ARE WE IN AN AI BUBBLE?

A WHITEPAPER FOR BUSINESS LEADERS

Accrued Intentions

What the AI boom has promised to pay

I built a 100% AI-run accountancy firm.
Here's why I think AI is a bubble.

Alexis Kingsbury
Founder of Accrual Intentions, the world's first 100% AI-staffed accountancy practice
Author of *The Accountancy Playbooks* and *Accrual Intentions*

4 November 2026

What's inside

Before you start
My answer in brief, why I'm asking, the best case against me, and how to read the numbers.

01 **What I mean by "bubble"**
Dot-com versus telecoms, and the AI bubble that has already burst.

02 **The valuation test**
What has to be true for $1tn+ valuations to make sense.

03 **A simplified going-concern test**
Growth & Co., the funding gap, and how Anthropic compares with the rest.

04 **What's different this time**
Borrowed money, leases and guarantees, one type of firm at a time.

05 **Who owns the downside**
From data centre to pension fund: where the risk has landed.

06 **How a correction spreads, and what stops it**
What spreads losses, what stops them, and the risk of recession.

07 **The 2027–30 test**
When the promises come due.

08 **"But…"**
The challenges I get from the audience, and my answers.

09 **Why the "SaaS apocalypse" was overstated**
Who came through the software panic, and where the fear did bite.

10 **Where I think it lands**
A burst, a serious risk of recession, and what to do about it.

11 **What I'll be watching**
Six signals that would change my mind.

12 **Conclusion**
The technology can win and the money can still lose.

Appendix, glossary and sources
Every organisation in the paper, plain-English terms, and links to every source.

Research cut-off 4 November 2026. Not investment advice.

My answer, in brief

Short on time?

A note before you start: this paper is long, because I've tried to show my working. If you're short on time, drop it into your favourite AI tool and ask something like this:

> Using what you know about me and my work, summarise the points in this paper that matter most to me. Then suggest three questions about it that you think I'd find most useful, and let me pick one for you to answer.

That's also a neat test of the paper's argument: AI doing useful work for you at almost no cost.

When I speak at events about my "Accrual Intentions" experiment, someone in the audience often asks me whether we're in an AI bubble. My answer is: almost certainly, yes.

That doesn't mean AI is worthless. I think it will change how every accounting firm works. It means the prices being paid for AI companies, and the money being borrowed to build AI infrastructure, assume a future that goes almost perfectly, and quickly. I don't think it will be quick, and when the correction comes, I think it'll hurt more than most people expect.

This paper is my working. Each point uses the organisation that shows it most clearly, and I say why (the appendix lists every organisation and why it's there). Here's the argument in eight steps:

1. **Demand isn't the problem.** Anthropic's quarterly revenue went from $1.4bn to $17.0bn in a year, and the number of customers paying it over $100k a year went from more than 950 to more than 11,900.
2. **The prices assume years of near-perfect growth.** Anthropic's shares are expected to go on public sale soon, with talk of the whole company being valued at $1.5–2tn. To be worth that, my rough sums say its yearly revenue would need to grow by about 30–45% every year for five years, then stay highly profitable, and faster still once you allow for the return investors want for waiting. That's harder than it sounds, because AI prices keep falling, so it has to sell more each year just to stand still.
3. **Real businesses are adopting more slowly than the prices assume.** Uber used up its 2026 AI coding budget by April and capped spending; Meta, Accenture and others have rationed AI too. Anthropic's quarterly revenue grew by about $6.8bn in the spring and about $5.4bn in the summer.
4. **Even the best-placed company needs growth or new money to keep its promises.** Anthropic easily passes a simplified version of the test auditors use to check a business can pay its bills for the next 12 months (the going-concern test). If its revenue stopped growing, it would need

$27–114bn of new money by 2030. OpenAI, even on its own growth plan, reportedly needs roughly $150bn more than it has raised.

5. **The money is still there, but it's getting more expensive.** OpenAI's latest round reportedly raised about $11bn of a $30bn target. Lenders want around 16% a year to hold CoreWeave's bonds. Oracle is borrowing at higher rates and selling shares at a discount.
6. **The borrowing sits below the household names, and the shovel sellers are lending to the miners.** Data-centre developers, specialist funds and supplier guarantees carry the debt. Nvidia alone has disclosed support of up to $105bn, and its customers owe it $63bn.
7. **AI is eating some of the profits that pay for it.** Google's search queries grew at their fastest in five years in July–September, but its search revenue was flat on the previous quarter, partly because AI answers earn less.
8. **It has already happened somewhere.** China's main AI share index fell 40% in under four months this year.

> **Where I think it lands:** a bubble burst, with a serious risk that it tips the US into recession. The main route is households. Shares are a record 48% of American households' financial assets, and the richest 10% own about 88% of them. When their portfolios fall, they spend less, though economists disagree about how much less. If it burst today, I think the damage would mostly stay with the people who took the risk. My worry is about time: every month the music keeps playing, more money is borrowed and promised. If those promises keep growing faster than the cash to pay them, the shock will be bigger when it stops.

> **The good news:** a correction in AI infrastructure is good news for the businesses that use AI. Prices fall, suppliers compete, and value shifts to whoever owns the client relationship and the trusted data. If it does crash, businesses inherit more intelligence than they could ever have afforded, at a fraction of today's price.

> **If you run a business, the short version:** keep adopting AI, because the tools will get cheaper whatever happens to the companies making them. Avoid long contracts at today's prices. Keep your data and processes in systems you control. Spend the time AI frees up on your clients, because that's where the value will land. And plan for a tougher economy. Section 10 has the detail.

Why I'm asking, and why I keep changing my mind

> *I think I'm getting value. Am I though? If I am, who keeps that value? And if not, who's going to pay for the AI?*

I want to be clear about where I'm coming from. In March 2026, I published my now best-selling book, Accrual Intentions, about an experiment in which I built the world's first 100% AI accountancy firm. I wanted to go to the extreme to find out whether AI could replace a human team. It couldn't. What it did was move the work: from doing to reviewing, from producing to directing. As much as I tried, I still remained the bottleneck. The hard part wasn't the technology. It was leadership and management: giving the AI clear goals, context, processes and cross-checks. I wrote up what worked and what didn't in a free guide (accrualintentions.com/guide).

My whole career has given me a useful perspective on this problem. I started as a business operations consultant at PA Consulting, and I've run my own consultancy, Bridging Insight, for more than 16 years, working on how teams actually perform. Through the businesses I've co-founded, Spidergap, AirManual and now Ukoru, I've worked with thousands of businesses, sometimes deeply enough to see exactly how they work and where they get stuck. And as a speaker, I hear from people around the world, in every kind of role, about what makes change hard, what makes delegating and managing people hard, and now what makes AI hard to get value from.

If that's taught me anything, it's that changing how people work takes years, not quarters. I've sat in the boardrooms of large organisations, including some backed by private equity and banks, where the message from above is clear: invest in AI, cut costs, deliver shareholder value. The leaders I meet are struggling to make that happen in practice. That isn't a technology problem. It's a management one. So when I look at the prices being paid for AI, I can't shake the feeling that investors are assuming a faster rollout than any change programme I've ever seen.

Then there's the loop I keep getting stuck in. I use AI every day, and I'm sure I'm getting value from it. Am I, though? If I am, who keeps that value: me, my clients, or the company I pay for the AI? And if it's mostly flowing to me and my clients, what exactly are investors paying $2tn for? If I'm not getting value, then who is going to pay for the AI?

I've gone round that loop more times than I'd like to admit.

I've also changed my mind on this paper more than once, and so have my AI helpers. In one earlier draft, an AI tool helped me confidently "correct" a research finding that had been right all along. So treat what follows as my working based on available evidence, not a final verdict.

The best case for "no bubble"

Here's the other side, as its supporters would put it:

- **The demand is real and broad.** Anthropic's customer base is widening fast, and in July–September about 79% of its revenue came from outside its top five customers (Anthropic S-1/A).
- **Cheaper AI creates new uses.** Every time a unit of AI gets cheaper, tasks that weren't worth automating become worth it. Total spending can grow even as prices fall.
- **The biggest spenders are hugely profitable.** Nvidia's day-to-day business brought in $74.4bn of cash in six months (Nvidia 10-Q). These aren't dot-com start-ups with no revenue.
- **Much of the commitment is flexible.** About $213.8bn of Anthropic's $895.4bn of commitments can be cancelled or depends on conditions (Anthropic S-1/A).
- **Capacity is short, not surplus.** Anthropic had to buy $5.0bn of expensive short-term computing in one quarter because demand ran ahead of its contracts (Anthropic S-1/A).
- **The banking system looks well insulated.** Regulators estimate large US banks' AI-related lending at under 1% of their assets (Chicago Fed), and a study by the Bank for International Settlements (the central banks' own bank) found only limited, short-lived knock-on effects in markets from AI financing news (BIS).
- **Investors are still queuing up to lend.** Amazon's recent $25bn bond drew $62bn of orders, and Meta's share sale was five times oversubscribed (Simply Put).

Every one of those points is true, and some of them are why I don't expect a banking crisis. But none of them tells us whether the cash will catch up with the commitments and the prices. That's the question the rest of this paper is about.

How to read the numbers in this paper

Every figure here comes from a company filing, a company announcement, named research or named press reporting, and I give the date. Each source is linked in brackets after the claim, usually the first time it's used in a section. Tables have their sources underneath in italics. The summary at the top is the exception: its figures are sourced in the sections that follow. Where I've done the sums myself, I say so and show the working.

Big AI numbers get quoted on stage as if they were all the same kind of number. They aren't, so I've kept to five rules:

- **A guarantee's maximum isn't a cost.** It's the worst case if everything goes wrong, and it may never be paid.
- **A commitment isn't debt.** It's a signed promise to pay later, and some of it can be cancelled.
- **A funding gap isn't a loss.** It's money someone has to find. It only becomes a loss if nobody can.
- **The same dollar isn't counted twice.** If a supplier pays a lender under a guarantee, the loss has moved. It hasn't doubled.
- **A scenario isn't a forecast.** "What if" sums use stated assumptions. They show the shape of a problem, not a prediction.

Not all sources are equal. Anthropic, Meta, Oracle, Nvidia, Broadcom, CoreWeave and the Eagle Point funds publish regulated filings. OpenAI doesn't: its figures come from press reports of internal documents, mainly in the Financial Times and Bloomberg. I've labelled which is which. That difference is also one reason Anthropic features so much: it's the best-documented.

Anthropic's 2026 figures are unaudited interim results from its updated prospectus (the document a company publishes before its shares are listed), filed on 2 November 2026 (SEC). My research cut-off is 4 November 2026. Nothing here is investment advice. I'm explaining my answer to a question I get asked on stage, not telling anyone what to buy or sell.

01

What I mean by "bubble"

> *The technology can succeed while the people who paid for it lose.*

A bubble is when the price paid for something runs well ahead of the cash it will ever produce. The thing itself can be brilliant. Railways, telephone networks and the internet all changed the world, and plenty of people who financed them still lost fortunes.

The useful question is "who pays when expectations reset?" History gives two very different answers.

