From Anecdote to Accounting: measuring the AI boom

All firm-level figures are computed from company filings available as at 12 June 2026.

The world is making one of the largest capital bets in history, and arguing about it almost entirely by anecdote. On their own guidance, the big cloud companies will spend several hundred billion dollars this year on the chips, data centres and electricity that train and run artificial intelligence. To grasp the scale, set it beside the last great boom: by some Federal Reserve staff estimates, AI-related computing has lately added more to America's economic *growth* than information technology did at the peak of the dot-com years. That is a share of growth, not the size of the sector. But it is that share a sudden stop would take away, which is why trillions of dollars now ride on it.

So: wisdom, or mania? The optimists point to the technology. It is plainly useful, it is improving fast, and the firms spending the money are among the most profitable in history. The pessimists point to the money: the outlays dwarf any revenue yet visible, they lean more and more on debt at the edges, and there is a faint, familiar smell of cash moving in circles.

This essay proposes a way to weigh this bet. Rather than providing a yes-or-no verdict on the question: "Is it a bubble?", it offers three measures, each built from public data and open to review and update over time. Together they profile the bet: where it stands and where it is heading. A public dashboard already today tracks all three measures.

One big question, decomposed into three

Why not answer the bubble question directly? Not because bubbles cannot be seen coming: the awkward lesson of the limits-to-arbitrage literature is that they often can. The hedge funds of 1999 largely *identified* the dot-com bubble and rode it up anyway, because betting against a rising market is ruinous even when you are eventually proved right. The deeper problem is that the data yield probabilities, not verdicts. Greenwood, Shleifer and You's survey of a century of industry booms, found that sharp price run-ups do not, on average, predict low returns. What they predict is a higher *probability* of a crash, with the odds set by things anyone can observe: how much new stock is issued, how volatile prices are, how steep the climb has been. So the question must be broken into parts, each tied to its own data.

Three essential worries hide inside the question. *How likely* is a correction? *How bad* would it be for the economy, not just for shareholders? And *who* would be hurt? The first asks whether the machinery is paying for itself, because a crash needs fuel: a gap between what is spent and what comes back. The second asks whether the world gains value even where investors do not, because that, not the investors' fate, is what a bust would cost society. The third asks how the building is financed, because a century of history says the damage a falling asset does depends less on its worth than on how much borrowed money was stacked beneath it. Call them the fuel, the salvage and the fuse.

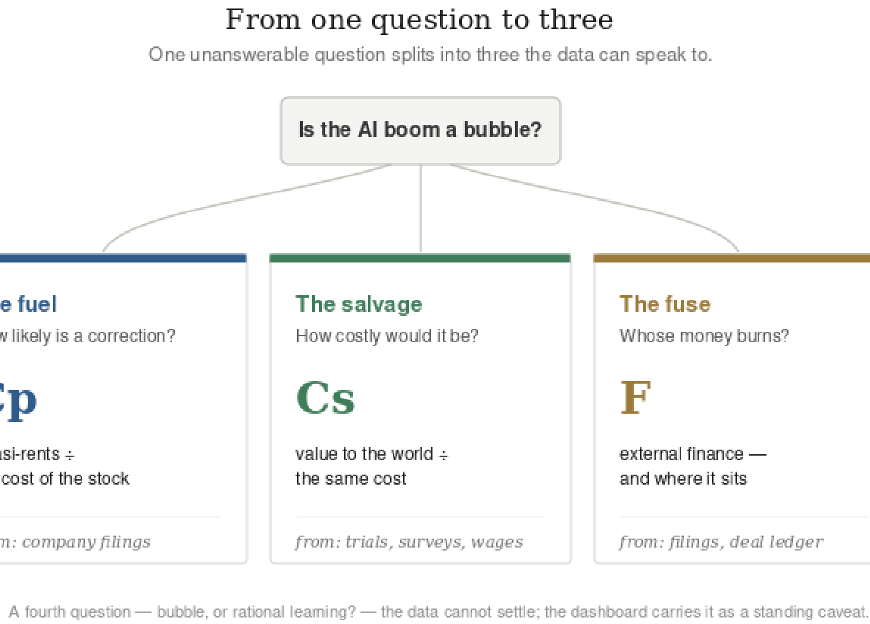

From one question to three

Figure 1. One unanswerable question becomes three that the data can speak to.

