central to everyday banking, both took around five years to reach broad adoption. While generative AI reached the same level in just twelve months.

The slower pace with earlier technologies allowed knowledge, governance, and safeguards to grow alongside them. With AI, the speed of adoption is itself a sign of fragility: systems spreading faster than our understanding can keep up with, and their stability depends entirely on what we know.

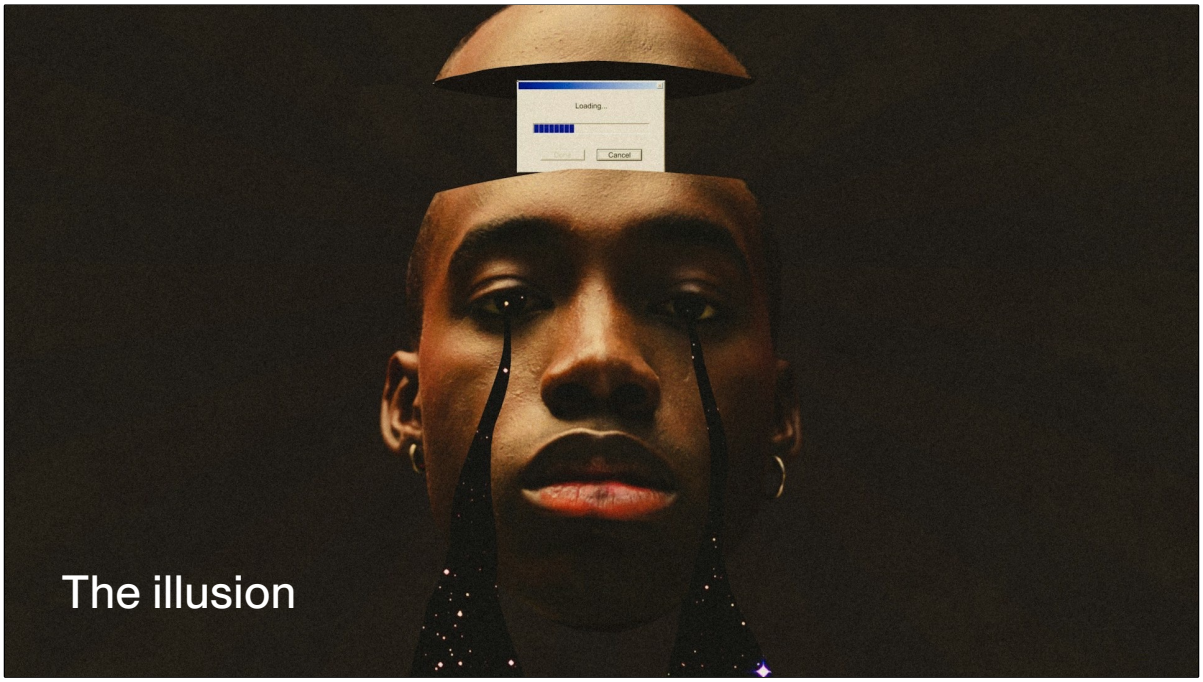

What concerns me about this is that pace is being driven largely by companies with enormous profit interests in us adopting their solutions.

While there's a whole world of possibilities and efficiencies AI can create in areas ranging from data analysis to customer service optimisation,

blind optimism and hype around the technology can ultimately have a counterproductive impact on a business," says David Kadio-Morokro of EY.

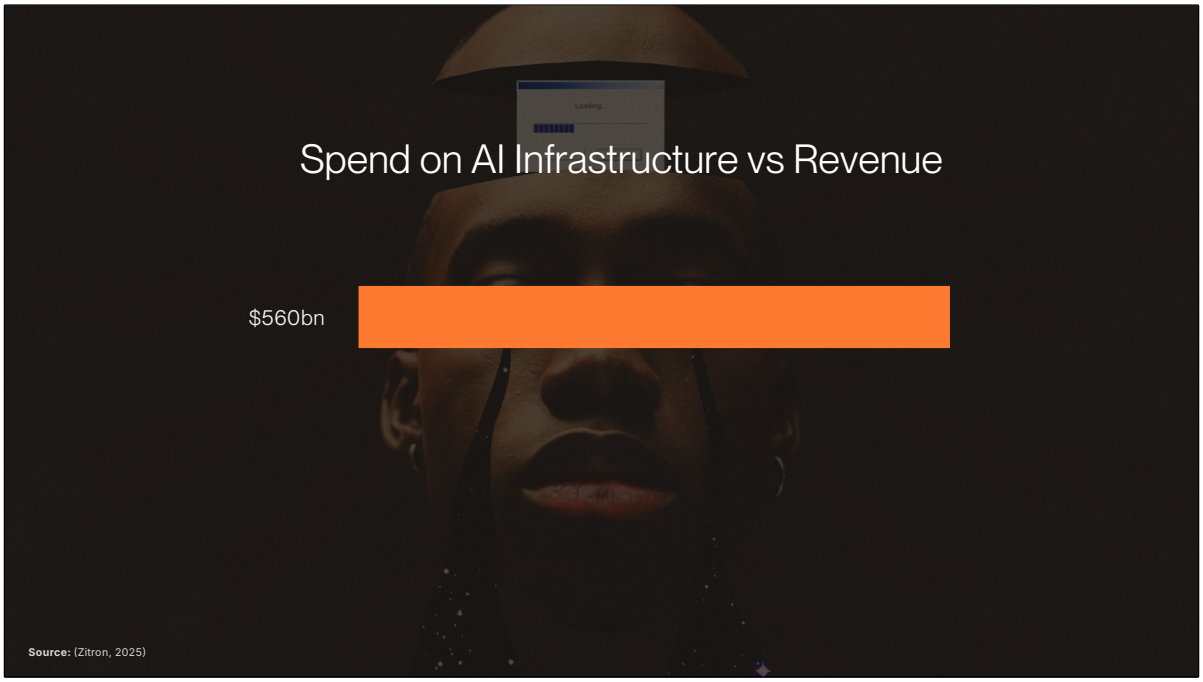

The evidence bears this out. As analyst Ed Zitron has noted, the largest tech companies are on track to spend more than 560 billion dollars on AI infrastructure in just two years.

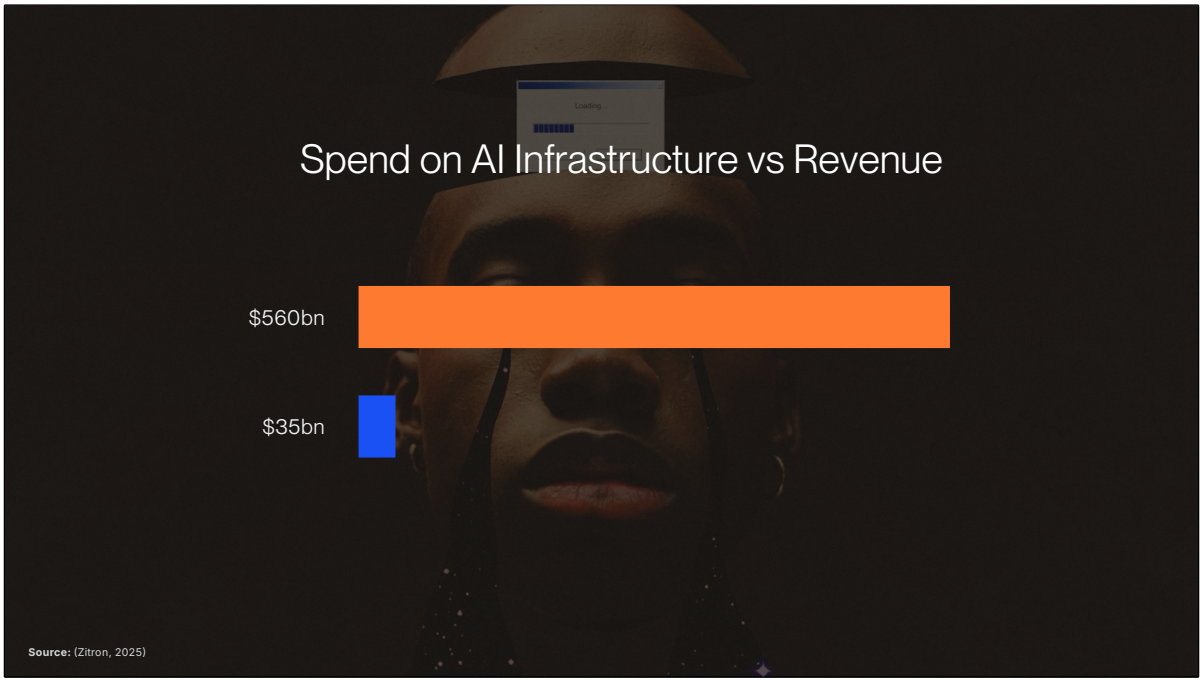

While revenues to date are closer to 35 billion. That is not resilience, that is fragility disguised as growth. When investment races ahead of knowledge, systems become brittle, and exposed to collapse when assumptions prove false (Zitron, 2025)

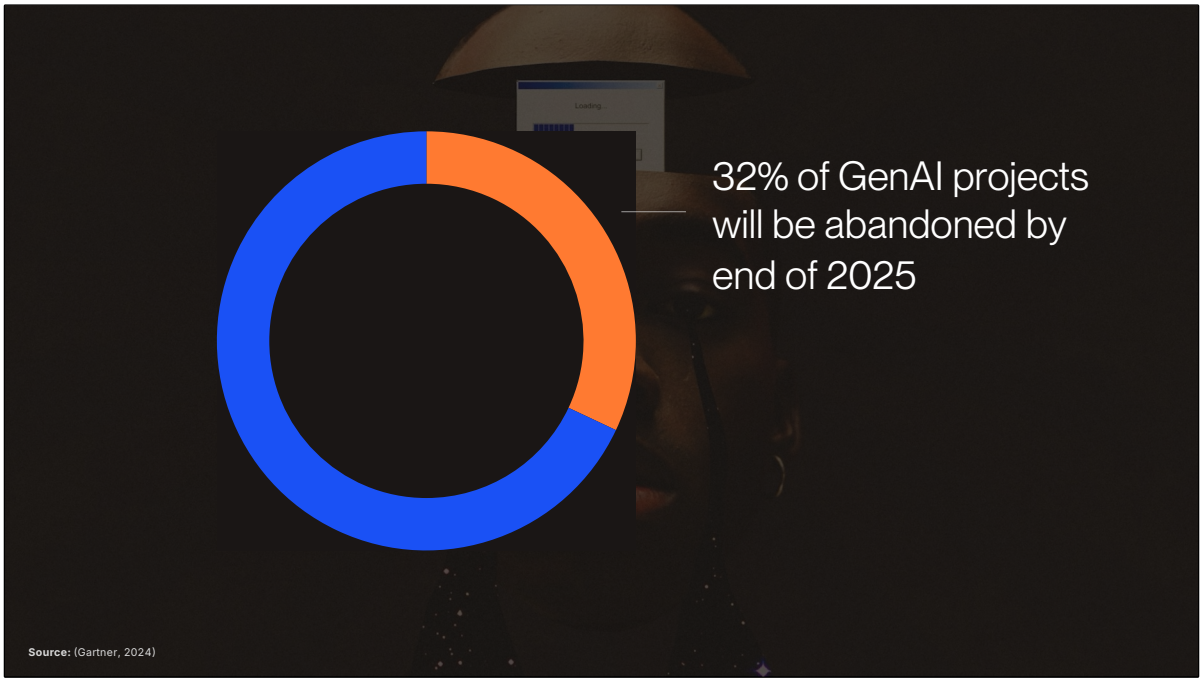

Meanwhile, Gartner similarly predicts that nearly a third of generative AI projects will be abandoned by the end of 2025 (Gartner, 2024). That’s four months away.