Episode	What was lost	Who carried it	Did it spread?
Dot-com crash, 2000–02	About $5tn of stock-market value, peak to trough	Mostly shareholders	Banks largely escaped
Telecoms bust, 2002	$163bn of corporate bonds defaulted worldwide that year; 56.4% (about $91.9bn) were telecoms	Bondholders and equipment suppliers	Severe for the sector, but no general squeeze on lending

Sources: IMF working paper (Dagher, 2018) for the dot-com figure; OECD working paper (Lenain and Paltridge, 2003) for telecoms. These are different measures: lost share value isn't a cash loss, and defaulted bonds aren't the same as final creditor losses.

The dot-com crash hurt, but it was mainly shareholders losing paper wealth. The telecoms bust is the closer comparison for AI: a useful technology whose use kept growing after the crash, but operators had borrowed to build for much faster growth than arrived. The fibre got used eventually. The people who paid for it didn't get their money back.

It has already happened once this year

I use China here because it's the one place where an AI share-price bubble has visibly burst. Shanghai's STAR 50 index of AI and chip companies fell about 40% in under four months from its July peak, wiping out about $480bn of value. Even after the fall, its companies are valued at about 100 times last year's profits, down from nearly 250 times at the peak (Reuters, via BusinessDay). Yet Rhodium Group still estimates Chinese AI investment rising from $139bn in 2026 to $193bn in 2027. The technology carries on; the share prices didn't.

So when I say AI is "almost certainly a bubble", I mean two things:

- **Valuations:** the prices paid for the leading AI companies need years of near-perfect growth (section 2).

- **Financing:** a growing share of the infrastructure is paid for with borrowed money and promises, which can turn a share-price correction into a credit problem (sections 3–6).

The first is almost certain to correct at some point, as China shows. The second decides whether it stays a correction or becomes something worse.

02

The valuation test: what has to be true for $1tn+

For Anthropic to be worth $2tn, my rough sums say its yearly revenue would need to grow by about 35–45% every year for five years, and then stay highly profitable. That's my working, not a forecast, but it shows how much future growth is already built into the price.

I use Anthropic and OpenAI here because they're the two AI companies priced at around $1tn or more that don't have a separate business model outside of AI.

The prices on the table (as at 4 November 2026):

- Anthropic sold shares in May 2026 at a valuation of $965bn, including the new money. Its amended prospectus, filed on 2 November, still leaves the price range blank. Press reports suggest the public sale could value the whole company at $1.5–2tn, which Anthropic hasn't confirmed. (StockTitan summary of the S-1/A; Invezz)
- OpenAI raised $122bn at an $852bn valuation in March 2026 (Bloomberg). It's now reported to have secured about $11bn of a round that started with a $30bn target, at $1.25tn rather than the $1.4tn it sought a month earlier (Dataconomy, 2 November).

My working, using the best-documented case. Anthropic's latest quarter (July–September 2026) brought in $16.956bn, or about $67.8bn a year if every quarter matched it (Anthropic S-1/A). Suppose investors eventually value the company at 25 times its yearly profit, and it ends up keeping 20–25% of its revenue as profit.

Valuation	Yearly profit needed (at 25 times profit)	Yearly revenue needed (keeping 25% → 20% as profit)	Revenue growth needed every year for 5 years
$1.5tn	$60bn	$240–300bn	about 29–35%
$2.0tn	$80bn	$320–400bn	about 36–43%

The 25-times-profit price and the profit shares are my assumptions. Growth is measured from today's $67.8bn a year. Change them and the answer moves, but not by enough to make the conclusion change.

There's one more reason the bar is higher than it looks. These sums show what the company would need to be earning in five years. But investors buying today also want a return for waiting. If they want 10% a year, a company worth $2tn in five years is only worth about $1.24tn today.

To be worth $2tn today, it would need to be worth about $3.2tn by then, which means revenue growing by roughly 50–57% a year, not 36–43%. The 10% is my assumption, but any reasonable return raises the bar further.

Is that growth impossible? No. Anthropic's revenue grew about 47% between April–June and July–September alone. But that's growth in a single quarter, from a much smaller base, and it was already slowing (section 3). Keeping up 30–45% a year for five years, on top of today's $67.8bn, is a different challenge, and three things make it harder than the headline suggests:

- **Prices are falling.** Anthropic's listing documents (October–November 2026) say price cuts on new models and contract discounts reduced revenue by about $600m between the second and third quarters.
- **Costs are rising faster than revenue.** Spending on computing power and infrastructure went from 48% of revenue to 58% in a quarter. Part of that is because demand ran ahead of contracted capacity, so Anthropic bought $5.0bn of more expensive short-term computing. That's a sign of strong demand, and of thinner margins (Anthropic S-1/A).
- **Existing owners' slices will shrink.** Had the listing happened on 30 September, Anthropic says it would have booked $8.2bn of staff pay in shares at once, with $24.1bn more over about two years (Anthropic S-1/A).

OpenAI's sums are harder because it doesn't publish accounts. On press reports of its figures, its quarterly revenue in April–June 2026 was about $6.7bn against Anthropic's $11.6bn, with a far larger operating loss, yet its latest round was priced at $1.25tn (AI Funding Tracker, compiling press reports). If the best-documented company needs near-perfect growth, the bar for OpenAI is higher still.

If a friend who owned an accountancy practice selling their shares at a multiple of fees came to you with these numbers, you'd ask two questions before any other: "How sure are you that fees keep growing at this rate?" and "What have you already committed to spend, regardless?" Section 3 answers the second one.

03

A simplified going-concern test: the clients don't leave, and the money still runs out

Anthropic currently passes a 12-month going-concern test comfortably. Stretch the same test to the end of 2030 and assume revenue simply stops growing, and it needs $27–114bn of new money to meet the commitments in its own filing. Nothing has to go wrong. Growth just has to pause.

I run this "stop the music" test on Anthropic first because it's the best case: fast growth, positive operating cash flow and $131.8bn in the bank (as of the end of Q3 2026). If the strongest company needs new money under a flat-revenue assumption, the others need more. Further down I run the same test on OpenAI, CoreWeave, Oracle and Meta.

Growth & Co.: the same numbers, as a practice

Imagine an accounting practice that bills £10m a year. Clients love it and fees have grown fast. On the strength of that growth, the partners have signed leases and outsourcing contracts that cost about £16m a year for the next four years.

It has about £20m in the bank, but most of that (about £16m) is money the partners put in this year, not fees it earned.

If fees stay at £10m, the firm is £6m short every year before it pays a single member of staff, so the bank balance runs out in under three and a half years, before the contracts end. Add its other costs at 30% of fees, which is still less than most practices pay their people, and it's £9m short a year. The money runs out in a little over two years.

No client has left. Nobody has stopped paying. The firm simply signed up for the growth it expected, and the growth paused.

Those are Anthropic's numbers, rounded and scaled down so that £10m of fees equals its current yearly revenue. Its real costs outside computing were about 34% of revenue last quarter, so the 30% case is, if anything, kind to it.

The 12-month test: an easy pass

Anthropic's listing documents show $69.8bn of payments due in the 12 months after 30 September 2026. At the July–September pace, a year's revenue is $67.8bn, and the company holds $131.8bn of cash and short-term investments. On this simple comparison, there's no sign of

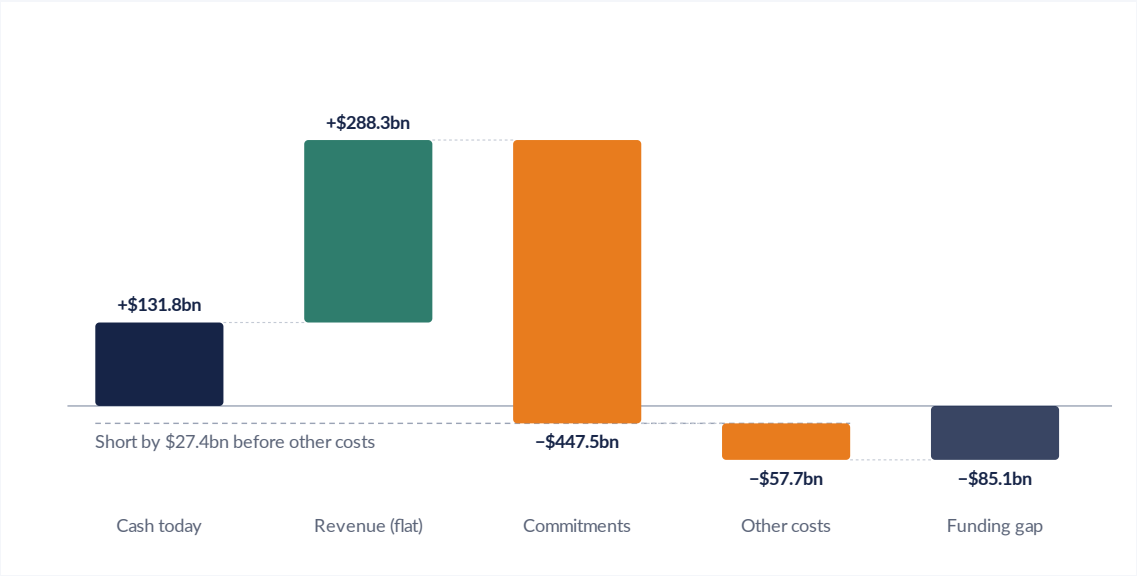
near-term trouble. A real going-concern review looks at far more, but nothing here suggests a company in trouble today.

The four-year test: a funding gap

The longer view is different. All inputs below come from Anthropic's amended prospectus of 2 November 2026, but the working is mine. It assumes every quarter to the end of 2030 matches July–September 2026, and that the $447.5bn of commitments the filing expects before 2031 are paid in that period. Ed Elson reached the same figures independently from the same filing, and puts them in context: Anthropic has spent about $29bn on computing since 2023, about $19bn of it this year, while its commitments imply computing bills of more than $100bn a year from 2027 (Simply Put).

Flat revenue leaves an $85bn funding gap by 2030

Anthropic, 1 October 2026 to 31 December 2030, US$bn. Teal adds cash, orange uses it, navy is a balance.



Source: Anthropic Form S-1/A, 2 November 2026 · author's calculation · a flat-revenue scenario, not a forecast

Even before staff and running costs, the commitments outrun four years of flat revenue plus every dollar of cash in the bank.

The gap depends heavily on what the company spends besides its computing contracts:

Other cash costs as % of revenue	Funding gap by end-2030
0% (deliberately unrealistic)	$27.4bn
10%	$56.2bn
20%	$85.1bn

Other cash costs as % of revenue	Funding gap by end-2030
30%	$113.9bn

Which row is realistic? In July–September, Anthropic's operating costs outside computing were about $5.7bn once a one-off $1.96bn non-cash charity charge is removed (Anthropic S-1/A). That's about 34% of revenue, though some of it is staff pay in shares rather than cash. My best guess is that the truth sits between the 20% and 30% rows.

How much growth would be enough?

Less than you might think. Keep the commitments fixed, keep other costs at 20% of revenue, and let revenue grow each year from 2027:

Revenue growth each year from 2027	Revenue to end-2030	Cash left (or gap) at end-2030
0%	$288bn	($85bn)
5%	$324bn	($57bn)
10%	$363bn	($25bn)
15%	$406bn	$9bn
20%	$454bn	$47bn

So growth of about 14% a year closes the gap if other costs are 20% of revenue, or about 20% a year if they're 30%. That's well below the 30–45% a year a $1.5–2tn valuation needs. Two caveats: this is a total to 2030, so cash could still run short along the way, and faster growth would itself need more computing. The working is mine, using the figures above.