The fuel: is the machinery paying for itself?

Start with the comparison everyone has made. On one side, headline capital spending: several hundred billion dollars a year. On the other, the revenue customers visibly pay for AI: a few tens of billions. The conclusion writes itself: the boom is burning roughly twenty dollars to earn one. It is the most-quoted fact in the debate, but it is the wrong sum, because both halves are mismeasured and the comparison premature. Repairing three inherent faults in this logic builds the first measure.

The first fault is the denominator. A boom is not paid for out of one year's spending; the proper charge is the annual *cost of owning* the capital already in the ground, the user cost: the return investors require on the money tied up, plus depreciation. For most infrastructure depreciation is an afterthought. For AI it is the main event. Railway track laid in the 1840s carried trains for a century; an AI chip is obsolete within a few years. Chips age like fish, not wine. An asset losing, say, a quarter of its value a year must earn its keep far faster than any boom before it, and the wear-out rate is itself disputed: company accounts imply five- or six-year lives, sceptics two or three.

The second fault is the numerator. What must cover the user cost is not revenue but *quasi-rents*: the revenue the capital earns once running costs, chiefly electricity and operations, are stripped out. Also, the visible subscription revenue is too narrow, because much of what AI earns comes indirectly, through assistants folded into office software sold by the seat, through internal productivity, through rented-out computing. Count only subscriptions and the boom looks madder than it is; forget to subtract running costs and it looks saner.

The third fault is the comparison itself. Every young general-purpose technology shows a wide gap between what is invested and what it earns; a gap alone proves nothing. What carries information is the direction of travel. A coverage ratio of a fifth that doubles every other year is a technology growing into its capital; the same ratio sitting still while the capital stock compounds is a problem. Pástor and Veronesi made the formal point: when a new technology's payoff is genuinely uncertain, high valuations can be rational, and a later fall in relevant share prices need not signal disappointment. It can be the price of risk: as a niche bet grows into something that moves the whole market, the premium investors demand rises. The bubble is clear only in hindsight.

So the first measure, *private coverage*, asks what fraction of the annual cost of owning the AI capital stock is met by the earnings that capital throws off. Read it not as a verdict but as a level held loosely, with the trajectory watched closely.

How we measure it. **Private coverage = quasi-rents ÷ user cost.** The user cost is $\boldsymbol{\Sigma\ (r + \delta) \times K}$: for each asset class (chips, buildings, power), the required return plus the depreciation rate, times the capital stock, from audited filings. Quasi-rents are AI-attributable revenue, margin-adjusted for running costs, each source tagged by the strength of its evidence. Depreciation enters as a band, never a point; the output is a range with a four-quarter trajectory.

The salvage: is the world getting value even if investors aren't?

A new technology can make society richer and its own investors poorer at the same time, and often has. Britain's railway mania of the 1840s ruined thousands of shareholders and left a network that powered the country for a century. The telecoms frenzy of the late 1990s wiped out hundreds of billions of dollars, and the "dark fibre" laid in the panic went on to carry the internet for almost nothing. Investor returns and social returns are not the same variable.

This idea is called appropriability and was set out by Arrow in 1962; Nordhaus, studying the United States from 1948 to 2001, estimated that innovators captured only about two per cent of the social value their innovations created. The other ninety-eight per cent leaked to customers and competitors as lower prices, time saved and better quality: real value created, often large, that cannot service any loan. Whether AI's owners can beat this historical low share will probably be shaped by the battles over closed models, bundling and lock-in we see unfolding today. The hours AI saves a nurse, an accountant or a shopkeeper does not directly appear on anyone's balance sheet.

This social value is the second measure. It does little to prevent a correction, but it is almost the whole story of what one such correction would cost the world. A bust that follows real value creation leaves cheap infrastructure, trained people and a more productive economy: the railway outcome. A bust with little underneath leaves holes in the ground. The trouble is that this quantity, perhaps the most important in the debate, is difficult to measure. The ingredients for a measurement are appearing: randomised trials of what AI does to productivity, data on who uses it in which job, and wages with which to value the time it saves, etc.