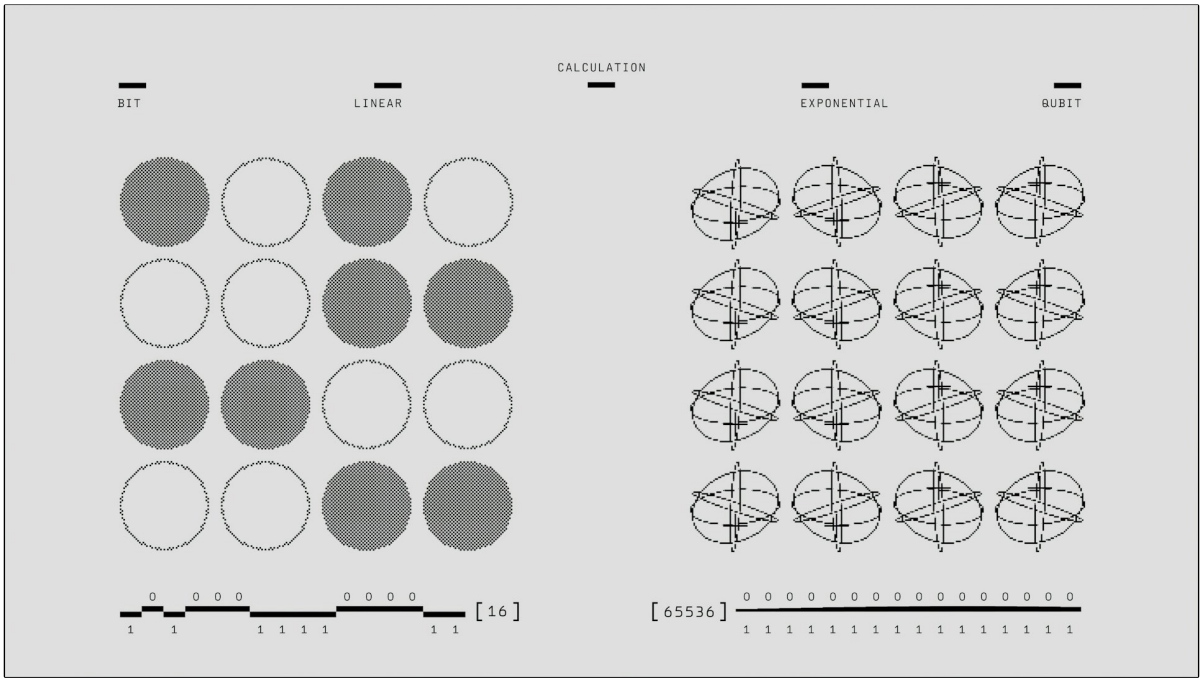

The urgency feels manufactured, the promises oversized, and the evidence remarkably thin at the scale and pace we are being asked to bet on.

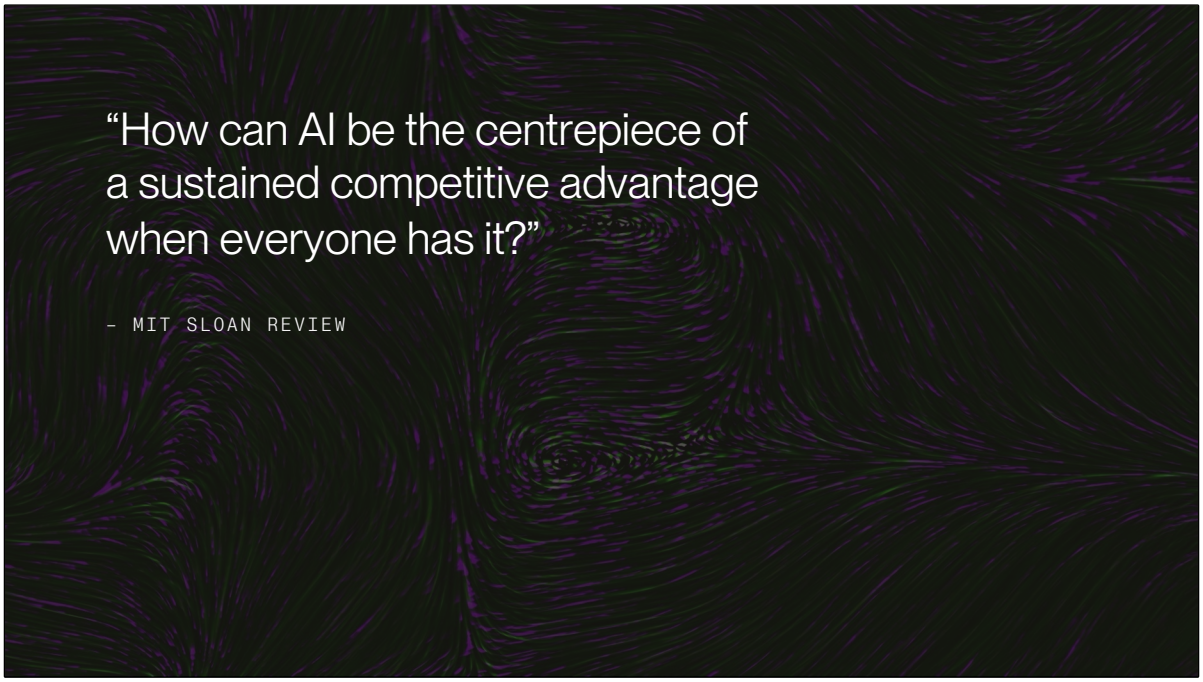

And this rush reveals a deeper issue. We are operating under the illusion that AI itself is competitive advantage. It is not, and it can’t be.

MIT put it plainly earlier this year: “How can AI be the centrepiece of a sustained competitive advantage when everyone has it?” They argue that it simply cannot. Far from being a source of differentiation, artificial intelligence will be a source of homogenisation.

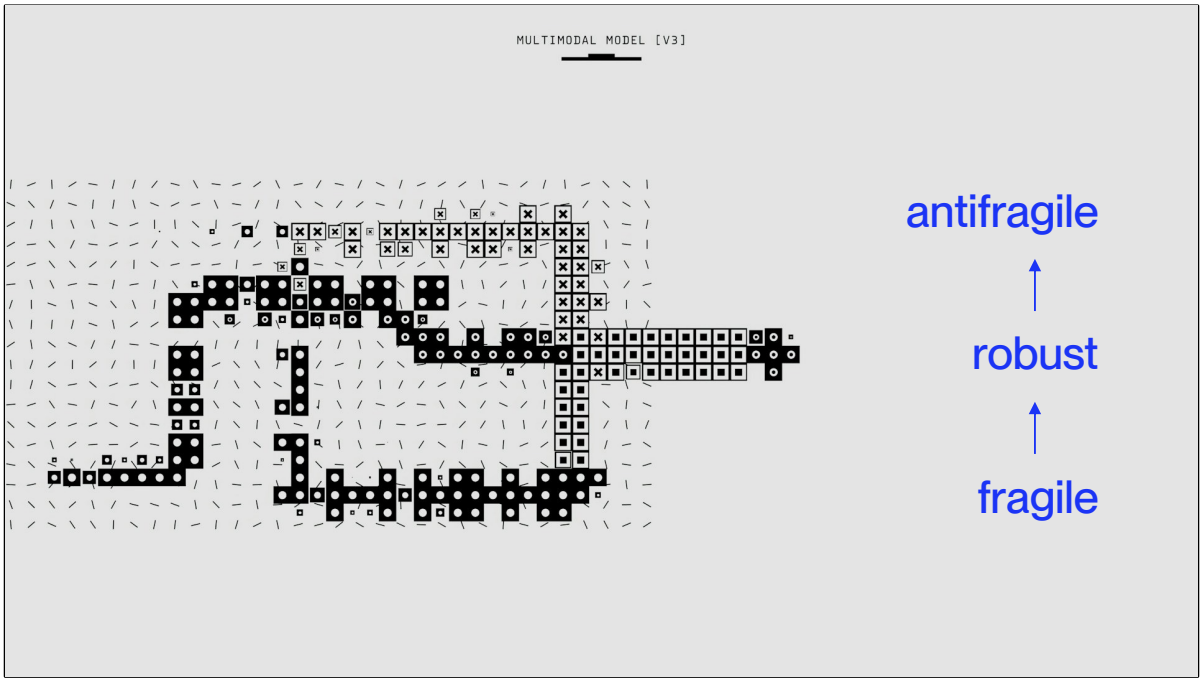

The real advantage lies in building what Nassim Nicholas Taleb has called *antifragile systems* (Taleb, 2012). Fragile systems break under stress. Robust systems survive stress. Antifragile systems are different: they grow stronger because of it. They learn, adapt, and turn disruption into a source of improvement.

An example comes from Toyota. Any employee on the production line has the authority to stop the entire line if they see a problem. A fragile system would let defects accumulate until it breaks. A robust system might absorb the error and continue. Toyota's antifragile system goes further: it uses disruption to learn and improve, so every stoppage makes the system stronger. That is what it means to turn stress into advantage.

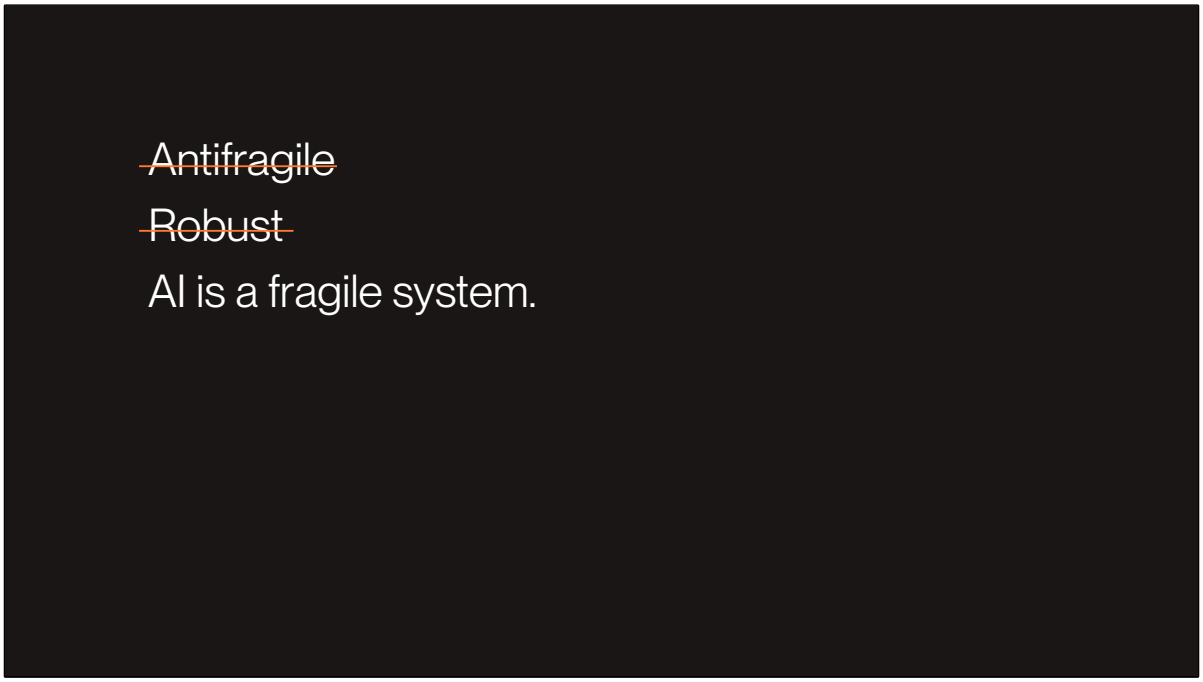

Which brings me to my simple thesis: AI is not an antifragile system, and nor is it robust. AI is, by its nature, a fragile system.

And that raises the most important question for every organisation: **why are we using it?**

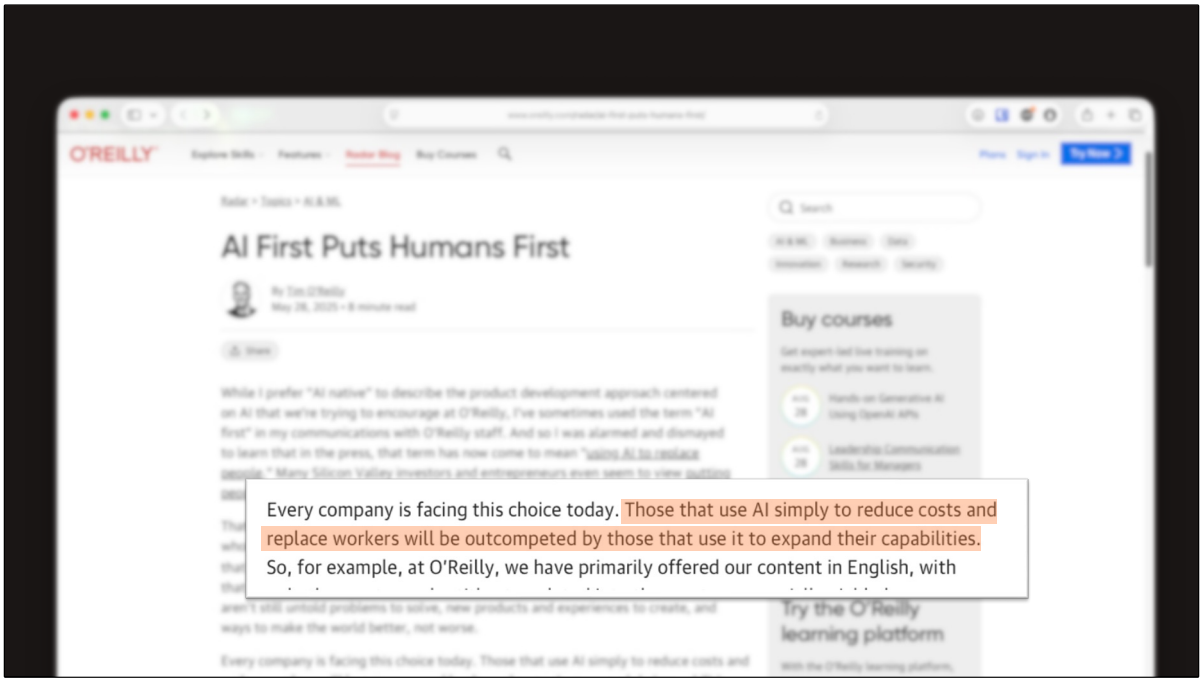

As Tim O'Reilly put it, "Every company is facing this choice today. Those that use AI simply to reduce costs and replace workers will be outcompeted by those that use it to expand their capabilities" (O'Reilly Radar, 2024). The organisations that thrive will be those that use AI to amplify human capability, not diminish it.