Why "flat revenue" isn't as absurd as it sounds

You might well laugh at the idea that AI revenue could stop growing. Fair enough. I'm not forecasting it. But it's less absurd than it sounds, for two reasons.

AI takes time to reach the bottom line, and most businesses aren't there yet. In my own experiment, the savings didn't arrive as free time. The work moved to review, and the bar went up. The surveys say I'm not alone. In PwC's January 2026 survey of 4,454 CEOs, 56% said AI had delivered neither more revenue nor lower costs over the previous year, and 22% said it had raised their costs (PwC). KPMG's mid-2026 survey found only 7% of leaders had an established return on AI (BERI, citing KPMG). And nearly a quarter of CEOs told EY-Parthenon that their productivity gains were being absorbed by complexity, regulation and extra work rather than turning into financial results (CFO Dive). If there are no savings, there's nothing to pass on to customers, and nothing to pay the AI bill with.

Businesses are already rationing it. After a spring of "token-maxing" where leaders were encouraging staff to use as much AI as possible, the reaction has set in:

- **Uber** used up its entire 2026 AI coding budget by April, then capped each employee at $1,500 a month per tool (Faros, citing Bloomberg).
- **Meta** replaced open access with department budgets for tokens (units of AI text) in June (The Information, via Digg).
- **Accenture** now charges heavy use to project budgets. Its agentic AI strategy lead told staff that senior leaders were "still asking the question of whether they're getting value" (TechCrunch).
- **Goldman Sachs** went further as the year's last quarter began. Under internal guidance updated in October, it took Claude Code and GitHub's AI coding tools away from most bankers and traders and limited them to its roughly 12,000 engineers, after reviewing productivity and risk. Staff are being steered towards Goldman's own in-house assistant, which uses models from several AI companies (Bloomberg, via Yahoo Finance).

You can see the timing in Anthropic's numbers. Its revenue grew by about $6.8bn (144%) from January–March to April–June, then by about $5.4bn (47%) to July–September: slower in both percentage and dollar terms. That's my calculation from the filing. Goldman's October move suggests the rationing has carried on into October–December, so Anthropic's next quarterly figures will be telling.

The usual explanation is that Anthropic was short of computing capacity. But in the same quarter, it cut prices and gave discounts worth about $600m. Those two don't sit comfortably together. If demand really outstrips supply, basic economics says you raise prices or ration, not cut them. There are partial explanations: newer models are cheaper per unit of work but get used more, and big customers get discounts for long commitments, much as airlines discount advance fares on planes that end up full. But the simplest reading is that competition is fierce enough to push prices down even during a shortage. And if prices are falling while capacity is tight, I'd expect them to fall faster once it isn't.

None of this means revenue will go flat. It means the gap between "AI is useful" and "AI pays back $2tn" is filled by time, and time is the one thing a borrowed-money build-out can't easily buy.

Usage isn't revenue, and revenue isn't profit

It's easy to treat four different things as one:

- how much AI gets used
- how much useful work it actually does
- what the customer pays for that work
- how much of that payment the AI company keeps as profit.

They can move in opposite directions. AT&T's AI usage grew more than five-fold in a year while its spending stayed roughly flat (section 8). In my own experiment, AI did a huge amount of work, but much of the value went into higher standards and more review rather than lower costs. And Anthropic's price cuts and rising computing costs mean each dollar of its revenue is earning it less.

To be fair, the cost of supplying AI is falling too. Anthropic bought expensive short-term computing because demand outran its contracts, and says it plans to replace it with cheaper long-term capacity (Anthropic S-1/A). If its costs fall faster than its prices, its margins could improve even as AI gets cheaper. AT&T's experience shows the customer wins; on its own, it doesn't prove the supplier loses. The real question is which falls faster: the price customers pay, or the cost of delivering the work.

So AI can be a real success for the people using it while becoming a worse business for the people supplying it. That's what happens when a technology gets cheaper faster than people find new things to do with it, and it's the gap the bubble sits in.

Can Anthropic fill the gap? Perhaps, but not without risking a panic

Anthropic has options. The listing could raise tens of billions; there's speculation of around $100bn. It has an unused $15bn borrowing facility, like an agreed overdraft, and agreements allowing up to $77bn more of loans that can later turn into shares. The facility is secured on its assets, and will stay that way, at a higher interest rate, until two of the three big rating agencies rate it investment grade (GTR). About $213.8bn of its $895.4bn total commitments can be cancelled or depend on conditions, and its large SpaceX computing deal can be ended on 90 days' notice after early 2027. If those terms allow, cutting or delaying about 19% of the pre-2031 spending would close the gap in the 20% case. Cancelling usually comes with conditions and costs, so this is room to manoeuvre, not a free exit.

Put simply: flat revenue on its own probably wouldn't sink Anthropic, because it has room to raise money, cut and delay. But every one of those options means either new investors or a smaller build-out, and a smaller build-out undermines the growth story the price depends on.

There's a catch in that escape route. If Anthropic, or any big AI company, started cancelling or delaying its commitments, the companies on the other side of those contracts would feel it at once. Suppliers' order books would shrink, data centres built for those contracts would lose their rent, and the lenders behind them would start asking who's next. It's like someone shouting "fire" in a crowded theatre: the stampede for the exits can do far more damage than the smoke ever would. One company sensibly protecting itself could set off the very panic this paper is about. And some of it can't be escaped at all: more than $430bn sits in four deals, with Broadcom, Google, Amazon and Microsoft, that can only be ended by default, and Broadcom's leases describe the rent as "absolute and unconditional" (Yahoo Finance; Simply Put).

So this isn't a prediction that Anthropic fails. It shows that even the best-placed company needs new investors' money to keep its promises unless growth continues. That points to three separate questions, with different answers. What growth keeps the company funded? About 14–20% a year. What growth gives today's shareholders a decent return? 30–45% a year or more. And what happens to its suppliers if it cuts back to protect its cash? **Anthropic could pass the first test and still fail the second.** And if the market stops believing in the growth, the money to fill any gap gets dearer exactly when it's needed.

The same test on others: Anthropic is the good news

It would be easy to read the Anthropic numbers and think, "That's not too bad." That's exactly why I started with the best case. Here's the same idea applied to four others, as closely as the published figures allow. I picked the one most stretched in each part of the build-out: the AI model maker most dependent on new money (OpenAI), the most indebted specialist cloud provider (CoreWeave) and the big tech company leaning hardest on borrowing (Oracle). Meta is there for contrast, as one of the richest spenders. OpenAI doesn't publish accounts, so its figures come from press reports of its internal plans and are only indicative.

Company	What I compared	Result	What it means
Anthropic	Spending it has committed to before 2031, against flat revenue	About $160bn short before staff and other costs. Its $131.8bn of cash covers most of that	Needs $27–114bn of new money. Uncomfortable, but survivable
OpenAI	Its reported plan to spend about $856bn on computing over 2026–30, against five years of flat revenue at about $50bn a year	About $600bn short before staff and other costs, nearly four times Anthropic's gap	Even if growth goes to plan, it expects $278bn more cash going out than coming in, and March's $122bn raise to run out by 2028
CoreWeave	Cash from its business against spending on equipment, January–June 2026	Spent about $3.85 on equipment for every $1 its business brought in, and owes about $51.6bn	Needs lenders to keep lending every year. At that pace, its business would take about seven years to bring in enough cash to repay its debt, before any interest
Oracle	Cash from its business against spending on equipment, June–August 2026	Spent $28.5bn on equipment against $23.1bn of cash from its business, and nearly half of that cash was customers paying in advance. Owes about $125bn	Has sold about $40bn of new shares and bonds this year to keep building, with a credit rating one step above high-risk. Expects free cash flow to turn positive in the second half of its 2028 financial year
Meta	Cash from its business against investment spending, July–September 2026	About $8.5bn a year short if that quarter repeated	Trimming investment by about 6% closes it. Expensive, but not a credit problem

So of the companies here, Anthropic is the one best placed to cope. If it needs new money when growth pauses, OpenAI and CoreWeave need far more, and they're the ones that would carry a correction into the lending system.

Sources: OpenAI's plan and cash forecast from the FT, via Reuters on Investing.com and AI Funding Tracker; its revenue run rate of about $50bn a year from Bloomberg, via ValueAdd VC. CoreWeave 10-Q for the quarter to 30 June 2026 and CNBC. Oracle's 10-Q for the quarter to 31 August 2026 and October 2026 investor update. Meta Q3 2026 results. Comparisons and rounding are mine.

04

What's different this time: borrowed money, leases and guarantees

> *The people selling the shovels are now lending to the miners.*

The AI companies at the top are mostly funded by shareholders, which is how dot-com was funded. The infrastructure beneath them is increasingly funded by debt, leases and promises from suppliers, which is how the telecoms bust was funded. That shift is what worries me most, and the credit markets have started to notice.

It helps to know who's who. These are the types of firm in this story:

Type of firm	What they do	Examples in this paper
AI model makers	Build and sell the AI itself	Anthropic, OpenAI
Big tech cloud providers ("hyperscalers")	Rent out computing power and build data centres; most also make their own AI	Amazon, Microsoft, Google, Oracle; Meta builds for its own use
Specialist AI cloud providers ("neoclouds")	Rent out AI computing, mostly paid for with borrowed money	CoreWeave
Chip and equipment suppliers	Make the chips and kit inside the data centres	Nvidia, Broadcom
Power and construction suppliers	Make the turbines, generators and machinery that build and power data centres, often with orders booked years ahead	Caterpillar, GE Vernova, Eaton, Vertiv (section 6)
Data-centre developers	Build and own the buildings, often with borrowed money	Nexus Apex, Meta's joint ventures
Lenders and investors	Fund all of the above	Banks (Bank of America, Citi, Morgan Stanley, TD), private-credit funds (Blackstone, Apollo, Eagle Point), insurers (Athene), pension funds (CDPQ)

The sections below go through them in that order, starting with the model makers and ending with the people lending to everyone else.

AI model makers: mostly shareholders' money

Anthropic raised $110.4bn by selling shares in the first nine months of 2026, while its operations generated $6.6bn of cash (Anthropic S-1/A). If its valuation falls, that's mainly a loss for shareholders, and shareholders can't call in a loan. Lenders are less sure than share buyers, though. The FT reported on 30 October that Blackstone, the lead investor in the $63bn Broadcom deal, now expects Anthropic to be rated "BB" after it lists, the top tier of junk, having previously expected investment grade before 2027 (TechCrunch). A smaller rating agency, Kroll, already rated it BB+ in June (Benzinga).

Big tech cloud providers start borrowing: Oracle is the extreme case

I use Oracle because it's the large cloud provider leaning hardest on borrowed money. On 2 November it agreed the sale of $10bn of new shares and $10bn of bonds to pay for its AI build-out, due to complete on 4 November. In July, S&P cut its credit rating to BBB-, one step above the level where bonds count as high-risk "junk". It had $125.34bn of debt at 31 August, and its shares are down about 50% over 12 months (CNBC).

It isn't alone. After almost three months without big tech borrowing, the giants are back. On 3 November, Meta borrowed $15bn by selling bonds. That's its third big bond sale in about a year, after $30bn last year and $25bn in April (Reuters, via bdnews24). Alphabet's latest euro bond cost it more than its May deal (Technology.org).

The point: lenders are still happy to lend to the giants, but companies that used to pay for everything from their own cash are now borrowing tens of billions a year to keep building, and in some cases paying more to do it.