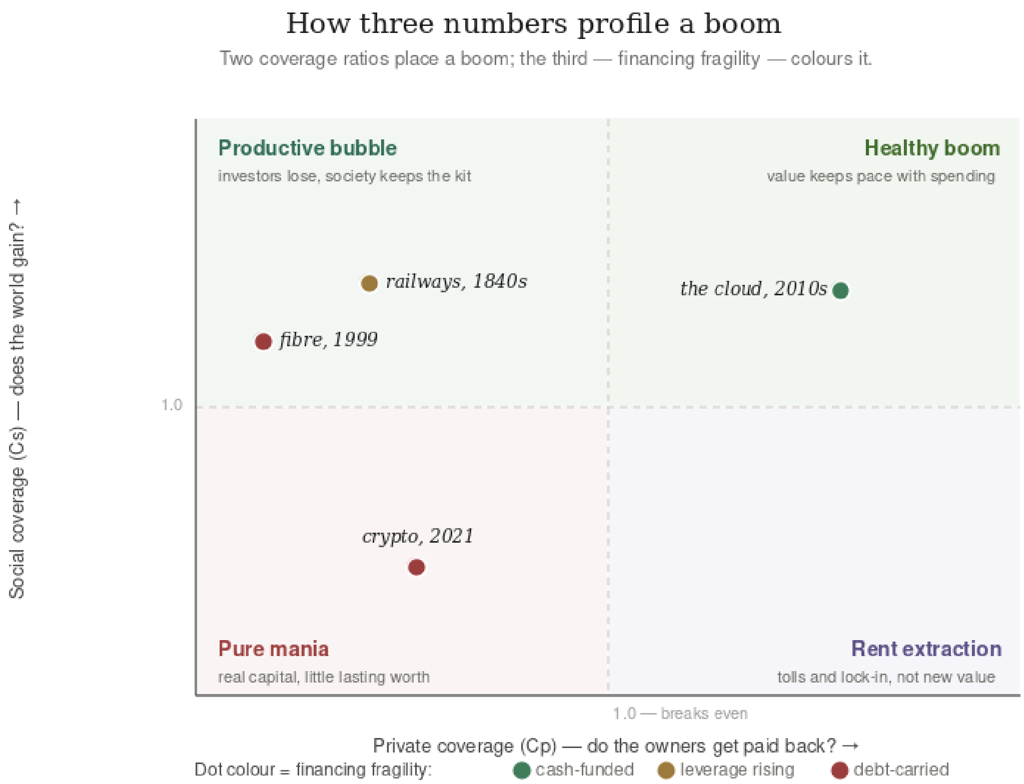

How three numbers profile a boom

Figure 2. Two coverage ratios locate a boom on the map; the financing fuse colours it.

How we measure it. **Social coverage = social surplus ÷ user cost**, against the same denominator. The surplus is summed across occupations as **Σ (productivity gain × hours × wage × adoption share)**, applying the measured gain from the trial literature to the hours and wages of each job, scaled by the share using the tools. It is deliberately gross, before harms such as displacement and energy use, which are tracked beside it rather than netted in. So it leans high, and it is the least precise number on the dashboard. It still, however, beats the blank it replaces.

The fuse: whose money burns?

The third question, on the evidence, decides whether the fuel ever catches fire. The dot-com crash destroyed trillions of dollars of share value and the real economy barely broke stride; the 2008 housing crash, smaller on paper, almost brought down the world's financial system. The difference was not the asset that fell but the way the assets were financed. Jordà, Schularick and Taylor, tracing 140 years of crises, put it plainly: asset booms financed by credit produce deep and lasting recessions, while booms financed by equity produce disappointed shareholders and little else. Their evidence is household and bank credit, so the parallel to corporate computing is imperfect, but it cuts towards reassurance: the equity-funded boom is the survivable kind. What matters is where the borrowing sits, and whether it is creeping from cash flow towards credit. Good technology is no protection: in the 1920s electricity was a roaring success, yet the pyramid of holding-company debt stacked on the utilities collapsed all the same in 1932.