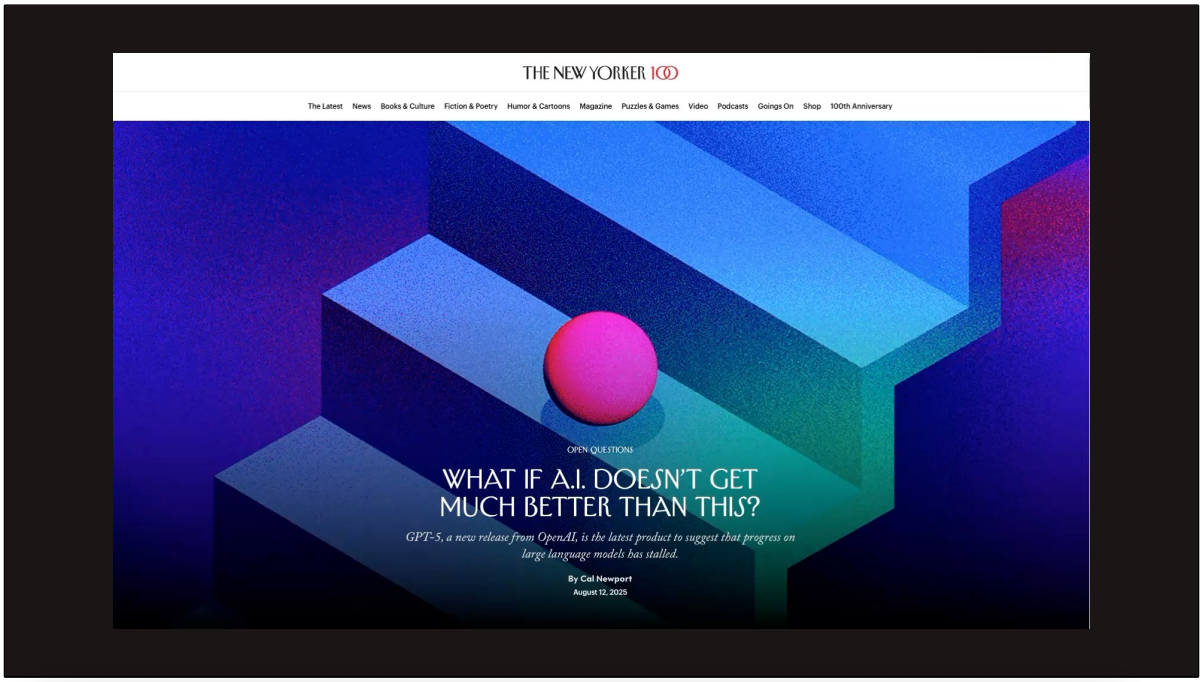
THE NEW YORKER 100
The Latest News Books & Culture Fiction & Poetry Humor & Cartoons Magazine Puzzles & Games Video Podcasts Goings On Shop 100th Anniversary
OPEN QUESTIONS
WHAT IF A.I. DOESN'T GET MUCH BETTER THAN THIS?
GPT-5, a new release from OpenAI, is the latest product to suggest that progress on large language models has stalled.
By Cal Newport
August 12, 2025

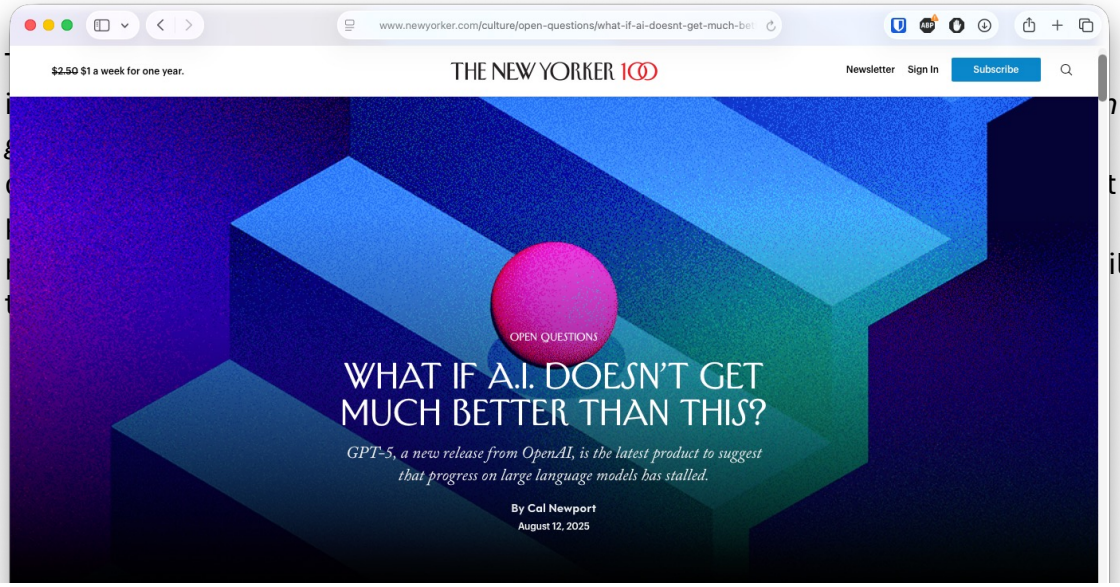
www.newyorker.com/culture/open-questions/what-if-ai-doesnt-get-much-bet
$2.50 $1 a week for one year.
THE NEW YORKER 100
Newsletter Sign In Subscribe
OPEN QUESTIONS
WHAT IF A.I. DOESN'T GET MUCH BETTER THAN THIS?
GPT-5, a new release from OpenAI, is the latest product to suggest that progress on large language models has stalled.
By Cal Newport
August 12, 2025

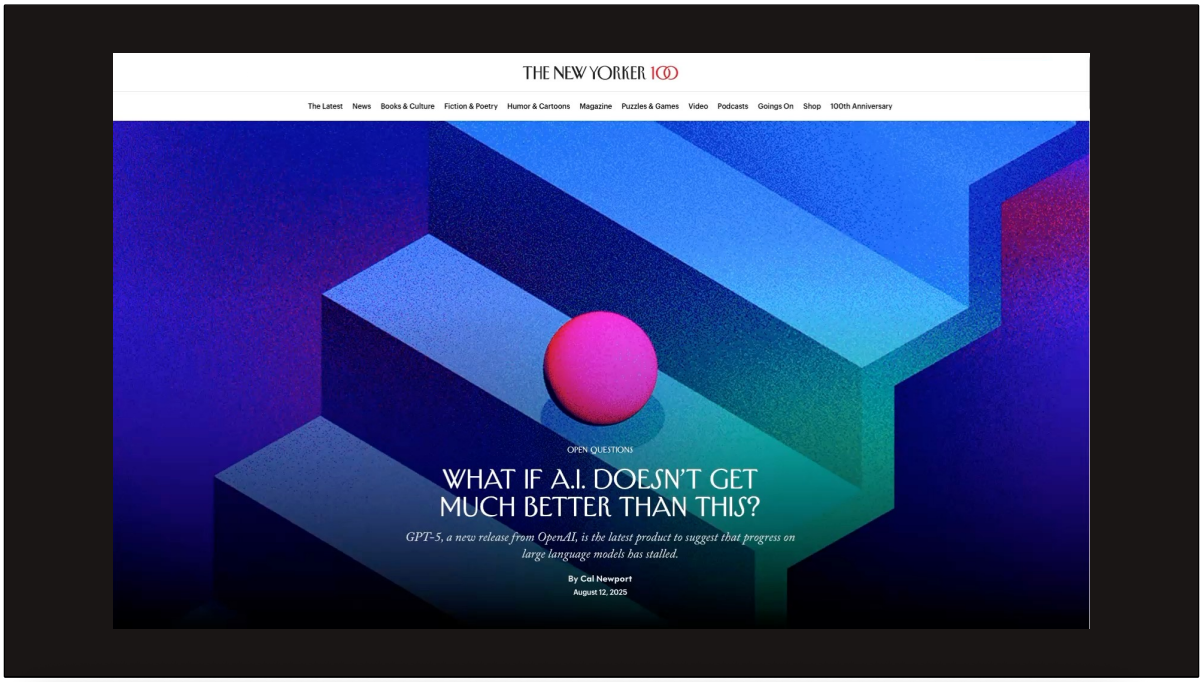
THE NEW YORKER 100
The Latest News Books & Culture Fiction & Poetry Humor & Cartoons Magazine Puzzles & Games Video Podcasts Goings On Shop 100th Anniversary
OPEN QUESTIONS
WHAT IF A.I. DOESN'T GET MUCH BETTER THAN THIS?
GPT-5, a new release from OpenAI, is the latest product to suggest that progress on large language models has stalled.
By Cal Newport
August 12, 2025

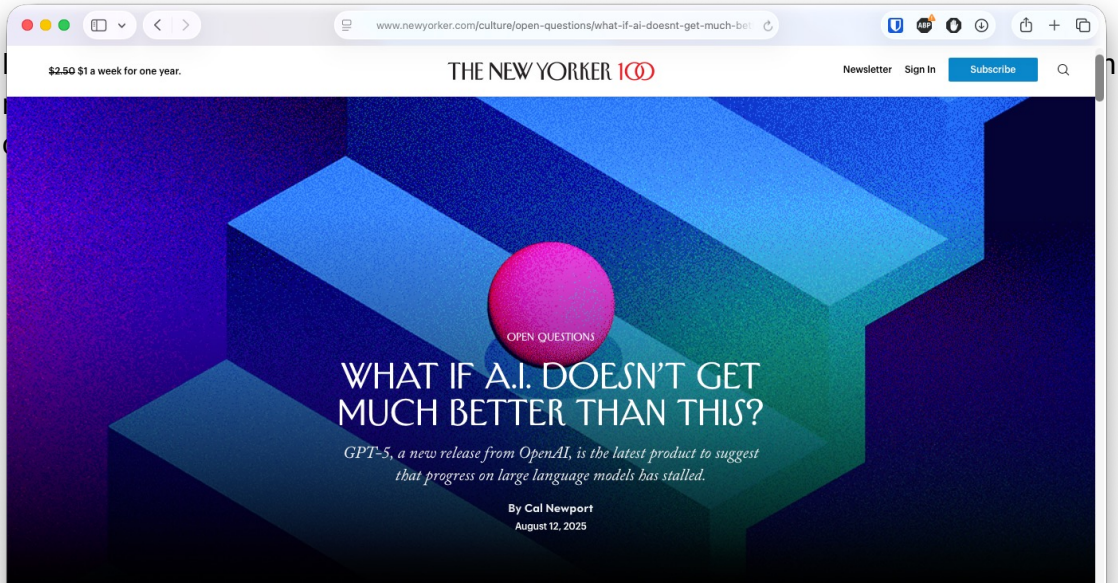
www.newyorker.com/culture/open-questions/what-if-ai-doesnt-get-much-bet
$2.50 $1 a week for one year.
THE NEW YORKER 100
Newsletter Sign In Subscribe
OPEN QUESTIONS
WHAT IF A.I. DOESN'T GET MUCH BETTER THAN THIS?
GPT-5, a new release from OpenAI, is the latest product to suggest that progress on large language models has stalled.
By Cal Newport
August 12, 2025

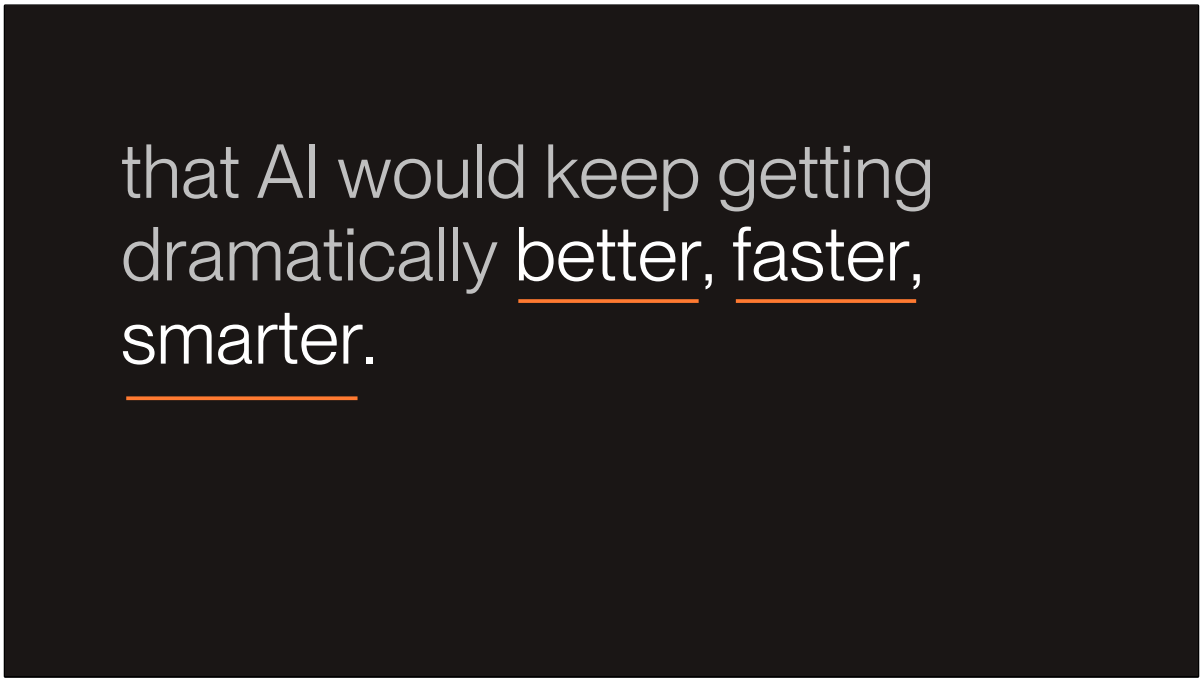

that AI would keep getting dramatically better, faster, smarter.