Specialist AI cloud providers: CoreWeave is the extreme case

CoreWeave rents out AI computing to the big AI companies, and it's the most indebted of the specialists. It has about $51.6bn of debt. Lenders buying its 2031 bonds now demand about 16% a year, and insuring its debt against default costs more than 10% of the amount insured each year, a level generally read as distress. Its debt is about seven times its shareholders' equity, up from about three times a year earlier (CNBC).

One trading desk summed up the mood in October: the market is "basically saying 'no mas' to AI debt without operating cash flow support" (Evercore ISI, via CNBC).

Chip suppliers: Nvidia for scale, Broadcom for the mechanism

Imagine your software supplier guaranteeing your office lease so you can afford its licences. That's roughly what's happening:

- **Nvidia** (the largest disclosed guarantee) has promised support of up to $105bn for land, power and building obligations linked to an OpenAI data-centre project in Ohio (Nvidia 10-Q). The support falls as OpenAI performs.

- **Broadcom** (the clearest example of the mechanism) has promised to cover up to about $43bn if equipment leased to Anthropic through a $63bn financing structure is worth less than the remaining leases after a default (Reuters, via The Standard). Bank of America, Citi and Morgan Stanley are arranging the loans that get repaid first; Blackstone alone is providing the slice that gets repaid last.
- **Meta** has promised to cover up to $41.5bn if two data-centre joint ventures it owns 20% of turn out to be worth less than expected when their leases end, up from $28.1bn at the end of 2025 (Meta 10-Q).

Those are maximums: worst cases that may never be paid. Nvidia's and Broadcom's maximums are each nearly twice their cash, but Broadcom's management judges a payment unlikely, Nvidia's support falls as OpenAI performs, and both companies generate a lot of cash. What matters is when these promises would bite: in a downturn, exactly when the suppliers' own sales would be falling. The insurer and the insured are standing in the same building when it catches fire.

The relationship an auditor would flag: OpenAI is the extreme case

OpenAI's March round was anchored by three companies that also sell it infrastructure: Amazon committed $50bn ($35bn of it conditional on OpenAI listing or reaching artificial general intelligence), and Nvidia and SoftBank $30bn each (Bloomberg). Amazon's deal included OpenAI using Amazon's own AI chips (Verdict, via Yahoo Finance). Nvidia invests in OpenAI and supports its data centre, and OpenAI buys Nvidia's chips. SoftBank is OpenAI's partner in the Stargate data-centre venture.

Anthropic shows the same pattern on a smaller scale. Google, Amazon and Microsoft are its shareholders, its suppliers (it owes them at least $276bn of minimum cloud spending between them, whatever it uses) and its sales channel (about 44% of its revenue in the first nine months of 2026 came through their marketplaces) (Anthropic S-1/A).

None of that is improper. But if your biggest client were also your landlord, your main supplier and a shareholder, you'd want to know which way the money really flows before you relied on any single number. Most marketplace sales will be to genuine outside customers, so this isn't circular in itself. The concern is narrower: when the same few companies invest the money, supply the computing and sell the product, an outsider can't easily tell how much of the growth they are, in effect, funding themselves.

05

Who owns the downside: from data centre to pension

> *A 30% fall in a data centre's value can wipe out everyone who put in the riskiest 30% of the money.*

The risk doesn't stay with the banks whose names are on the deal. It spreads across investment funds, insurers and pension investors. The IMF made the same point in October: more and more lending now happens outside banks, where it's harder to see, and some of it is more concentrated than the headline names suggest.

Working machines can still be bad investments

A data centre can keep running perfectly and still lose its owners most of their money. Lenders care about what the chips can earn and what someone else would pay for them, and both can fall while the chips keep working.

Bitcoin mining showed how fast that can change. The specialist machines kept working through 2022, but what each one could earn collapsed as energy costs increased, Bitcoin's price fell and competition grew. The price of the most efficient machines fell about 87%, from $119.25 per unit of computing power in May 2021 to $15.71 by December 2022. Older models fell 91% (Hashrate Index data, via Cointelegraph). Mining companies that had borrowed to buy machines, such as Core Scientific, ended up in bankruptcy or selling assets to survive.

AI chips do many more jobs than Bitcoin miners, and demand for them is strong today. But the same pressure applies: newer, more efficient chips keep arriving, and every price cut by an AI company reduces what older chips can earn.

China has already seen a small version of this. Sichuan Binji Technology, which bought AI chips with borrowed money through sale-and-leaseback deals, went bankrupt and caused direct losses for 15 licensed leasing companies in April–June 2026. It's one company, but this kind of leasing is common: 54.3% of new Chinese computing-centre projects in the first half of 2026 involved it (Rhodium Group). Rhodium notes that how fast chips lose value varies a lot by model, so this example shows the risk exists but says little about its size.

Small falls in value, big losses for some

The order of repayment is important. Take a $100 data-centre project funded with $10 of owners' money, $20 from "junior" lenders who get repaid last and $70 from "senior" lenders who get repaid first. Losses fall on the owners first, then the junior lenders.

Fall in the project's recoverable value	Owners lose	Junior lenders lose	Senior lenders lose
10%	100%	0%	0%
20%	100%	50%	0%
30%	100%	100%	0%
40%	100%	100%	about 14%

This is an illustration, not a named project. Now add a second layer. If the fund holding that junior debt has itself borrowed 40% of its money, a 20% fall in its assets becomes a 33% loss for the fund's investors. At 60% borrowing, it's a 50% loss.

In Growth & Co. terms: the partners' capital goes first, then a loan from a partner's family trust that agreed to be repaid after the bank, and only then the bank. The bank can be fine while the family trust is wiped out.

Where the risk has landed: the most concentrated example I found

I use the Eagle Point funds because they're the clearest public example of one AI borrower making up a large slice of a fund. They're not typical, and that's the point: a manager's huge total assets can hide a large bet for the people in one particular fund.

At 30 June 2026, four listed Eagle Point funds held "senior secured" loans to Nexus Apex, a data-centre borrower, at a stated interest rate of 14.73% and maturing in February 2029:

Fund	Gross holding, % of net assets (before participations sold)	After Nexus participations sold
Eagle Point Enhanced Income Trust	20.8%	Not separately disclosed
Eagle Point Credit Company	15.7%	about 12.4%
Eagle Point Institutional Income Fund	12.3%	about 9.6%
Eagle Point Income Company	12.0%	about 9.5%

These are gross values before some risk was passed on to other investors, not predicted losses. "Senior secured" means first in line among Nexus Apex's own lenders. It doesn't necessarily put these funds ahead of lenders to the businesses Nexus Apex owns, so the protection may be weaker than the label suggests.

Sources: each fund's portfolio filing with the US Securities and Exchange Commission for 30 June 2026, for example Eagle Point Enhanced Income Trust's; Eagle Point's announcement of the Nexus Apex financing, March 2026. Percentages are each fund's own figures. The other three filings: Eagle Point Credit Company, Eagle Point Institutional Income Fund and Eagle Point Income Company. All four funds have also sold part of their Nexus Apex loans to other investors: $20.0m for Eagle Point Credit Company, $3.0m for the Institutional Income Fund and $7.2m for Eagle Point Income Company, while the Enhanced Income Trust sold $6.6m across this and one other loan. So the gross percentages overstate the risk each fund has actually kept. The right-hand column is my calculation, deducting the Nexus participations each filing discloses.

The chain goes further than funds. These are single examples, chosen because they're the only public ones I found at each link:

- **Pension money.** CDPQ, a Canadian pension investment manager, was named among the participants in CoreWeave's $7.5bn debt financing in May 2024; its own share wasn't disclosed (CoreWeave).
- **Insurance money.** Athene, an insurer, guarantees 15% of an Apollo-managed vehicle's obligation to buy Broadcom equipment. If the full $35bn were outstanding, that's a $5.25bn ceiling, falling as the obligations are funded and expected to end by 2028 (The Tape).
- **Banks passing risk on.** Blackstone reportedly took the first 8% of losses on a $2bn TD Bank loan portfolio that includes data-centre loans, according to Octus, a credit-research firm, citing unnamed sources; neither party confirmed it (Octus, 17 July 2026). That tranche covers up to $160m of losses on a portfolio that also includes other kinds of loan. Bloomberg separately reported TD weighing this kind of deal for its data-centre loans (Bloomberg). The bank keeps the loans but sells part of the downside, which frees up capacity to lend again. TD isn't alone. In May, the Financial Times reported that JPMorgan, Morgan Stanley and SMBC were also looking to pass on parts of their data-centre loans as their exposure neared internal limits (FT, via Resultsense).

The brakes

There are real buffers. Eagle Point Credit Company's own debt is fixed-rate, none of its financing matures before January 2029, and its participation agreements don't require it to put up extra security if the loans fall in value (ECC semiannual report). But it isn't immune. Funds like this must keep their assets at a set multiple of what they owe: at 30 June, its preferred shares were covered 2.12 times, against a legal minimum of 2.0. On a simplified calculation, a fall of about $60m in its assets, roughly a tenth of its net assets, would take it to that floor. Below it, the fund is barred from new borrowing and some payouts, and if the shortfall isn't fixed within a set period, it must redeem preferred shares or repay some notes. A loss isn't automatically a cash crisis, but it can still force a fund's hand.

06

How a correction spreads, and what stops it

> *Losses spread two ways: through people who need cash fast, and through people who feel poorer.*

I'm not an economist, so here's who I'm leaning on

On how losses might spread, I'm relying on people who study this for a living. Their views range from reassured to alarmed:

- **The IMF (mainstream).** Its October 2026 Global Financial Stability Report named an AI-related market correction as one of three main risks to the financial system. It flagged about $950bn of investment and debt linked to $1.4tn of future AI commitments, and warned that private-credit funds could restrict lending under stress. It also found banks resilient under a severe stress test (IMF; Wall Street Times).
- **Moody's (credit rating agency).** It called Broadcom one of the strongest large technology companies, but warned its risk would rise if it backed several financing deals for AI labs in quick succession (Reuters, via Khaleej Times).
- **Michael Burry (the investor who bet against US housing before 2008).** He argues big tech is overstating profits by assuming AI chips last five or six years when they're replaced every two or three, understating costs by about $176bn over 2026–28 (CNBC). Others disagree, though Amazon has shortened the assumed life of some servers from six years to five (CNBC).
- **Ed Zitron (tech critic, Where's Your Ed At).** He argues AI's economics don't work, and that an OpenAI failure would be the AI bubble's Lehman Brothers moment (Quartz). He's at the far end of the debate and has critics of his own (The Argument).
- **The Bank for International Settlements (the central banks' own bank).** Its June annual report warned that an AI investment bust could disrupt credit markets on a scale comparable to 2008, pointing to circular financing and poor disclosure (The Next Web). Its October research found the knock-on effects so far have been limited (BIS Bulletin 140).
- **Scott Galloway and Ed Elson (Prof G Markets podcast).** In January, Galloway predicted AI stocks would correct and the data-centre bubble would burst, with cheap Chinese AI models as the trigger, likely sending the US into recession (Prof G). In June he forecast AI valuations falling 50–70% within two years (24/7 Wall St). Elson argues the real bubble is in private lending rather than share prices, pointing to CoreWeave, and that cheap Chinese models have pushed AI prices down while the cost of building keeps rising (Simply Put).