So the third measure is financing fragility. At this moment the picture is, so far, at the centre, reassuring: the hyperscalers pay for most of their building out of enormous operating cash flow. But the margin is shifting. Measured from filings, capital spending now runs at roughly 60% of operating cash flow at Microsoft, Alphabet and Meta, around 100% at Amazon, and well above 100% at Oracle. At the specialist “neocloud” layer it stops being reassuring: in the first quarter of 2026 CoreWeave spent \$7.7bn of capital against \$3.0bn of operating cash flow, the gap covered by new debt and equity.

It is important to note that we only capture in our financing fragility measure what filings make legible: each firm’s capital spending against its cash flow, plus a ledger of disclosed financing deals. What it does not yet capture is the more worrying web of circularity now forming between the chipmakers, the data-centre operators and the AI labs. A chipmaker takes a stake in a customer that spends the money on that chipmaker’s chips. Capacity is carried off the balance sheet as long-term leases. Debt is raised against chips and against future computing contracts. When a dollar of vendor financing can return as a dollar of vendor revenue, an outsider reading only the filings cannot cleanly separate real demand from financed demand. None of it is improper, and each piece is disclosed somewhere; but it is only flagged, not yet summed in. Folding it into our measures is among the most important work we still have ahead.

How we measure it. **Fragility is taken from the most exposed part of the structure, not the average:** roughly **F = max over segments of (external finance ÷ capital spending)**, raised when the financing is circular, because the danger sits at the edge. Per firm we track trailing capital spending against operating cash flow; beside it a deal ledger gives each externally financed arrangement one row: parties, amount, instrument, and a mandatory link to a primary document.

What the dashboard shows today

Take the three measures in order of confidence. The *fuse* is firmest, because it rests on filings alone: capital spending has crossed the whole of operating cash flow at two of the six largest spenders, and every specialist AI-cloud quarter we have parsed shows external finance, not cash flow, carrying the build. The deal ledger’s most telling early result concerns timing: several of the boom’s famous financing arrangements could be tied to a primary document only weeks after the fact, a measure of how thin the public record still is.

The *fuel* and *salvage* gauges are harder, and provisional while their attribution is stress-tested in public. But the shape is clear. Private earnings recover only a small fraction of the annual cost of the capital: about three-tenths, within a wide band from a tenth to seventenths. Social value runs ahead of that, as it does in every productive boom, but it clears the capital cost only over a long horizon.

It is important, as emphasized earlier, to see the levels as measures to treat with substantial caution; what we expect to matter, over time, is how they move.

AI revenue is growing fast. Microsoft's AI run-rate is up 123% in a year, Anthropic's from one billion dollars to seven. So the earnings ratio is more likely than not to rise over the coming year, even as the capital stock grows another third. If it does, that is evidence for the "growing into its capital" case, and the dashboard will say so. The test bites the other way too: if by mid-2027 the ratio has not risen on a frozen method while the stock has grown by a third, the optimists will have lost their best evidence. Either way the record will be public, and the argument can be had on shared facts.