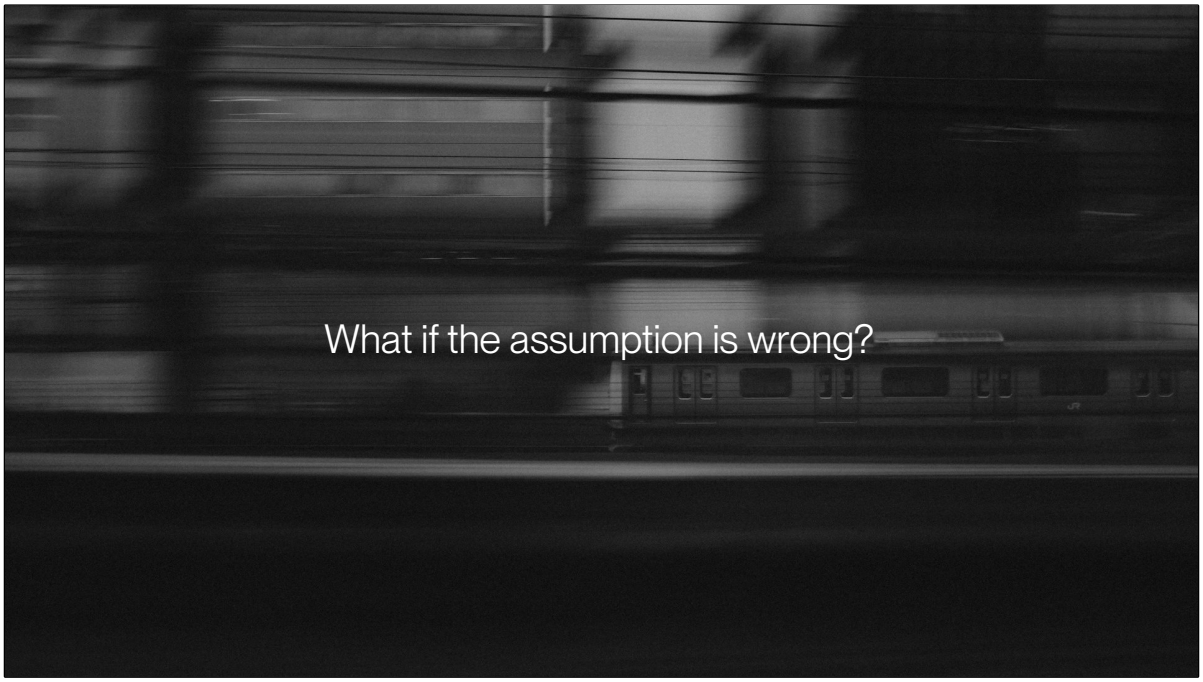

But let’s pause for a moment and ask: what if that assumption is wrong? What if the AI capabilities we have today are roughly what we will have five years from now? Ten years from now? Forever?

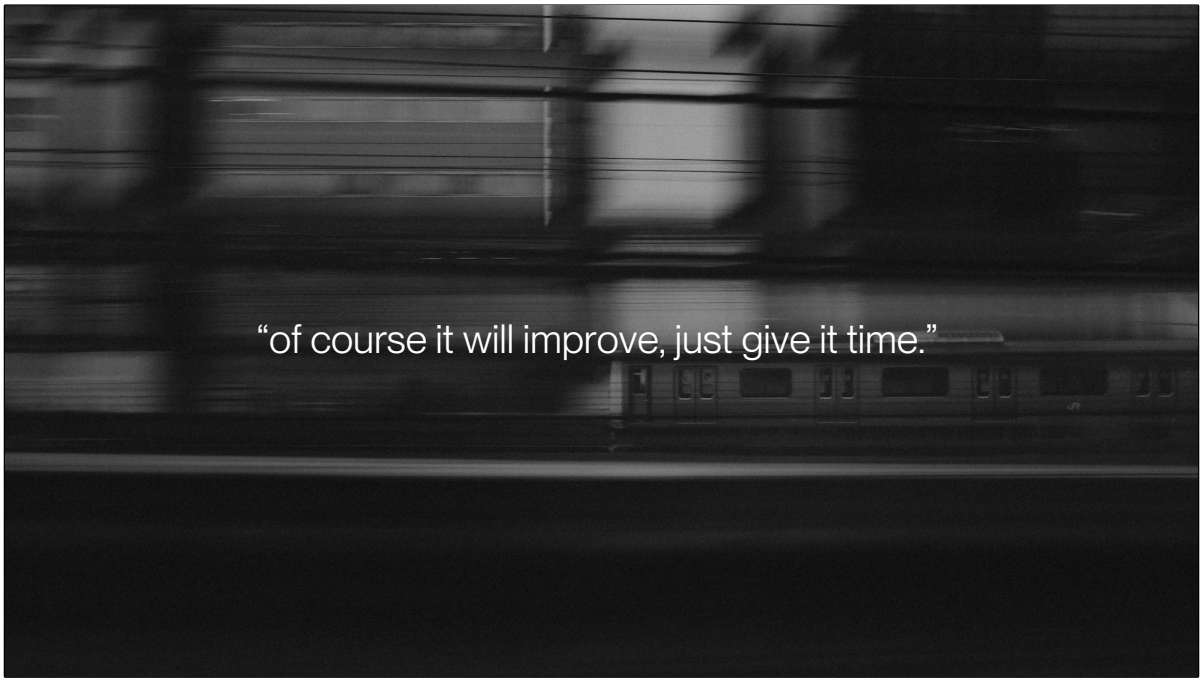

Some of you might well be thinking right now, “*of course it will improve, just give it time*.”

And maybe it will. But as a futurist and rational optimist, I find it more prudent to assume the opposite: to plan as though today’s capabilities are as good as it gets.

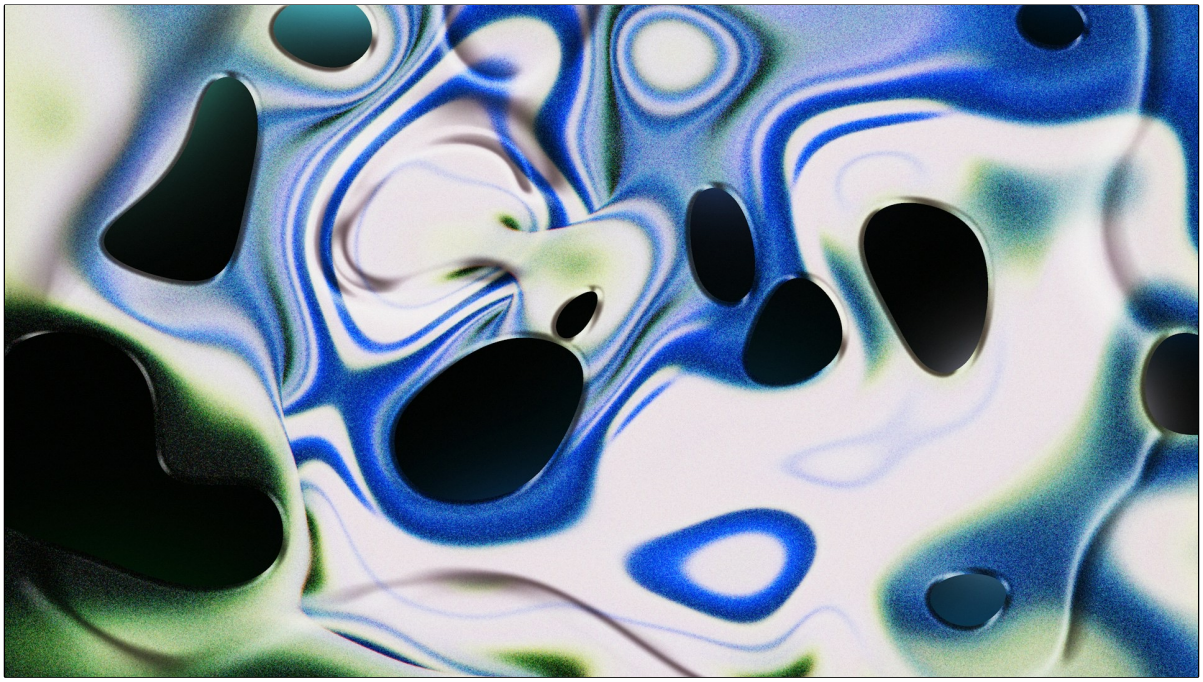

If I am wrong, I'll be pleasantly surprised. If I am right, I'll be ahead of the game.

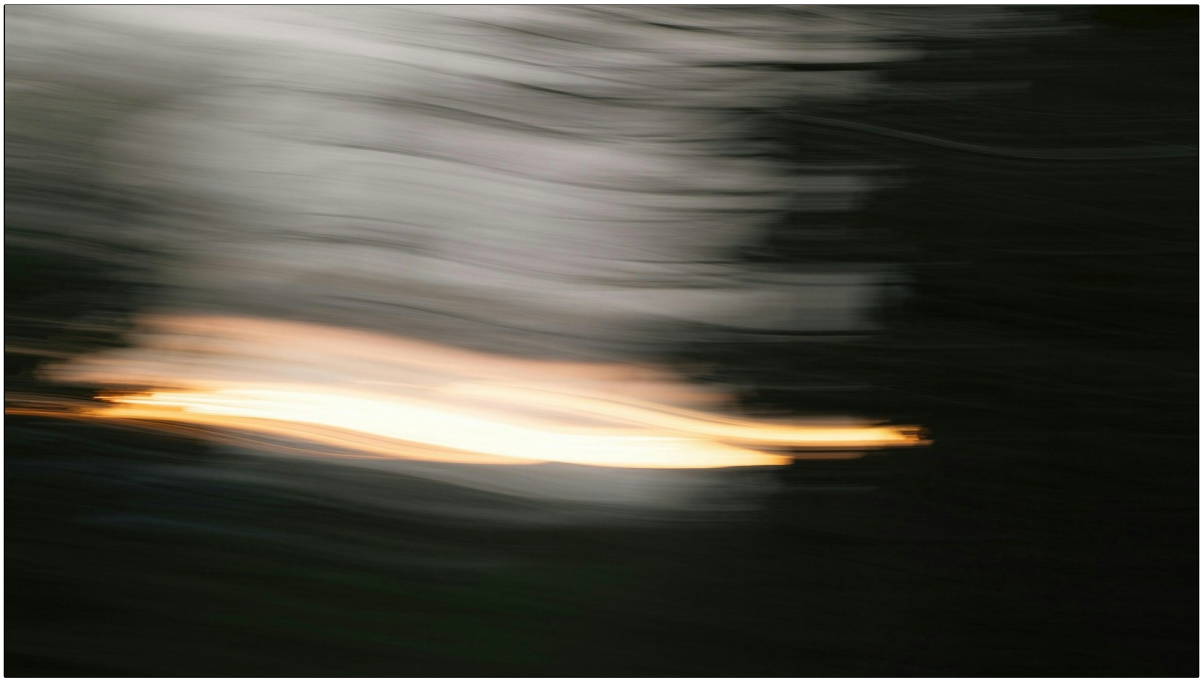

This kind of scepticism is healthy. It protects us from illusions, keeps us from buying into manufactured urgency, and prevents us from mistaking marketing for reality, and being prepared for realistic futures.

But I hear you, scepticism alone doesn't build anything. It is like a smoke detector: it warns of danger but does not make the house stronger.

Trust

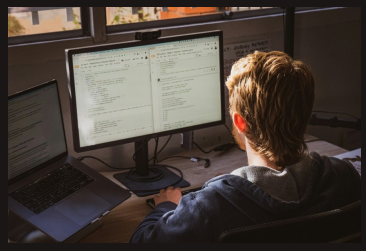

Contextual intelligence

Human insight

So let us turn to what *does* build strength.

The organisations that will thrive, whether AI accelerates or not, will be those that invest in three enduring foundations: **trust, contextual intelligence, and human insight.** These are what transform AI from a fragile tool into part of a system that grows stronger under stress.