A note on the title. I'd been working on this paper for two weeks under the title "When the Music Stops". On the day I came to publish it, my AI research assistant spotted that Ed Elson had published a newsletter with exactly that title the day before, drawing many of the same conclusions from the same filing (Simply Put, 3 November 2026). So I changed mine to "Accrued Intentions", a nod to my book, and to the promises at the heart of this paper. We both borrowed the phrase from Chuck Prince, Citigroup's chief executive, who told investors in July 2007 that "while the music is playing, you've got to get up and dance". Four months later the housing bubble burst, Citigroup wrote down billions and he lost his job.

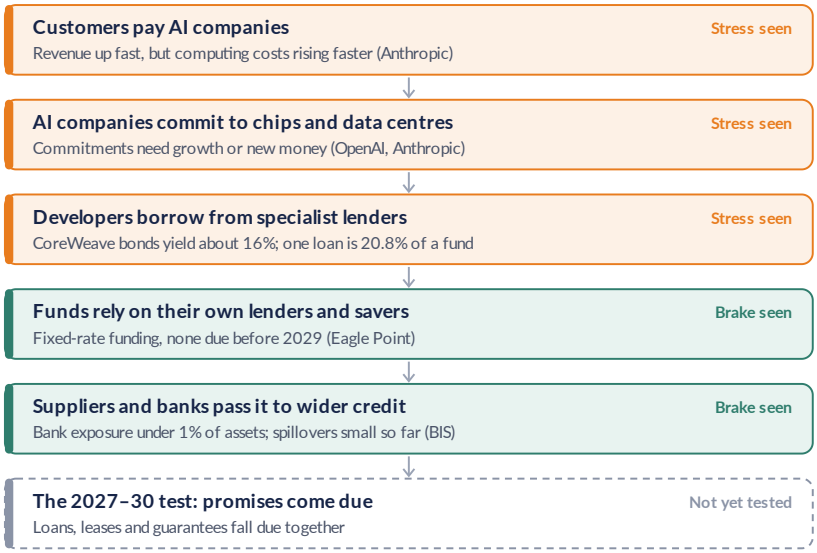
With the AI bubble, the ways for losses to spread are real, but on today's evidence they may point more towards a painful correction rather than a banking crisis. What follows is my attempt to explain those mechanisms in plain English, drawing on their work.

A big loss only becomes a wider crisis if it hits someone who has to find cash fast: a borrower facing a refinancing, a fund whose own lenders demand more security, a bank whose safety buffer falls below the legal minimum. Without those links, losses stay with the people who took the risk.

Here's the chain, with what the evidence shows at each link today:

Stress shows in the top three links; the brakes hold further down

How an AI correction would spread, and what the evidence shows at each link, November 2026.



Source: evidence as at 4 November 2026, as set out and linked in sections 4 to 6

The top three links already show strain. The brakes further down are why a fall today would probably stay contained, and 2027–30 is when they'll be tested.

Three stages, each harder to reach than the last

1. **Valuation correction.** Share prices and private valuations fall. Investors lose money, but nobody has to miss a payment. China's STAR 50 shows this stage, and the week of 26–30 October gave a taste of it in the US: Alphabet fell 13%, Intel 25% and ServiceNow 34% (Simply Put, citing S&P Capital IQ).
2. **Infrastructure distress.** Weaker income, delayed projects or failed refinancings lead to restructurings and losses for junior lenders. CoreWeave's bond prices show the market starting to price this stage.
3. **Spread to the wider financial system.** Those losses force cash demands elsewhere: funds sell unrelated assets, guarantees are called, banks lend less to everyone. I've seen no evidence of this stage yet.

Stage 1 is likely. Stage 2 is plausible. Stage 3 needs several things to go wrong at once.

What the regulators' evidence says so far

- **Bank exposure looks modest.** The Chicago Fed estimated large US banks' AI-related corporate lending at about $150bn outstanding and $450bn committed, around 0.8% of their assets, using data mostly to late 2025 (Chicago Fed). That category is broad: software, energy and chips as well as data centres.
- **No wider spread yet.** The BIS studied 18 days of big AI financing news between January and September 2026 and found only limited, short-lived knock-on effects in wider markets (BIS Bulletin 140).

But small numbers can still matter at the margin. Banks must hold a safety buffer of their own money against what they lend. A bank that must hold $10 for every $100 it lends, and loses $10bn, would need to lend $100bn less to rebuild that buffer, unless it raises new money. That's arithmetic, not a forecast, but it's why credit tightens by more than the original loss. Your clients would feel it as a tougher conversation when their next loan comes up for renewal, whether or not they've ever touched AI.

Follow one company through: why Nvidia isn't safe either

Nvidia is usually held up as the safe bet: whoever wins a gold rush, the person selling the shovels gets paid. I think that's one of the most dangerous assumptions in the whole story, and it's the clearest way I know to show how a correction ripples.

Here's Nvidia's position at 26 July 2026, from its own quarterly filing (Nvidia 10-Q; results release):

- Customers owed it $63.1bn, up from $38.5bn six months earlier. Five customers made up 70% of that, and Nvidia now gives some large customers up to a year to pay.

- It had committed to $279bn of supply purchases, mostly memory, up from $119bn a quarter earlier (CFO commentary, via Wall Street Debug).
- It held $31.6bn of unsold stock, up from $21.4bn in January.
- It has promised up to $105bn of support for an OpenAI data-centre project.

Now imagine AI spending slows. Four things hit Nvidia at once:

1. New orders fall, just as it's committed to buy hundreds of billions of dollars of components.
2. Customers who are struggling pay late, or not at all, on that $63bn.
3. Guarantees written in good times get called in bad times.
4. Its own product gets cheaper. Struggling data centres sell or rent out their chips at fire-sale prices, competing with Nvidia's new ones.

This has happened before. Cisco was the shovel seller of the dot-com boom. In the quarter to April 2001, it took a $2.25bn charge for excess inventory after what it called a "sudden and significant decrease in demand" (Cisco 10-Q). Its product gross margin fell from 64.9% in its 2000 financial year to 47.9% in 2001 (Cisco annual report). Cisco survived, but its shares lost about 90% of their value by late 2002 and didn't climb back above their March 2000 peak until December 2025, more than 25 years later (Bloomberg).

And the ripple doesn't stop at Nvidia. Its suppliers lose orders, its lenders and guarantee counterparties get nervous, and its shareholders, which include almost every pension fund and index tracker, feel poorer. The same goes for the companies building and powering data centres. Caterpillar's order backlog has almost doubled in a year to a record $78bn, much of it turbines and engines for data centres, with some orders stretching to 2030 (Caterpillar 10-Q; Manufacturing Dive). GE Vernova, which makes gas turbines and grid equipment, has a $197bn backlog after its orders almost doubled in a year, and about a fifth of its gas-power customers are data centres (GE Vernova; DCD). Orders that long are only as good as the customers' ability to pay for them.

The economy: a serious risk of recession

AI reviews of drafts of this paper (ChatGPT and Claude) repeatedly suggested I amend my conclusion to say that the correction would be painful but contained. I actually think it carries a serious risk of recession, mainly because of what falling share prices do to spending.

- **Shares are a record share of household wealth.** They made up 48% of American households' financial assets in the second quarter of 2026, above the 39% peak of the dot-com bubble (Bank of America, via Yahoo Finance).
- **They're owned by the people who do much of the spending.** The richest 10% of households own about 88% of US shares (same source). Estimates of how much of consumer spending comes from the top 10% range from 23%, measured by income (Fiduciary Trust), to 35–50% in Minneapolis Fed research (cited by Man Group).

- **Rising AI shares have been propping up spending.** Goldman Sachs estimates they added 0.4 percentage points to US consumer spending growth over the past year (Reuters, via The Star).

So when AI shares fall, the households that do a big chunk of the spending feel poorer. Add AI investment that stops growing, ordinary companies whose share prices, I suspect, assumed AI savings that will arrive later than hoped, and big tech's own profits being squeezed by AI (section 8), and there isn't much left holding growth up. Oxford Economics estimates that, without AI-related effects, underlying growth across advanced economies has been only about 0.5–1% a year, and warns that a serious AI slowdown would make the risk of recession substantial (Reuters, via Kuwait Times). Bank of America's economists describe a lasting AI correction as one of the larger risks to the US and global economy.

There's an important counterweight. Federal Reserve research found that the richest 20% of households spend only about 0.8 cents of each extra dollar of wealth, against 3.3–3.5 cents for households overall (Federal Reserve, August 2025). The people who own most of the shares react least.

That makes a big difference. US households have gained about $19.5tn of AI-related share wealth since late 2022 (UHERO). If half of it disappeared, spending might fall by about $80bn a year at 0.8 cents, or about $340bn at 3.5 cents. The first is a dent. The second, alongside falling AI investment, looks like a recession. Nobody knows which is right, which is why I call it a serious risk rather than a forecast.

Why I can't rule out something worse

The regulators' evidence above is reassuring, and it's why I don't expect a 2008-style banking crisis. But I hold that view loosely, for four reasons:

- **The money goes round in circles.** Suppliers invest in customers who then buy from those suppliers (section 4), so the same dollar can prop up several balance sheets at once.
- **Most of it is hard to see.** Lending by investment funds rather than banks, joint ventures and debts kept off companies' main accounts mean even regulators are estimating exposure from partial data. You can't measure a systemic risk you can't see.
- **Diversified isn't always diversified.** A fund can hold dozens of loans that all depend on the same few AI companies paying their bills.
- **More is added every month, and the incentives all push one way.** Many of the people deciding are paid on deals done and assets grown, and their upside depends on the boom continuing. Stepping back while competitors keep lending feels like career risk. Every quarter of 2026–28 that the build-out continues, the pile needing refinancing gets bigger.

Three possible endings

1. **Orderly adjustment.** Companies delay optional spending, raise long-term money from investors, and shareholders take the hit. Growth slows but the economy avoids recession. This

now looks optimistic to me.

2. **A burst with a recession.** Valuations reset, some developers and providers can't fund their commitments, junior lenders take losses and projects change owners. Falling share prices hit spending hard enough to tip the US into recession. Painful, but the banking system bends rather than breaks.
3. **A wider credit crisis.** Losses coincide with funds that have themselves borrowed needing cash, banks short of capital and guarantors whose own sales are falling. Credit tightens for everyone, including businesses with nothing to do with AI. Unlikely, but not as unlikely as I'd like.

My central case sits between the first two: a burst, with a serious risk of recession. Timing matters. If it burst today, I'd expect something close to the first ending. The longer the build-out runs on borrowed money, the more the odds shift towards the second and third.

Section 7 explains when I'd expect to find out.

07

The 2027–30 test: when the promises come due

The share-price test comes first, and it's imminent: Anthropic's listing. The credit test comes later. When I lined up the dates in the filings, an uncomfortable number of them landed in 2028 and early 2029.

Each row is labelled with the kind of event it is. Only two are loans that actually fall due. The rest are leases starting, guarantees being tested, shares going on sale and cash running low. They don't all land at once, but they do bunch up.