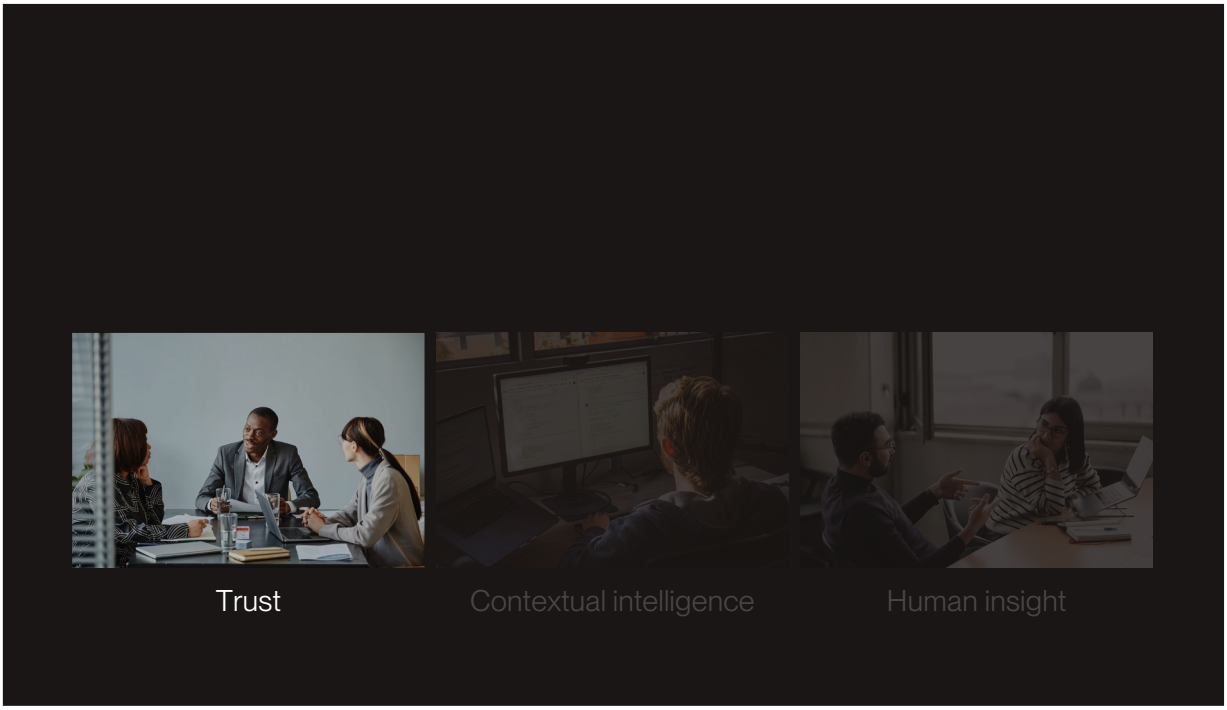

The companies betting everything on faster, or just more, AI are fragile. Those investing in trust are positioning themselves to win regardless of what happens with the technology. Trust is what makes any system usable at scale. Without it, no technology endures, however sophisticated it may be.

A recent field experiment, *My Advisor, Her AI and Me* (HEC Paris and Goethe University Frankfurt, 2025), illustrates this vividly. Researchers partnered with a large German savings bank to compare pure AI-generated financial advice with a human–AI collaborative model, where bankers had the final say.

They found that human involvement changed the advice but did not reduce its quality. What it did change was customer behaviour: people were far more likely to follow recommendations when a human was in the loop, especially for riskier decisions.

And that mattered, because greater follow-through led to significantly better financial outcomes, with realised payoffs around 45 percent higher than under pure AI advice.

The mechanism was not accuracy. It was trust. Customers interpreted the presence of a human as a social and emotional assurance, a signal of accountability.

In other words, the value of human oversight was not in making the AI smarter, but in making people confident enough to act.

And this is the broader pattern: people often say they trust AI in principle, but hesitate to rely on it without human accountability. Adoption doesn't fail because the models are weak — it fails because the trust equation hasn't been solved. Deloitte's Trustworthy AI framework makes this explicit: customer-facing systems require a "human in the loop" with the ability to challenge and override outputs (Deloitte, 2024).

So, customers are not rejecting technology; they are demanding that it earn their trust through visible human validation. And this creates stability: consistent usage produces reliable data, which improves AI performance, which builds more trust. It is a reinforcing cycle that can make systems stronger over time.

Here's what's particularly interesting: while we have been debating whether AI will replace human judgment, regulators are insisting on the same principle. In REP 798: Beware the Gap, ASIC found that many licensees lacked meaningful human oversight.

It warned that governance must include genuine human accountability throughout the AI lifecycle (ASIC, 2024).

The message is clear: trust requires human accountability. And trust is what transforms AI from an experiment into business infrastructure.

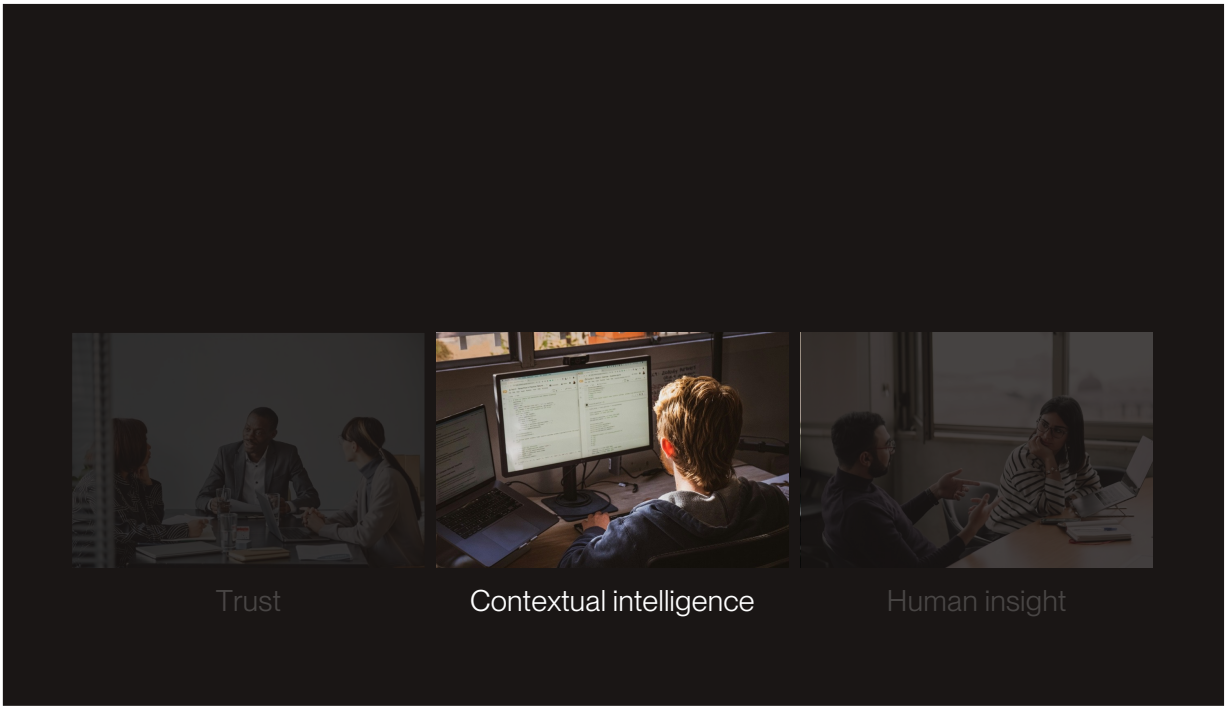

Human contextual intelligence is what enables organisations to adapt when algorithms miss the nuance of real situations. AI excels at recognising patterns; humans excel at understanding what those patterns mean in context.

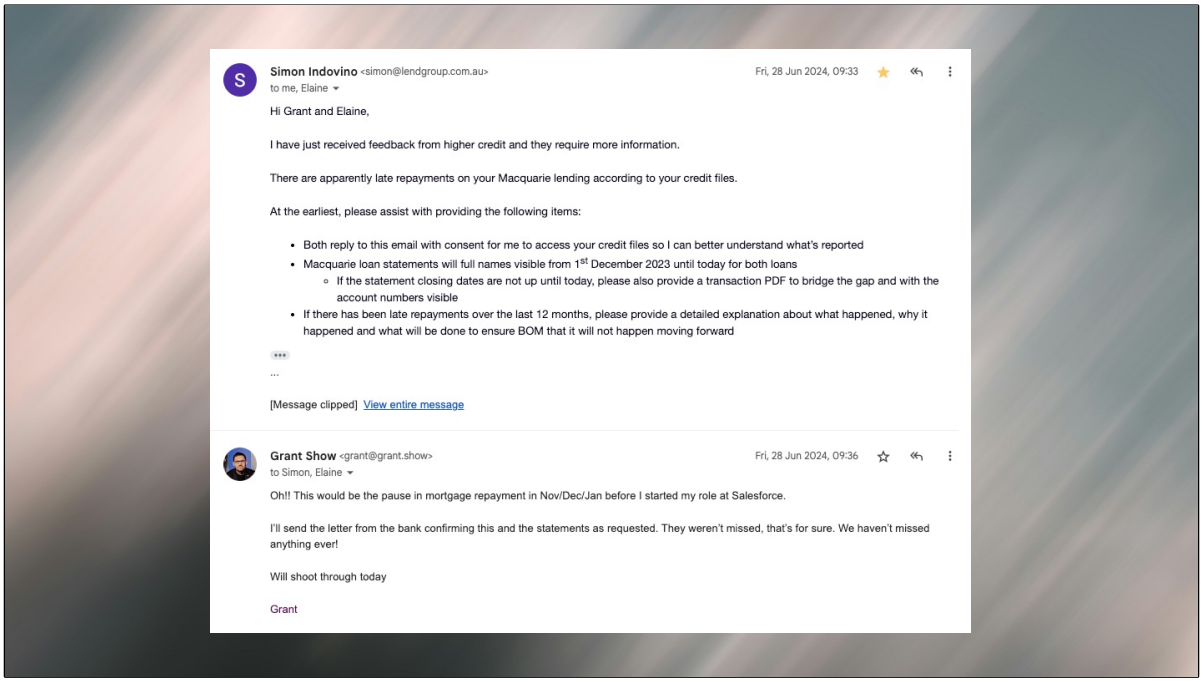

Let me share a personal story. My wife and I bought our home in 2021 at an outstanding interest rate of 1.9%. Glorious. In 2024 our fixed rate term was coming to and end.

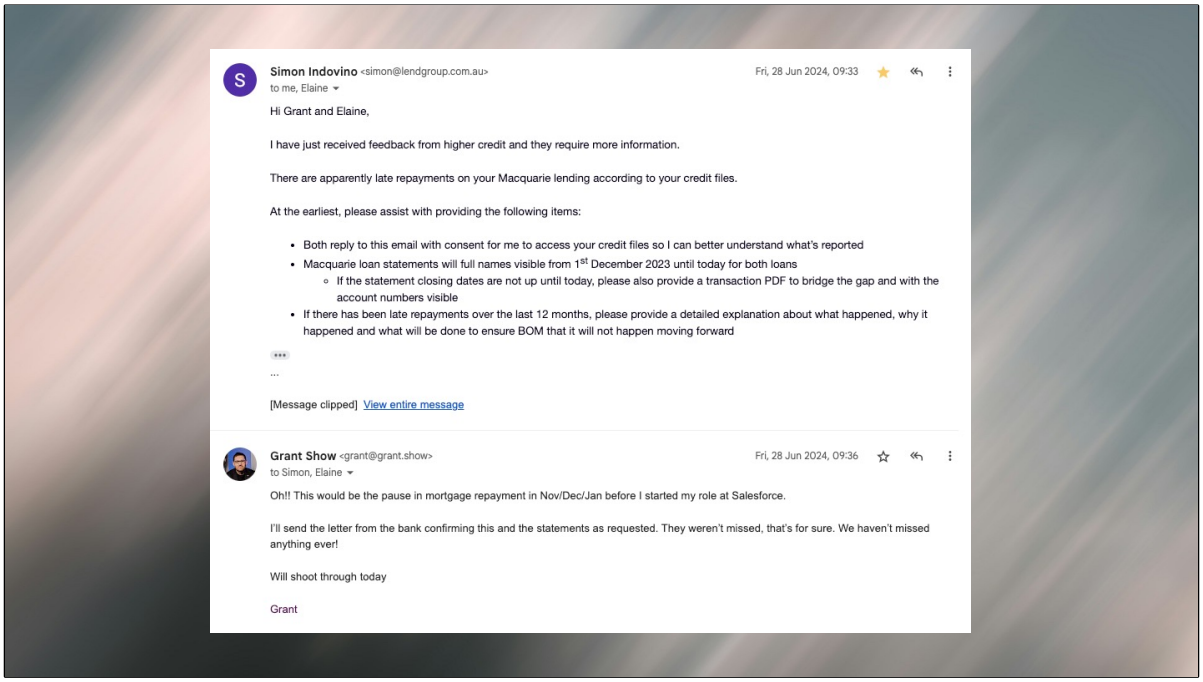

Simon Indovino <simon@lendgroup.com.au> Fri, 28 Jun 2024, 09:33
to me, Elaine

Hi Grant and Elaine,

I have just received feedback from higher credit and they require more information.

There are apparently late repayments on your Macquarie lending according to your credit files.

At the earliest, please assist with providing the following items:

- Both reply to this email with consent for me to access your credit files so I can better understand what's reported
- Macquarie loan statements will full names visible from 1st December 2023 until today for both loans
  - If the statement closing dates are not up until today, please also provide a transaction PDF to bridge the gap and with the account numbers visible
- If there has been late repayments over the last 12 months, please provide a detailed explanation about what happened, why it happened and what will be done to ensure BOM that it will not happen moving forward

...

[Message clipped] View entire message

Grant Show <grant@grant.show> Fri, 28 Jun 2024, 09:36
to Simon, Elaine

Oh!! This would be the pause in mortgage repayment in Nov/Dec/Jan before I started my role at Salesforce.

I'll send the letter from the bank confirming this and the statements as requested. They weren't missed, that's for sure. We haven't missed anything ever!

Will shoot through today

Grant

When I applied for a home loan in 2024, I was rejected because the technology flagged missed repayments. The technology in whatever form saw risk and recommended decline. But when escalated and a human reviewed the case, they saw something the system could not:

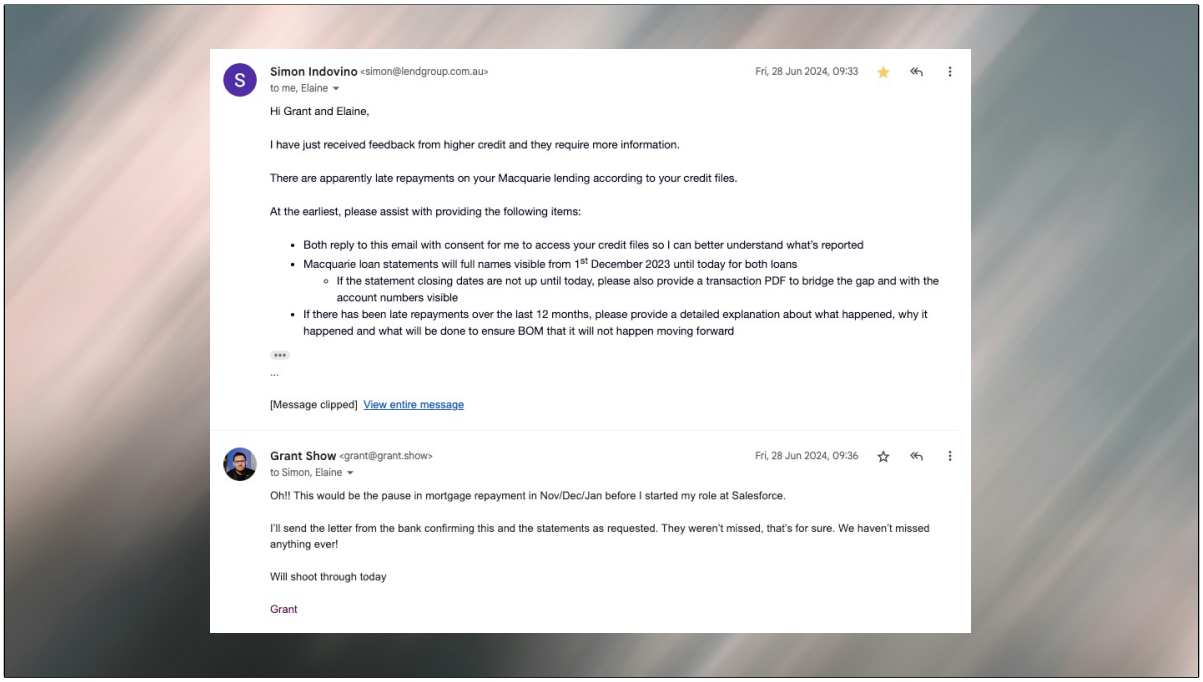

I had paused repayments during the pandemic while under employed, using a feature designed to protect both borrower and lender.

They could see that we had money in the bank, no signs of financial distress, and that we could easily make repayments. More than that, they appreciated that I was making intelligent financial decisions at the time. Something to be celebrated, not punished.

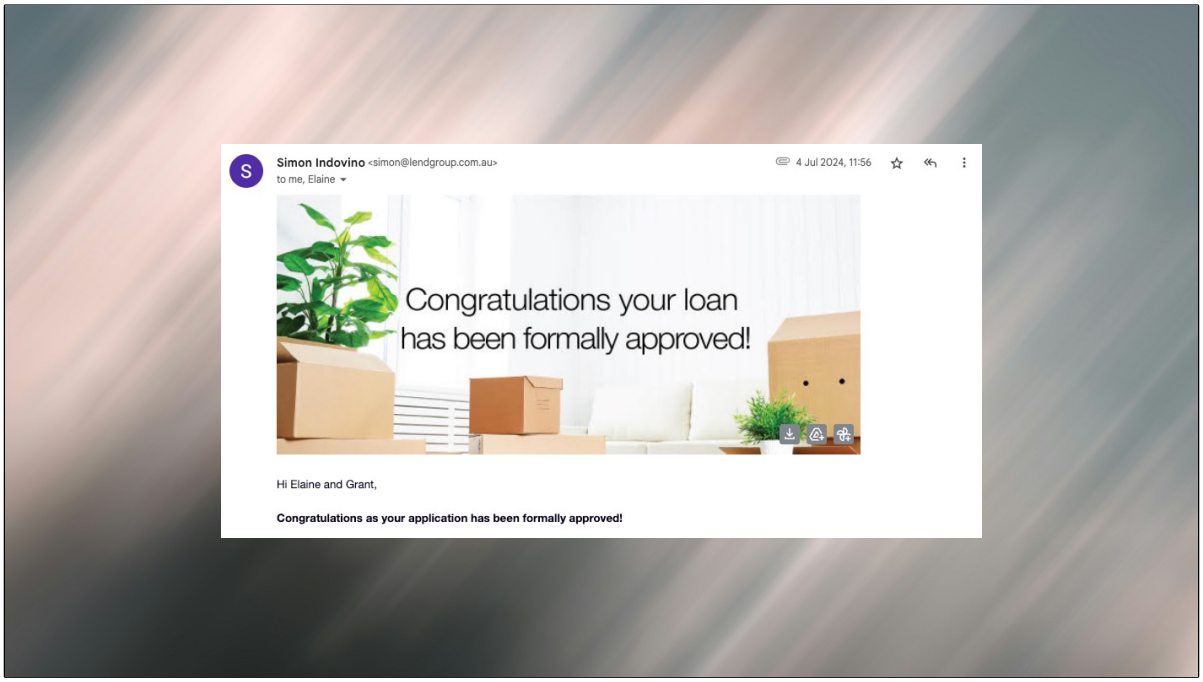

Our loan was approved.

This wasn't just about a loan. It revealed how technology can misinterpret responsible behaviour as risk, while human contextual intelligence recognises the difference between financial distress and good judgement.

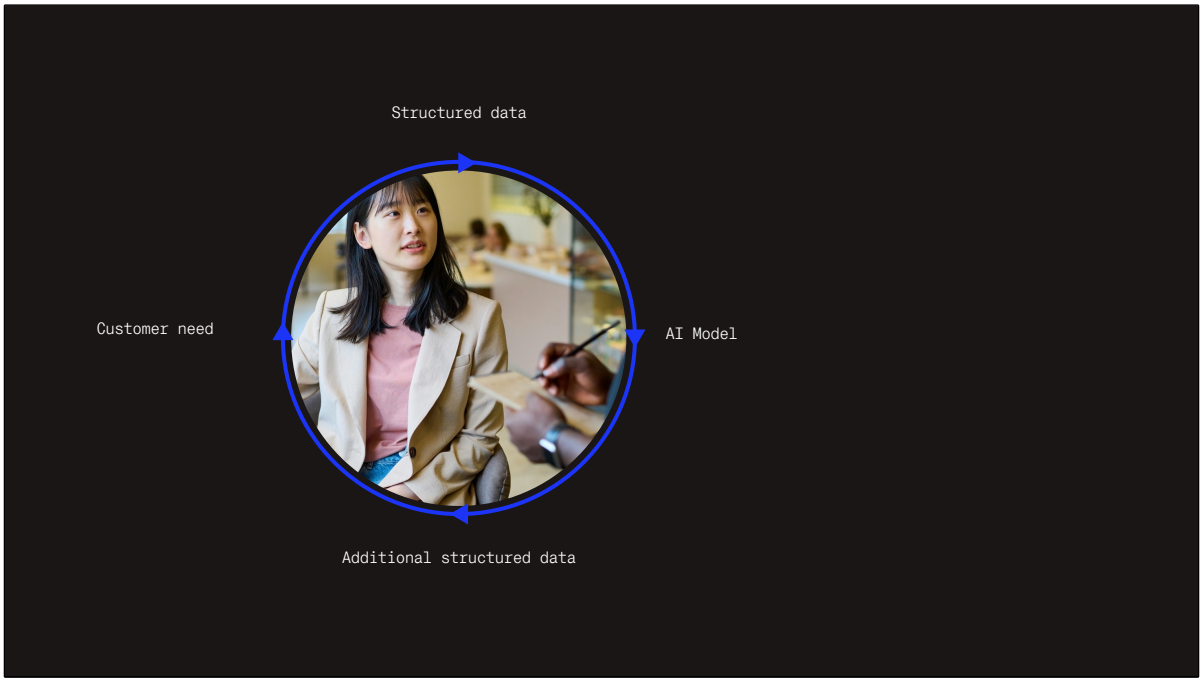

Most AI today is designed for stable contexts.

A customer need gets translated into structured data, processed by a model, and generates more structured data. In predictable environments, that’s fine. But it assumes tomorrow will look like yesterday. Neat, predictable, but fragile.

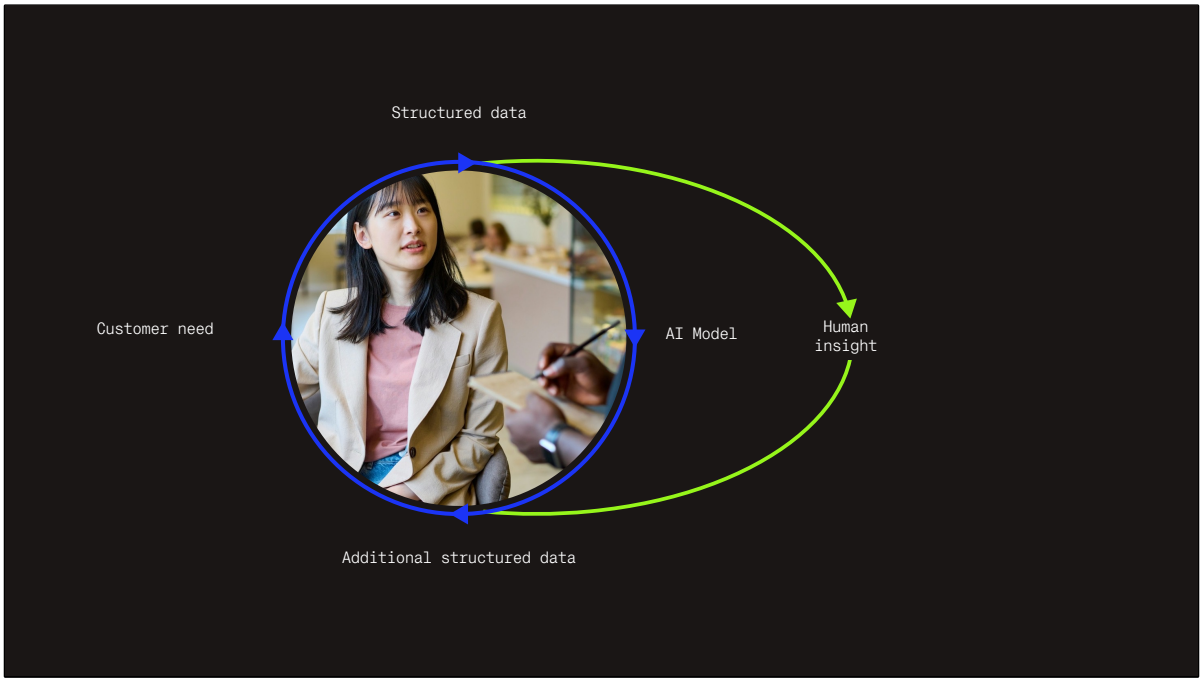

Real life is not stable.

Customer needs shift. Data loses predictive power. Models fail. What makes a system resilient is human judgement: the context that interprets, refines, and strengthens outcomes. That turns a fragile pipeline into a learning loop, centred on customer outcomes.