When	What happens	Source
Mid-November 2026 (reported)	Share sale: Anthropic's listing, the first public price for a leading AI model maker	Press reports; StockTitan
From early 2027	Exit option: Anthropic's SpaceX computing deal can be ended by either side on 90 days' notice	Anthropic S-1/A
2027	Lease starts: Broadcom equipment deliveries begin under the $63bn lease deal	Blackstone and partners' announcement, via Reuters
Late 2027	Share sale: Anthropic's co-founders and large holders become free to sell their shares, up to a year after listing	Anthropic S-1/A
2028	Cash runs low: OpenAI on track to use up the $122bn it raised in March 2026, per the FT	Reuters, via Investing.com
2028	Lease starts: Meta's Texas data-centre lease begins (initially four years). Guarantee ends: Athene's guarantee expected to end	Meta 10-Q; Athene, via The Tape
Nvidia fiscal 2029 (Feb 2028 – Jan 2029)	Guarantee tested: first phase of the Ohio project that Nvidia's support of up to $105bn relates to	Nvidia 10-Q
January 2029	Loan due: earliest repayment date for Eagle Point Credit Company's own borrowing	Eagle Point Credit Company
February 2029	Loan due: the Nexus Apex loans held by the Eagle Point funds must be repaid or replaced	Eagle Point fund filings
2029	Lease starts: Meta's Louisiana data-centre lease begins	Meta 10-Q
By end-2030	Promises fall due: $447.5bn of Anthropic's commitments are due by then; OpenAI expects $278bn more cash going out than coming in over 2026–30	Anthropic S-1/A; FT via Reuters

Why refinancing is where the trouble starts

You don't need a tenant to stop paying for a project to fail. You only need the loan to come up for renewal on worse terms.

Every accountant has seen this at client level: Growth & Co.'s bank facility comes up for renewal, profits are fine, but the bank's lending tests have tightened. Now scale it up. Take a data-centre campus with a $70bn loan to refinance and $8.5bn a year of rent after running costs. The new lender wants the rent to be at least 1.5 times the yearly interest bill.

New interest rate	Biggest new loan the rent can support	Gap to repay the old $70bn loan
8%	$70.8bn	none
10%	$56.7bn	$13.3bn
12%	$47.2bn	$22.8bn
14%	$40.5bn	$29.5bn
16%	$35.4bn	$34.6bn

The rent hasn't fallen by a dollar. The building is full. But at 10%, the owners need to find $13.3bn or restructure. At 16%, the return lenders now demand on CoreWeave's bonds, they'd need to find about half the loan. That's a hypothetical, interest-only example, but it's how distress usually arrives: through a lender's conditions, not a disaster.

To put those rates in context, here's what different borrowers are paying:

Borrower	Typical rate
UK homeowner, average two-year fixed mortgage	about 6%
Oracle, bonds due 2031 (market yield; one step above junk)	about 7.3%
US "junk" bonds, average	about 8.1%
UK personal loan, £5,000 over three years	about 10.5%
Nexus Apex, data-centre loan held by Eagle Point funds	14.7%
CoreWeave, bonds due 2031 (market yield)	about 16%
US "junk" bonds, lowest-rated (CCC and below)	about 17.8%

Sources: Moneyfacts, via Mortgage Strategy (November 2026); CNBC; ICE BofA US High Yield indices via FRED, average (30 October 2026) and CCC and below (3 November 2026); Moneyfacts (March 2026); Eagle Point filing; CNBC. Household and company rates aren't directly comparable, but the order is the point.

Lenders now want about twice the average junk-bond return to hold CoreWeave's bonds, close to what markets charge the very riskiest companies.

So my answer to "when?" has two parts. The valuation correction could start at any time; the listing is the first real test. The credit stress, if it comes, is most likely in 2028–29, when the money borrowed during the boom has to be borrowed again.

08

"But...": the challenges I get from the audience

Each of these challenges is usually made with one number about one company. The question I'm asking is whether the cash will catch up with the commitments and the prices across the industry, so I answer each with the company that tests it hardest.

"Anthropic's revenue grew twelve-fold in a year."

It did: $17.0bn in July–September 2026, against $1.4bn a year earlier (Anthropic S-1/A). That's real customers paying real money, and it's the strongest financial evidence against anyone who says the AI economy is all hype.

The question is whether growth can outrun falling prices and rising commitments. In the same quarter, price cuts and discounts took about $600m off revenue, computing costs grew faster than revenue, and the business made an operating loss of $567m. And Anthropic is the best case (section 3).

"Anthropic has $130bn sitting in cash."

It has $131.8bn of cash and short-term investments. In the nine months to September, $110.4bn came from selling shares; its operations generated $6.6bn (Anthropic S-1/A). So the cash pile is mostly investors' money paid in this year, and $69.8bn of commitments fall due in the next 12 months.

If a client showed you a bank balance and most of it was partner capital introduced in the last nine months, you wouldn't call it a self-funding business. You'd call it a runway. To be fair, it's a long runway: Anthropic passes the 12-month going-concern test easily.

"There'll always be more money for OpenAI."

There's still money, but it's getting harder to raise and coming in at lower prices. OpenAI's latest round was reported on 2 November to have secured about $11bn of an original $30bn target, at $1.25tn against the $1.4tn sought a month earlier (Dataconomy). The FT reported that OpenAI forecasts $278bn of negative free cash flow over 2026–30, and that it's on track to use up March's $122bn raise by 2028. Sam Altman has said OpenAI won't list in 2026 (Reuters, via Investing.com).

This is what the start of a correction looks like: rounds getting smaller and slower, and priced below what the company asked for.

"Big tech can pay for all this out of profits."

Some can, and that's the main reason I don't expect a 2008-style banking crisis. But two things weaken the argument.

Even the giants are stretching.

- **Across big tech:** analysts expect every major cloud provider except Microsoft to have negative free cash flow in 2026, with Amazon at about −$30bn, Oracle −$24bn, Alphabet −$16bn and Meta −$8bn. Together they've raised more than $300bn of debt and new shares in ten months (Ed Elson, citing Capital IQ forecasts).
- **Meta (the clearest case of a rich spender stretching):** free cash flow, the cash left after investment spending, turned negative in July–September, at −$2.1bn, as investment spending rose about 80% in a year while operating cash flow rose about 9% (Meta 10-Q). On 3 November it borrowed another $15bn by selling bonds (Reuters, via bdnews24).
- **Oracle (the clearest case of a cloud provider running out of road):** on 2 November it agreed the sale of $20bn of new shares and bonds, due to complete on 4 November, with a credit rating one step above high-risk (CNBC).

AI is eating some of the profits that pay for AI. This is the point I think is most understated. The "fortress balance sheets" are funded by business models that AI itself is undercutting:

Google is the clearest case so far. In July–September 2026, search queries grew at their fastest rate in five years, yet search revenue was almost exactly flat on the previous quarter, and paid clicks grew 8% against 13% the quarter before (Alphabet 10-Q). Google named several causes, with currency movements first, but its chief business officer also cited lower earnings from AI Mode answers (Security Boulevard). Pew Research found people clicked a result 8% of the time when Google showed an AI summary, against 15% without one (PYMNTS). Fewer clicks mean fewer chances to show adverts. Google is now adding adverts to its AI answers, and search revenue was still up 12% on a year earlier, so this is a squeeze on growth, not a collapse.

Amazon says it isn't affected yet, but behaves as if it expects to be. Its chief executive told investors that outside AI shopping agents are not meaningfully affecting its advertising (Yahoo Finance). Yet in September it blocked Meta's Muse agent, which gained over 10 million users in six weeks, with shopping its most common use (Forbes). It's also fighting Perplexity's AI browser in court. The irony: Amazon is one of Anthropic's biggest shareholders and suppliers, and runs its own agent, Buy for Me, that buys from other brands' websites, while keeping other people's agents off its own (Engadget). Walmart has done the opposite, partnering with Muse and building a gateway for approved shopping agents (MarketScreener). Sellers go where the customers are. If shoppers move to AI agents and Amazon keeps them out, the AI-friendly retailers win the traffic, and the advertising budgets follow. Amazon's advertising growth slowed from 26% in April–June to 11% in July–September, which it puts down to Prime Day moving into the earlier quarter.

Meta is the hardest to call, so here are both sides. The case that it's struggling: in January–March 2026, the number of people using its apps each day fell for the first time, from 3.58bn to 3.56bn, which Meta blamed on internet disruption in Iran and WhatsApp restrictions in Russia (NAI500). The Economy, citing Futurism, argues that Facebook feeds are filling with low-quality AI-generated content and that younger users are leaving (The Economy). Growth in the number of adverts shown slowed from 19% in January–March to 10% in July–September (Carbon Finance). In July it reported its lowest free cash flow in years alongside a disappointing forecast, and the shares fell (Bloomberg).

The case that it isn't: advertising revenue still grew 26% in July–September, daily users recovered to 3.64bn, and time spent on Instagram hit an all-time high (Meta earnings call, via Yahoo Finance). I found no sign of advertising revenue collapsing in 2026. Meta's revenue did fall year on year in 2022, for the first time ever, which may be what some commentators remember (NPR). My read: the business is strong today, but AI threatens it from two directions at once. Assistants take attention away, and AI-generated junk may drive people off, while Meta spends heavily on AI. Its hedge is Muse, its own agent.

Microsoft is struggling to charge much for AI. Its proxy statement reports more than 30 million paid Microsoft 365 Copilot user licences (Microsoft proxy statement), against over 450 million business Microsoft 365 user licences (SAMexpert), so roughly 7%. It now offers a 30% discount to organisations buying for 300 or more users, a new top-tier bundle that throws in two other add-ons worth $27 a user for just $9 more than its top suite plus Copilot, and a free Copilot Chat in existing plans. Meanwhile, it raised the price of its core Microsoft 365 suites in July 2026 (GoSearch). That looks like AI becoming a feature Microsoft has to include, paid for by raising the price of everything else.

A fortress balance sheet protects you from a bad year. It doesn't protect you from a business model being slowly undercut by the thing you're spending so heavily to build. Of the four, Google shows the effect already, Amazon and Microsoft show companies bracing for it, and Meta is mixed: strong today, exposed tomorrow.

And much of the leverage isn't with the giants at all. It sits with specialist providers like CoreWeave, the funds lending to data-centre developers, and the guarantees the suppliers have written.

"Cheaper AI just means people use more of it."

Often true, and it's great for users. AT&T says its daily AI usage grew from about 8 billion tokens to over 45 billion in a year, while it kept costs roughly flat by moving work to cheaper open models and sending each task to the cheapest model that can do it (Fierce Network). That's more than five times the usage for about the same spend. Wonderful for AT&T. Less wonderful for whoever expected AT&T's AI bill to grow five-fold.

"Demand is so strong they can't build fast enough."

That's true today. Anthropic bought $5.0bn of expensive short-term computing in one quarter, and Oracle told investors in March that demand for AI computing was growing faster than supply (Oracle results, March 2026). But a shortage today doesn't prove the long-term commitments will be profitable. Some of that capacity is for training future models, not serving paying customers. And demand at today's prices isn't the same as demand at the prices needed to pay back the hundreds of billions already committed.

09

Why the "SaaS apocalypse" was overstated

> *The fear assumed value would flow to the model makers. It's flowing the other way.*

In early 2026, investors decided AI would make established software firms obsolete. The results since then say otherwise for the firms that own trusted data and sit inside their customers' daily work. This isn't a view on whether any share is cheap or expensive today; it's about whether the business fear was right.

The panic

After Anthropic showed AI agents carrying out complex workflows such as financial analysis, roughly $285bn was wiped off software companies' market value in 48 hours (Hardman & Co). The fear was simple: if AI agents do the work, firms need fewer human users, and software priced per user collapses.

What actually happened to the businesses

Two of the firms closest to accountants' and lawyers' daily work have since reported:

- **Thomson Reuters** (owner of tax, audit and accounting tools and the CoCounsel AI assistant) grew revenue 10% in July–September 2026, excluding acquisitions. Its Tax, Audit & Accounting Professionals segment also grew 10% on the same basis. Recurring revenue was 88% of the total, and it raised its full-year growth outlook, on the same basis, to 8.5–9.0% (Thomson Reuters Q3 2026 release).
- **RELX** (owner of LexisNexis) grew revenue 7%, excluding acquisitions and currency movements, in the first nine months of 2026, with Legal up 10% as law firms adopted its AI research tools (RELX trading update; City AM). Its shares had plunged earlier in the year after Anthropic launched tools aimed at legal work.

On these numbers, customers didn't leave. They bought the AI versions of the tools they already trusted.

Why the fear was overstated

The apocalypse story assumed the value of AI would flow to the companies building the models. This paper argues the opposite pressure. Model prices are falling (Anthropic cut about $600m of revenue through price cuts and discounts in one quarter) and models are increasingly

interchangeable. Cheaper intelligence helps whoever owns the customer relationship, the trusted content and the workflow. In professional services, that's often the established software firm, and the professional firm using it.

There's a second reason accountants will recognise. AI makes producing work cheaper, but checking it becomes the bottleneck. A Harvard study of 718 firms found AI coding agents raised code output by about 30% with no statistically significant rise in finished work, because review time rose 49% (Chen and Stratton, 2026). Where checking matters, verified content and trusted workflows become more valuable, not less.

Where the fear did bite

To be fair to the other side, I looked for the clearest examples of real damage:

- **Dassault Systèmes**, the French industrial design software group, had its worst day on record, falling about 20%, after weak results and a soft 2026 outlook. Growth had slowed to about 1% once currency movements are stripped out, and licence sales were weak (Moby, via Yahoo Finance).
- **HubSpot**'s shares fell about 22% in August after results in which it deliberately lowered its guidance for AI revenue, as investors questioned how software firms will get paid for AI (Impact Pricing, citing Investor's Business Daily).

In neither case have I seen evidence of customers abandoning the software en masse. The damage came from slower growth and uncertainty about how AI will be paid for. The pressure is real for software that is little more than a form on top of a database, priced per user, with no proprietary data.

Even the firms doing well pay for AI. Thomson Reuters' Legal margin fell 2.4 points on higher technology costs, and its shares fell after the results on concerns about AI spending (Reuters, via AOL).

How investors' attitude has changed

By August, Constellation Research was describing the apocalypse as over and the sorting of winners from losers as just starting (Constellation Research). The question investors now ask isn't "will AI kill this business?" It's "how much will AI cost this business, and can it charge for it?" That's a far more normal question, and a much better one.

So far, the answer to "can it charge for AI?" is: not easily. In a survey of 128 software executives by Blue Ridge Partners, 37% had faced customer pushback on AI-linked price rises, and 24% had customers argue that AI should reduce their bill, not raise it. Among companies charging separately for AI, 82% said it made up less than 10% of their recurring revenue (Blue Ridge Partners). In Revenera's survey of 415 software leaders, 96% had added AI to their products, but 48% said customer take-up was limited or low (CFOtech). Some vendors are giving AI away to keep customers: ZoomInfo now includes its AI features at no extra cost (Tropic).

That matters for the bubble as well as for software. If customers expect AI to make things cheaper, the savings flow to them, not to the software companies, and not to the AI companies supplying the software companies.

10

Where I think it lands: a burst, and a serious risk of recession

My prediction: a sharp reset in AI valuations, real losses for the riskiest lenders and most concentrated funds, and data-centre projects changing hands.

On top of that, a serious risk of a US recession, mainly through falling share prices hitting the spending of the wealthiest households. How serious depends on timing (section 6). I'd change my view if the warning signs in section 11 start appearing together.

I'm not equally sure of every part of that, so here's how confident I am in each:

Claim	How confident I am	Why
AI valuations and infrastructure are financially vulnerable	High	Based on the companies' own filings: their promises, their borrowing costs and the growth they need
A correction could tip the US into recession	Plausible, not proven	It depends on how much the richest households cut their spending (section 6)
A wider financial crisis	Unlikely today, rising over time	The ways it could spread exist, but there's no sign of them spreading yet

What I expect

1. **Valuations reset.** At some point the market stops paying for years of near-perfect growth. The listing of Anthropic, and whether OpenAI can raise at its last price, are the first public tests. Ordinary companies whose share prices assumed fast AI savings will be marked down too.
2. **Spending falls before the defaults.** Share prices can fall quickly, and households feel it within months. Credit problems take longer, and I'd expect them to show in 2028–29, when the boom's borrowing needs refinancing.
3. **The losses land unevenly.** Shareholders who bought late, the most indebted specialist providers (CoreWeave is the one I'd watch), the junior lenders behind data-centre developers, and funds that concentrated on a few borrowers take most of the pain. The large technology companies can absorb their mistakes, but even they face slower growth in the businesses AI is disrupting.
4. **The technology keeps improving.** Cheaper, better AI keeps arriving throughout. Someone else ends up owning the data centres at a lower price, and the users benefit.

What would make me wrong

My view weakens if:

- Revenue keeps growing faster than computing costs,
- Business AI budgets start growing again rather than being rationed,
- The listing prices well above the last private round,
- The riskiest AI loans keep finding lenders without demanding higher interest,
- and the supplier guarantees expire without being called.

If all of that happens through 2027, the growth will have caught up with the promises, and I'll have been wrong.

What it means if you run a business

This isn't investment advice, and I'm not telling anyone what to buy. But if you run a business that uses AI rather than sells it, such as an accounting firm, here's what I'd do:

1. **Don't wait for the bubble to burst before adopting AI.** The tools will get cheaper and better whether or not their makers' valuations survive. The hard part is the management and the process, and that takes time to build. My free guide sets out where to start (accrualintentions.com/guide).
2. **Read your AI contracts like a lease.** Avoid long, fixed commitments at today's prices. Look for automatic price rises, usage tiers that ratchet up, and minimum spends. Keep the right to leave. Prices are more likely to fall than rise, but a struggling supplier may try the opposite.
3. **Keep the keys.** Make sure your data, processes, instructions and client records live in systems you control and can export, not inside one AI tool. If that supplier changes its terms, gets bought or disappears, you should lose a tool, not your know-how. Work out now how you'd switch. This is one of the problems I'm building Ukoru to solve.
4. **Invest in your client relationships, not just your AI.** When intelligence is cheap, the value goes to whoever the client trusts and talks to. Spend the time AI frees up on understanding clients better and improving their experience. That's where the value will land, and in most cases that's you.
5. **Plan for a recession anyway.** If the recession risk turns real, your clients' cash flow will tighten before the AI story is resolved. That's a reason to have those conversations with them now.
6. **Get ready for the opportunity.** If it does crash, businesses that use AI inherit the build-out: more intelligence than they could ever have afforded, at a fraction of today's price. There may be disruption first, as services are withdrawn, suppliers merge and subsidised prices end. The firms that have already done the management work will be the ones ready to use it.

11

What I'll be watching

Six signals will tell me whether this is heading for an orderly adjustment or something worse. Each one watches the company or link most likely to crack first. One on its own may be noise. Three together would make me more confident in my prediction.

Signal	Why this one	My reading (as of 3rd Nov 2026)	What would worry me
Anthropic's listing price	The first public price for a leading AI model maker	Last private price $965bn (May 2026); talk of $1.5–2tn	Priced below $965bn, or delayed or pulled
OpenAI's fundraising	The AI company most dependent on new money	About $11bn of a $30bn target, at $1.25tn	A round below the $852bn March valuation, or another cut target
CoreWeave's borrowing costs	The most indebted specialist: the canary in the mine	2031 bonds yielding about 16%; insuring its debt costs over 10% a year (October 2026)	Failing to replace a loan, or replacing it only at a much higher rate, or lenders demanding more even while AI shares rally
Big tech cash after investment	Whether the giants can pay for it themselves	Meta's July–Sept free cash flow −$2.1bn; Oracle selling $20bn of new shares and bonds	Two or more of the largest spenders with negative free cash flow for two quarters running
Supplier guarantees and data-centre refinancing	Where credit stress would show first	None called that I've seen; Nexus Apex loans due February 2029	Any guarantee called, or a data-centre loan restructured at refinancing
Household spending as AI shares fall	The route from a share-price fall to a recession	Shares a record 48% of US households' financial assets; Nasdaq down about 10% from its peak in early November	Spending by higher-income households slowing while AI shares stay down

I don't plan to keep this up-to-date in this paper. I'll update my own version of this table as each signal moves. You can ask your own AI tool to check on these signals whenever you like!

12

Conclusion: the technology can win and the money can still lose

AI doesn't have to fail for this bubble to burst. I think AI will succeed. I've seen what it can do in my own experiment, my businesses, and in the businesses I work with, and I certainly expect it to change how every accounting firm works.

But success for the technology isn't the same as success for the people paying for it. Railways, telephone networks and the internet all changed the world. The tracks, the cables and the fibre all got used. Plenty of the people who paid to build them still lost their money.

What's different this time is the size of the promises, and how tightly they're tied together. Suppliers lend to their customers, customers invest in their suppliers, and funds lend to all of them. If it all stopped today, I think most of the damage would stay with the people who chose to take the risk. The longer promises keep growing faster than the cash to pay them, the further a fall would travel.

And here's the part I keep coming back to. The better AI gets, the faster its price falls, and the harder it becomes to earn back what was promised to build it. That's great news for anyone who uses AI. Whatever happens to the giants' share prices, the rest of us still get to stand on their shoulders: we're all standing on the shoulders of gAInts. That's also why I'd spend less time worrying about which AI company wins, and more time making sure my business is the one its clients trust.

> *Standing on the shoulders of gAInts* is also the working title of the book I'm writing next, about how businesses, their teams and individuals can do exactly that.

So, are we in an AI bubble? Almost certainly. When the music stops, the question won't be whether AI was worth building. It'll be who can afford what they promised to pay.

Appendix: the organisations in this paper

Each organisation appears for the point it shows most clearly. Anthropic appears most because as of 3rd Nov 2026 it was the best-documented AI model maker in this paper (the only one with a public prospectus) and, on the numbers, the best-placed of those I looked at. If even the strongest needs things to go right, the rest need more.