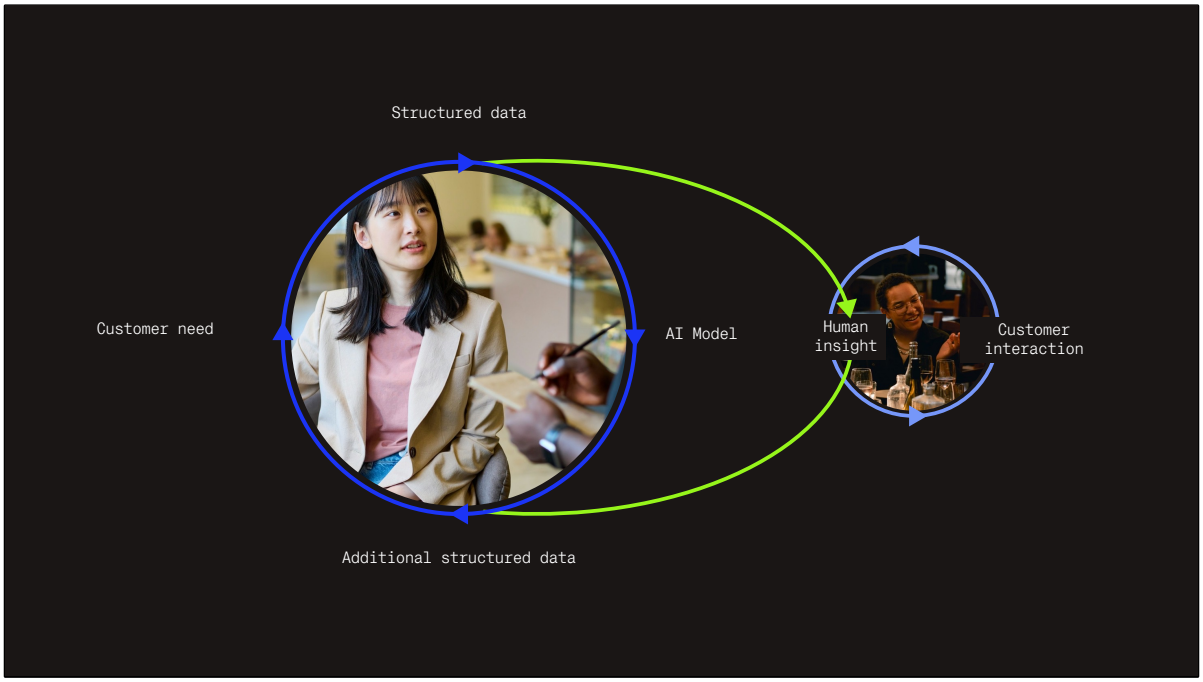

And the loop goes even further. Human insight doesn't just interpret data; it comes from direct interaction with customers. Every conversation, every complaint, every request adds nuance that AI alone cannot capture. When that context feeds back into the system, both the AI and the organisation grow stronger.

This is the flywheel of contextual intelligence:

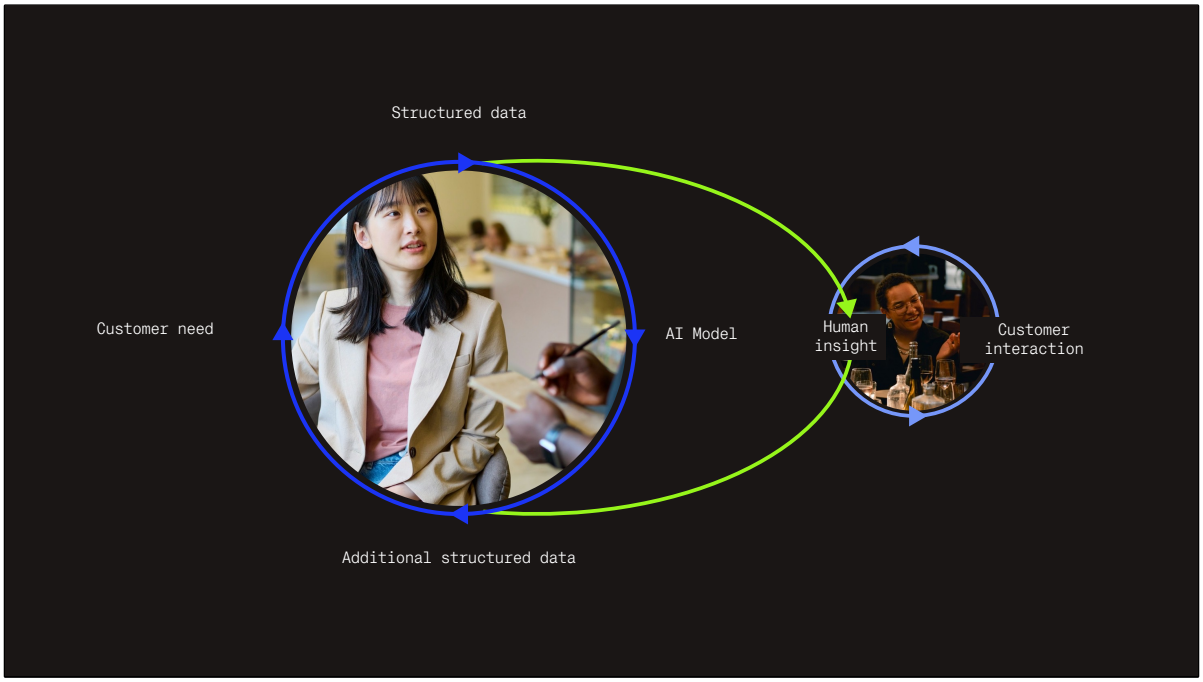

1. Customer needs generate data.
2. AI recognises patterns.
3. Humans co-interpret and refine those patterns.
4. Customer interactions add further insight.
5. That insight feeds back into the system.

The more this loop turns, the stronger the organisation becomes. Fragile pipelines give way to resilient systems. And what drives that resilience is not just data, it is context and customer interaction.

And there is another effect. Freeing up your team to focus on customers doesn't just deliver better commercial outcomes, it also creates more fulfilling roles. Happier, more engaged employees make for stronger organisations.

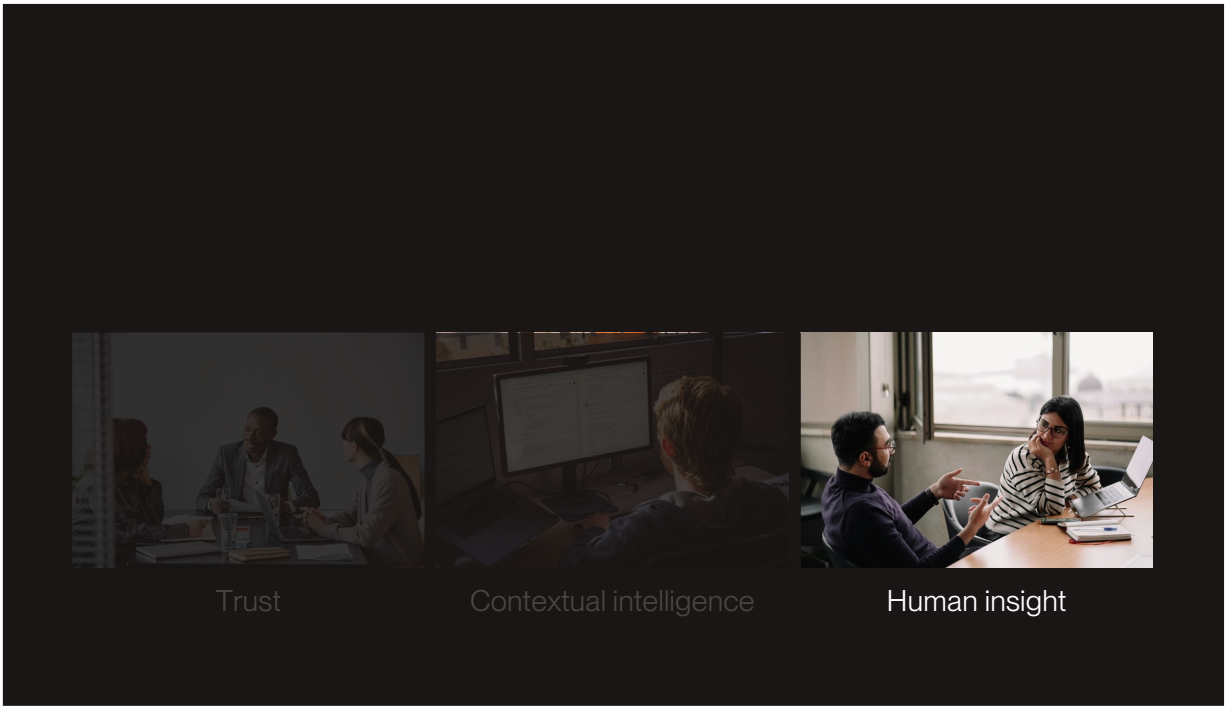

Learning from customers, however, is not simple. Customers do not speak in numbers alone.

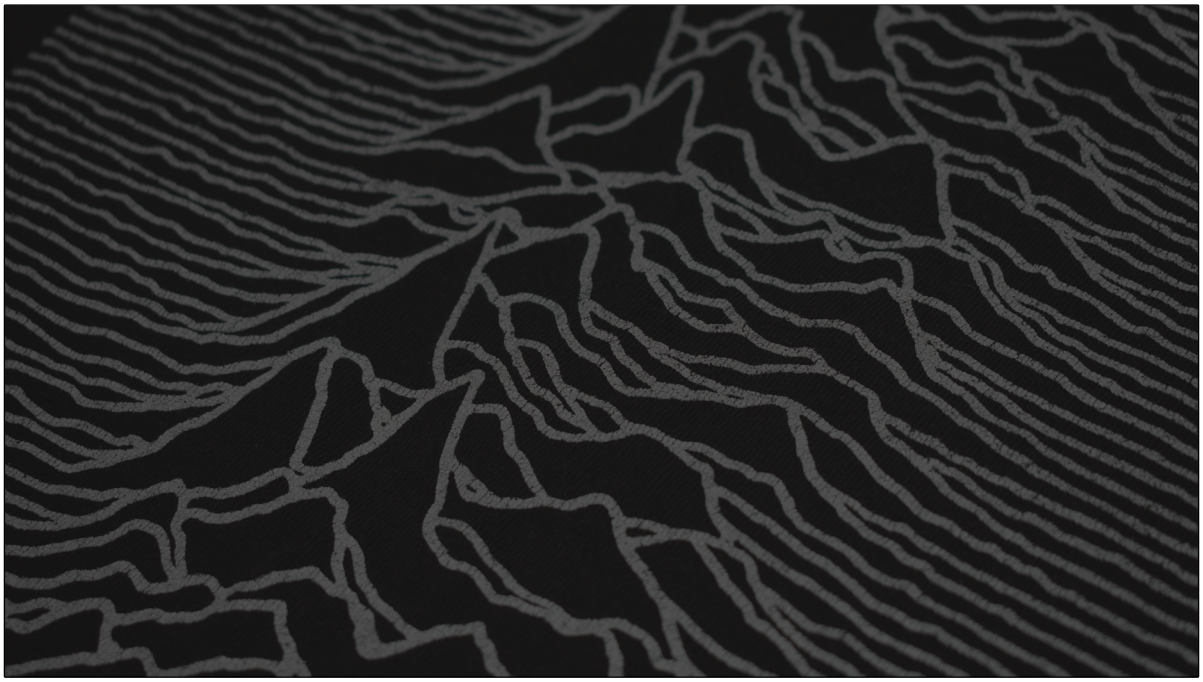

They speak in stories, rich in context, complexity, and mistakes. Every conversation, every complaint, every request is qualitative data that only human insight can truly interpret. And when those insights are fed back into your systems, they strengthen both the AI and the organisation.

Human insight is the true strategic advantage.

It enables organisations not just to identify new patterns but to understand what those patterns mean. AI excels at spotting trends at scale, but human insight interprets them with judgment, trust, and context.

Consider what would happen if a pandemic-level disruption hit tomorrow.

Every AI model trained on current data would become unreliable.

Historical spending patterns, risk profiles, and customer behaviours would all shift overnight. The data our AI depends on would lose its predictive power.

The organisations that would adapt fastest aren't those with the most sophisticated AI. They will be those with the closest relationships to their customers – in a sort of mutualism based relationship.

Teams that combine qualitative signals such as conversations, feedback, and frontline experience, alongside quantitative data, will understand not only what is happening – but why.

This is where organisations become antifragile. They do not just adapt to disruption, they use it to build stronger capabilities. Customer service representatives become the frontline source of intelligence, explaining behaviours that the data alone cannot. Risk analysts who once relied on historical models bring in real-time judgement.

Disruption reveals the value of human insight that strengthens the organisation.

Fraud is one of the clearest examples. Every new scam, every failed attack, adds intelligence to detection systems.

The system gets better not in spite of adversaries, but because of them. Human investigators bring context to those signals, feeding them back into AI, which makes the whole network stronger. That is antifragility in action.

Here is the crucial point. AI can identify patterns and anomalies, but it cannot inspire confidence or guide decisions in shifting contexts. Human insight brings the judgment and trust that turn information into action.

And when paired with AI at scale, that combination becomes a tour de force.

This creates resilience that grows stronger through volatility. Each disruption becomes a source of advantage instead of a threat. When everyone else's AI models falter, human insight combined with AI grows stronger and more valuable.