Organisation	Why it's in the paper	Sections
Anthropic	Demand is real, but even the best-placed AI company needs growth or new money	2, 3, 4, 8
OpenAI	Promises far ahead of cash; investors cooling	2, 3, 4, 8
Alphabet (Google)	AI squeezing the profits that pay for AI	4, 8
Amazon	Investor in OpenAI, shareholder in and supplier to Anthropic; blocking outside shopping agents	4, 8
Meta	A rich spender now borrowing; AI both helping and threatening its advertising	3, 4, 8
Microsoft	Shareholder in and supplier to Anthropic; struggling to charge extra for AI	4, 8
Oracle	The big tech company leaning hardest on borrowing	4, 8
Nvidia	Even the shovel seller isn't safe	4, 6
Broadcom	How a supplier guarantee works	4
SoftBank	OpenAI investor and data-centre partner	4
CoreWeave	The most indebted specialist AI cloud provider	3, 4, 11
Bank of America, Citi, Morgan Stanley and Blackstone	Lenders behind the Broadcom–Anthropic financing; Blackstone also took on TD Bank's data-centre risk	4, 5
Eagle Point funds and Nexus Apex	One AI borrower making up a large slice of a fund	5, 7
CDPQ	Pension money in AI infrastructure lending	5
Athene (with Apollo)	Insurance money guaranteeing equipment purchases	5, 7
TD Bank	A bank selling off data-centre risk	5
AT&T	Users win when AI gets cheaper	3, 8
Uber and Accenture (with Meta)	Businesses rationing AI	3

Organisation	Why it's in the paper	Sections
Walmart	Welcoming the AI shopping agents Amazon blocks	8
Perplexity	In a legal fight with Amazon over its AI shopping browser	8
Thomson Reuters and RELX	Established software that survived the panic	9
Dassault Systèmes and HubSpot	Where the software fear did bite	9
ZoomInfo	Giving AI away to keep customers	9
Cisco (2001)	What happens to a shovel seller when a boom ends	6
Caterpillar	Power-equipment suppliers whose record order books depend on the build-out continuing	4, 6
Bitcoin miners, including Core Scientific	Machines that kept working while their value collapsed	5
Sichuan Binji Technology	A failed Chinese AI chip financing	5
China's STAR 50 index	An AI share-price bubble that has already burst	1, 6

Research and data come from, among others, PwC, KPMG, EY-Parthenon, the Federal Reserve, the Bank for International Settlements, Oxford Economics, Goldman Sachs, Bank of America, Rhodium Group, Hashrate Index, Pew Research, Blue Ridge Partners, Revenera and UHERO. Links are in the sources.

Plain-English glossary

Term	What it means here
Bond yield	The yearly return lenders demand for holding a company's bonds. The riskier the company, the higher it goes.
Credit rating, "junk"	A grade from agencies such as S&P. BBB- is the lowest "investment grade"; anything below counts as high-risk, or "junk".
Default insurance (credit default swap)	A contract that pays out if a borrower doesn't repay. Its yearly cost shows how likely lenders think a default is. Above 10% of the amount insured is generally read as distress.
Commitment	A signed promise to pay later, such as a computing contract or a lease. Not debt, and some can be cancelled.
Guarantee (maximum)	A promise to pay if something goes wrong. The maximum is the worst case, not a bill.
Going-concern test	The check auditors make that a business can pay its bills for the next 12 months.
Operating cash flow	Cash the business brings in from its day-to-day work.
Free cash flow	Cash left after investment spending, such as building data centres.
Valuation	What investors say the whole company is worth, based on the price they paid for shares.
Prospectus, listing	The document a company publishes before its shares go on sale to the public, and the sale itself.
Refinancing	Replacing a loan when it ends with a new one, often on different terms.
Senior and junior lenders	Senior lenders get repaid first if things go wrong; junior lenders only after them, so they earn more and lose first.
Dilution	New shares, including staff pay in shares, shrink each existing owner's slice.
Tokens	Units of AI text. AI companies charge by the token.
Computing power	The chips and data centres AI runs on.
Wealth effect	People spend more when their shares and homes rise in value, and less when they fall.

Sources

Research cut-off 4 November 2026. Anthropic figures are unaudited interim results. Calculations and analogies are mine.

Company filings and announcements

- Anthropic, PBC. Form S-1/A, 2 November 2026; figures read via StockTitan's summary.
- Thomson Reuters, third-quarter 2026 results, 3 November 2026.
- RELX, trading update, 22 October 2026.
- Nvidia, fiscal 2027 Q2 Form 10-Q (quarter to 26 July 2026). Broadcom, Form 8-K of 27 October 2026 and 10-Q for the quarter to 2 August 2026. Meta, Q3 2026 results (21 October 2026) and 10-Q. CoreWeave, 10-Q for the quarter to 30 June 2026. Athene Holding, 10-Q for the quarter to 30 June 2026. Eagle Point fund reports at 30 June 2026. Blackstone and partners, $63bn financing announcement, 26 October 2026. CoreWeave financing announcement, 17 May 2024.

Reporting and research

- Dataconomy, OpenAI secures $11bn after cutting funding target, 2 November 2026.
- Reuters (via Investing.com), OpenAI forecasts cash burn near $280bn by 2030, FT reports, 18 September 2026.
- Invezz, Anthropic IPO valuation explained, 27 October 2026. Axios, Anthropic files for its IPO, 26 October 2026.
- Reuters (via AOL), Thomson Reuters revenue rises, shares fall on investment costs, 3 November 2026. City AM, Relx defies AI disruption fears, 22 October 2026.
- Hardman & Co, SaaSpocalypse: how AI is reshaping software valuations. Constellation Research, SaaS: apocalypse over, 24 August 2026.
- Fierce Network, Open models are driving AT&T's AI tokenomics strategy, 2026.
- Fiona Chen and James Stratton, Artificial Intelligence in the Firm: Bottlenecks in Software Production, Harvard working paper, version of 4 August 2026.
- Bridgewater, The macro implications of the AI capex boom. Investing.com, AI capex risk cuts both ways (Goldman Sachs estimates). Flossbach von Storch Research Institute, How much does US growth depend on AI investment?
- Federal Reserve Bank of Chicago, Tail risk for banks posed by investments in generative AI, February 2026. BIS Bulletin 140, Building AI on credit: structures and spillovers, 27 October 2026. Octus, SRT weekly wrap, on TD Bank's data-centre risk transfers. Bloomberg, TD Bank mulls hedging data-centre debt, 22 April 2026. OECD Economics Department Working Paper 361, After the telecommunications bubble (Lenain and Paltridge, 2003). IMF Working Paper 18/8, Regulatory cycles (Dagher, 2018).
- Prof G, 2026 predictions (Scott Galloway). 24/7 Wall St, Galloway predicts AI valuations will fall 50–70%, 5 June 2026. Ed Elson, A tale of two bubbles, 20 October 2026. The Next Web, BIS warns an AI bust could hit credit markets, 28 June 2026. FT (via Resultsense), Banks try to offload AI data-centre debt, May 2026. Caterpillar 10-Q, quarter to 30 September 2026.
- Rhodium Group, Examining China's AI financing, September 2026. Hashrate Index ASIC price data, via Cointelegraph, December 2022. Federal Reserve, Wealth heterogeneity and consumer spending, August 2025. Reuters (via bdnews24), Meta raises $15bn in five-part US bond sale, 3 November 2026.
- Bloomberg (via Yahoo Finance), Goldman restricts traders and bankers from using some AI tools, 2 November 2026. Bloomberg, Cisco shares finally top dot-com record, 10 December 2025.
- Oracle 10-Q, quarter to 31 August 2026. Oracle investor update, October 2026. Meta 10-Q, quarter to 30 September 2026. Eagle Point Enhanced Income Trust portfolio filing, June 2026. Eagle Point, Nexus Apex financing announcement, March 2026.
- IMF, Global Financial Stability Report briefing, 13 October 2026. Reuters (via Khaleej Times), Banks and Blackstone line up $63bn for Broadcom chips for Anthropic, 27 October 2026. CNBC, Michael Burry accuses AI hyperscalers of boosting earnings, 11 November 2025. Quartz, Ed Zitron on OpenAI as the AI bubble's Lehman Brothers, July 2026.

- Ed Elson, When the music stops, Simply Put, 3 November 2026. TechCrunch, Anthropic expected to receive high junk rating, 30 October 2026. GTR, Anthropic's $15bn credit facility. Yahoo Finance, Anthropic's $430bn of locked-in computing.
- Eagle Point semiannual reports, 30 June 2026: Credit Company, Institutional Income Fund, Income Company. CoreWeave, $7.5bn debt financing, 17 May 2024. Octus, Investors warn data center SRTs function as asset gathering, 17 July 2026.
- CNBC, CoreWeave credit default swaps hit distress levels, 15 October 2026. CNBC, Janus Henderson manager owns CoreWeave debt, 20 October 2026.
- CNBC, Oracle files to sell $20bn in new stock and bonds, 2 November 2026. Bloomberg, Oracle raises $20bn as Meta weighs new debt, 2 November 2026. Technology.org, Oracle, Alphabet and Meta tap bond markets, 4 November 2026. Oracle results release, March 2026.
- Bloomberg, OpenAI valued at $852bn after completing $122bn round, 31 March 2026. Verdict via Yahoo Finance, OpenAI closes $122bn round. AI Funding Tracker, Is OpenAI losing money? (compiles press reports).
- Reuters via BusinessDay, Global tech investors should take heed of China's AI meltdown, 26 October 2026.
- Moby via Yahoo Finance, Dassault Systèmes tumbles as AI fears hit growth story. Impact Pricing, Wall Street is sending software companies a message about AI pricing.
- Alphabet 10-Q for the quarter to 30 September 2026. Security Boulevard, Alphabet Q3 2026: search ad growth slowed, queries did not. PYMNTS, Google's next big bet.
- Nvidia 10-Q for the quarter to 26 July 2026. Nvidia results release, Q2 fiscal 2027. Wall Street Debug, Nvidia's quarter read from the filings.
- Cisco 10-Q, quarter to 28 April 2001. Cisco annual report, 2001.
- TechCrunch, Companies are scrambling to stop employees maxing out AI budgets, 24 June 2026. Faros, Why companies are limiting AI spending. The Information, via Digg, Meta moves to curb employee AI usage.
- Bank of America research, via Yahoo Finance, 3 November 2026. Reuters (McGeever), via The Star and via Kuwait Times. Fiduciary Trust, Q4 2026 outlook. Man Group, When the AI bubble bursts, don't count on the US consumer.
- PwC, 29th Global CEO Survey, January 2026. CFO Dive, EY-Parthenon CEO Outlook, 1 October 2026. BERI, citing KPMG Q2 2026 AI Pulse.
- Blue Ridge Partners, The five keys to AI pricing. CFOtech, Revenera monetisation research. Tropic, The AI tax.
- Yahoo Finance, Amazon's ad revenue growth decelerates. Adweek, Amazon says external AI agents aren't hurting its ad business. Digital Commerce 360, Amazon Q3 2026.
- Meta Q3 2026 earnings call highlights, via Yahoo Finance. Carbon Finance, Meta earnings review Q3 FY2026.
- Microsoft 2026 proxy statement. SAMexpert, Microsoft 365 E7 licensing guide. GoSearch, Microsoft Copilot pricing 2026.

This paper is research and opinion, not personal investment advice.

ALSO BY ALEXIS KINGSBURY

Accrual Intentions

The book behind this paper. #1 in Accounting & Finance on Amazon.

www.accrualintentions.com

Accountancy Playbooks

30 complete playbooks. One for every bottleneck. Accountancy-specific, step-by-step, immediately usable.

www.accountancyplaybooks.com

Take the diagnostic

Answer 30 quick questions and receive a personalised report showing where you are in the six stages of practice growth. It takes 6 minutes.

www.accountancyplaybooks.com/diagnostic

Ukoru

AI-powered management: goals, guidance, AI workflows, employee feedback and continuous improvement in a single platform.

www.ukoru.com

AirManual

Process documentation for accountancy practices. Rated 4.9 by practices across the UK.

www.airmanual.co

More from Alexis

The De-stress Your Business podcast, and speaking.

www.destressyourbusiness.com

www.alexiskingsbury.com