The regulatory environment reinforces this reality. Requirements for human oversight are not constraints, they are forcing functions that make organisations better.

While competitors chase autonomous AI systems that become brittle under stress, regulated institutions are developing human–AI collaboration that thrives under pressure and becomes more valuable in uncertainty.

So, how do we start to build the right foundations? It isn’t a choice of investins in either AI or human capabilities, we must invest in both. And what I propose, is investing disproportionally across three areas:

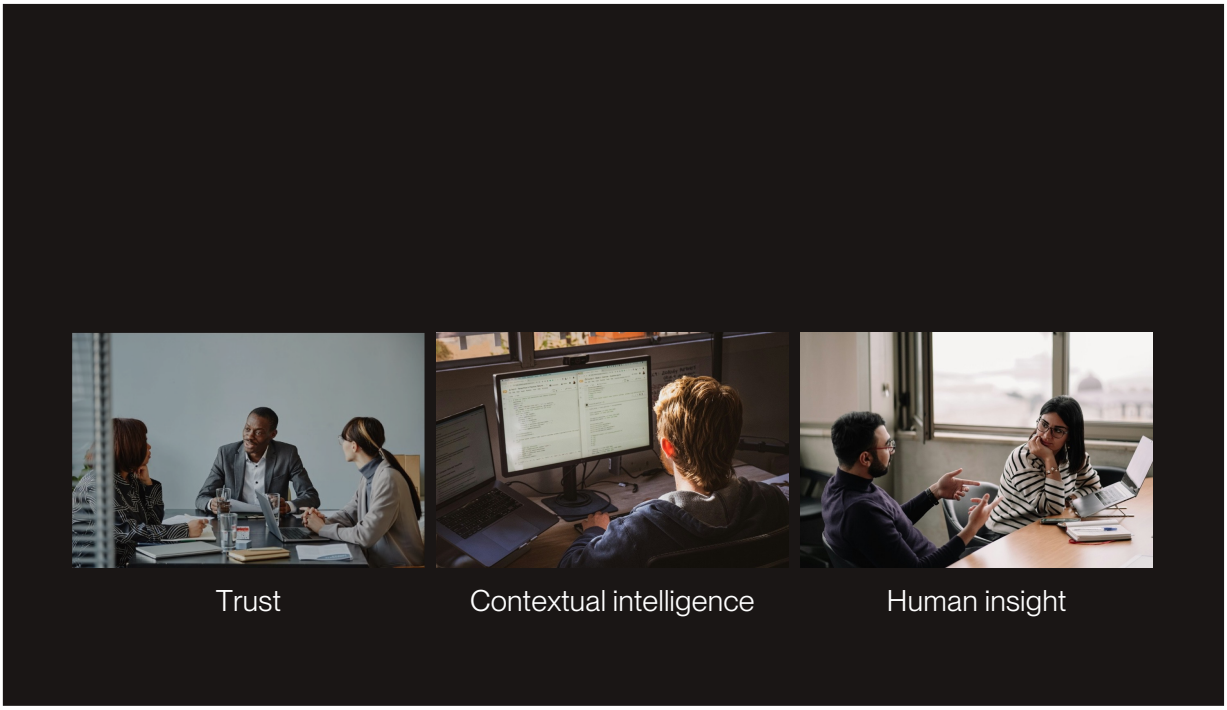

Trust, contextual intelligence, and human insight. These are what turn AI from a cost-saving tool into a capability-expanding force.

The most successful approaches I have seen follow an integrated strategy that assumes uncertainty is permanent.

Trust

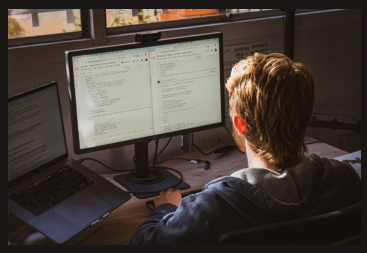

Contextual intelligence

Human insight

The key is that this is architectural: you are not bolting human failsafes onto AI systems, you are building living organisations that grow stronger through uncertainty, using AI as a force multiplier.

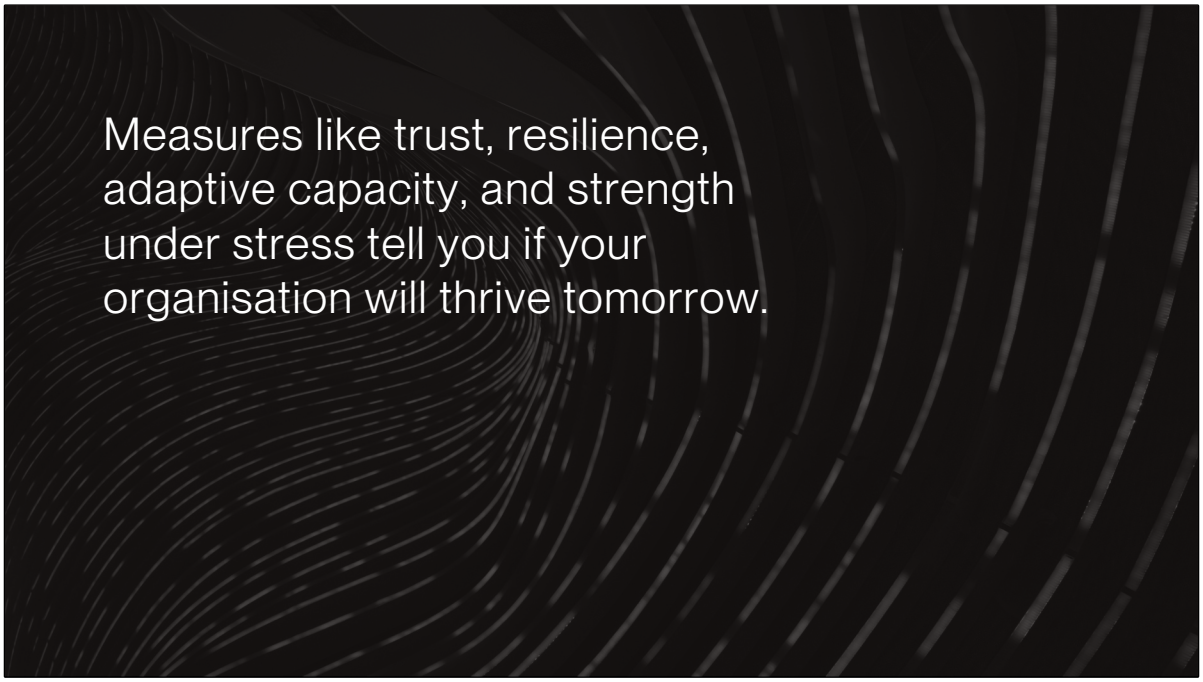

The first place to start is in developing new success metrics. Traditional measures like accuracy, speed, and cost reduction can tell you if your system works today. New measures like trust, resilience, adaptive capacity, and strength under stress tell you if your organisation will thrive tomorrow.

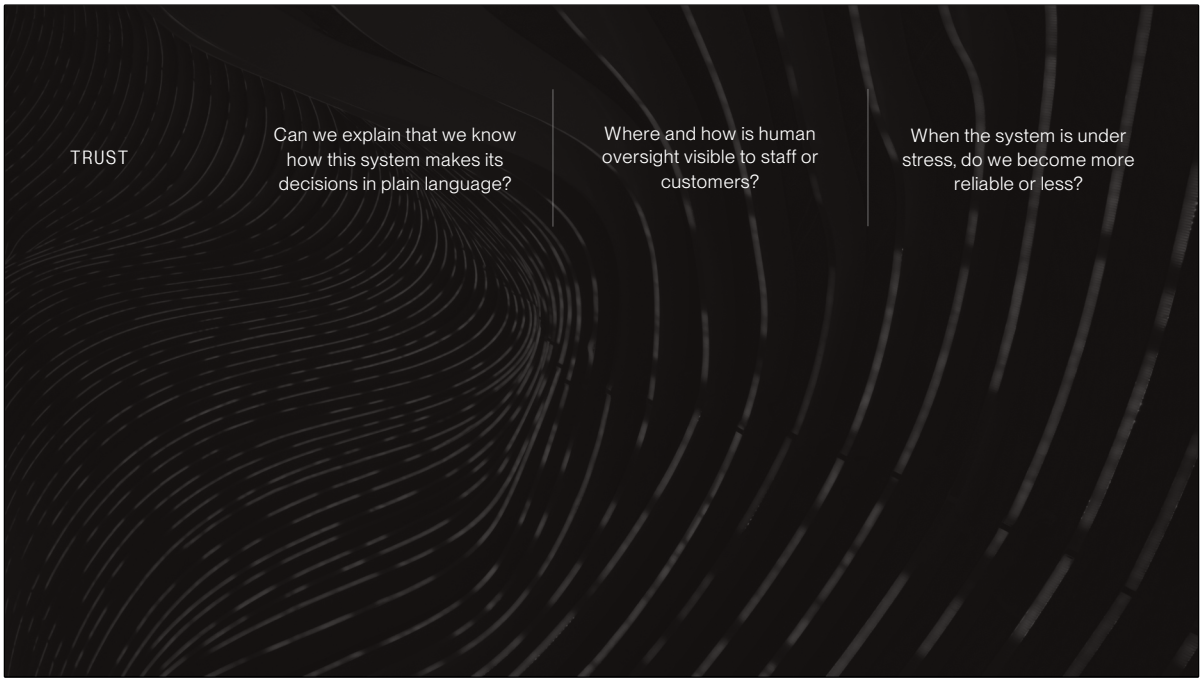

Trust turns technology into something people actually use. Systems need to be explainable, with human oversight visible, so customers know they are not at the mercy of a black box.

Contextual intelligence moves beyond patterns to meaning. Algorithms can spot correlations, but only humans can interpret intent and nuance.

Human insight allows organisations to treat shocks as information rather than threats, and to design decision-making processes that improve under pressure.

Customer experiences should not collapse in stress but get better because of it.

And that brings us back to where we began. The illusion.

In business strategy, we face a similar illusion: that sophisticated technology automatically creates competitive advantage, that AI advancement guarantees success.

It does not.

Enduring advantage is created through three foundations: trust that anchors stability in chaos, contextual intelligence that enables adaptability in change, and human insight that turns disruption into strength.

Winning organisations will not be those betting on breakthroughs alone. They will be those building capabilities that make them stronger under pressure.

And so, I invite you to see that because AI is a fragile system

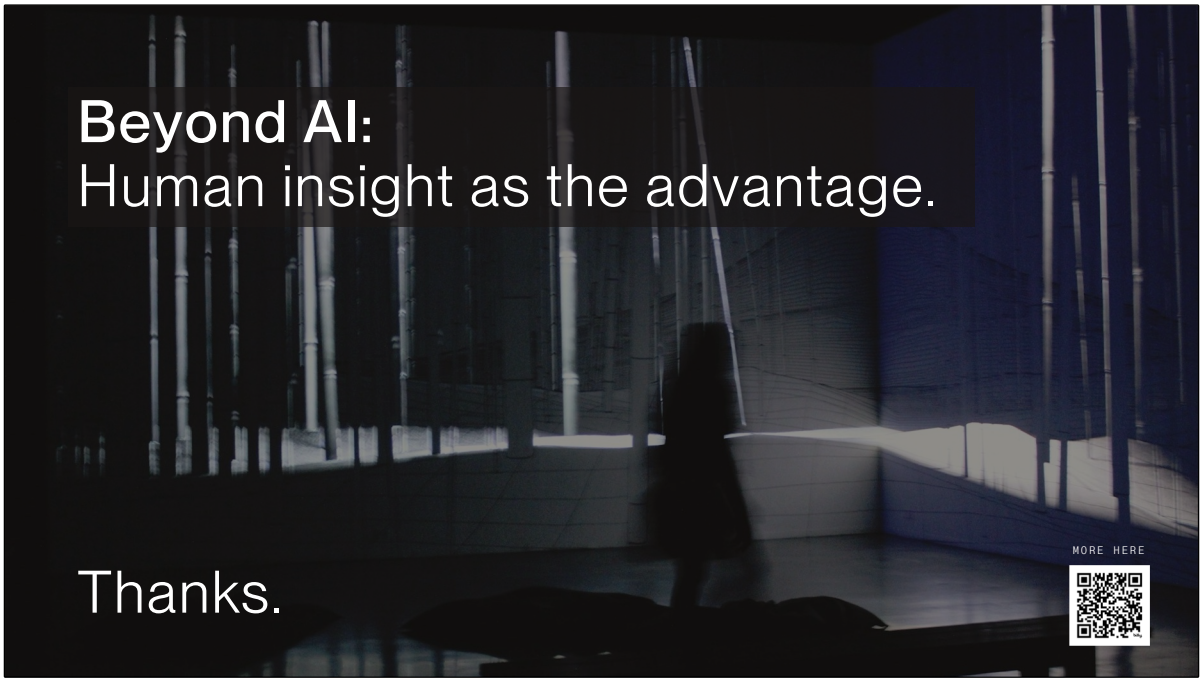

The real advantage lies beyond AI in human insight.