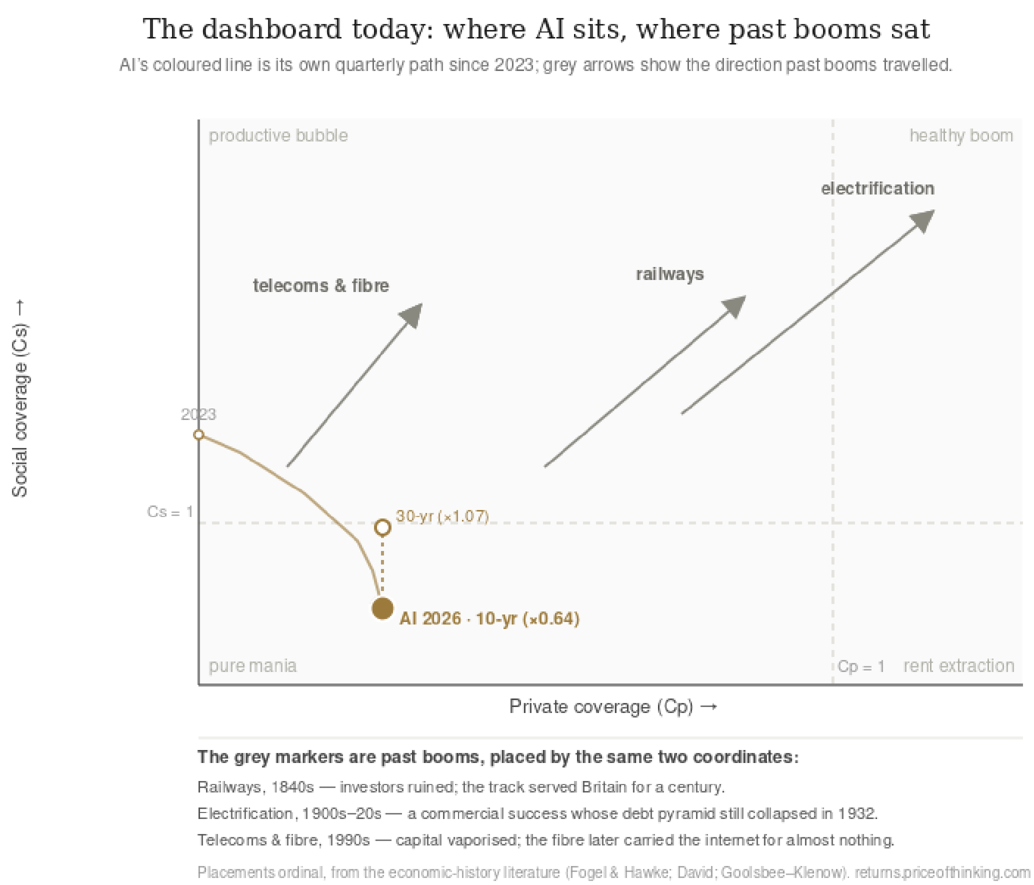

The dashboard today

Figure 3. AI today, in colour, against where past booms ended up, in grey.

The grey markers are past booms, placed by the same two coordinates. Railways sit low and to the right: ruinous for 1840s investors, lasting wealth for the society that inherited the track. Electrification sits high on both axes and still ended in a crash, a reminder that the fuse, not the fuel, did the damage in 1932. Telecoms and fibre sit low for the investors who paid and higher for the world that later used the glass. The placements are ordinal, drawn from the economic-history literature; their job is to show that today's AI marker sits inside a known cloud of outcomes, not off the map.

What measurement cannot do

A perfectly measured, stagnant private-coverage ratio would not on its own establish a high probability of a damaging bubble, because three innocent stories can produce this same picture. Investors may be *learning* what the technology is worth, rationally paying today for a wide spread of tomorrows. *Competition* may be doing the damage: open models and collapsing prices can crush private returns precisely while social value explodes, as in the American railway busts of the 1880s. Or the spending may be *defensive*: an incumbent buying insurance against disruption can rationally fund AI that never pays as a line of its own, because the payoff is a catastrophe avoided. Telling these apart in real time is beyond any measurement programme, and our dashboard does not pretend otherwise.

There are humbler limits too. The boundary of "AI capital" is a judgment, the depreciation rate a controversy carried as a band, revenue attribution inside conglomerates a negotiation with the accounts; levels are uncertain to perhaps a factor of two. What measurement *can* deliver is discipline on the *changes*. With the method frozen and every assumption versioned, the figures move only when the world moves, not when the methodology quietly does. The need is not hypothetical: midway through this project the government's business survey rebroadened its central question about AI, and the measured adoption rate roughly doubled overnight, with not a single firm having changed what it did. A number that can double on a definition is a number to handle with tongs. Consistency, not false precision, is the achievable virtue, and the one on which the national accounts have always rested.

The morning-after question

The Victorians who lost their shirts on railway shares never got them back; what Britain got was the railways. If this boom busts, the question that matters the morning after is not who called it, but how much of what was built still serves the world that paid for it. AI's version carries a twist its ancestors lacked: an unusually large share of its capital sits in the asset that wears out fastest, the chip. The durable residue, the power, the grids, the buildings and the engineers who know how to run them, may turn out to be the boring fraction of the bill, and the most valuable to society.

The bet of this project, offered to a profession that prizes the falsifiable, is that the bubble debate can move from anecdote to accounting without pretending to prophecy. If the coverage ratio closes as today's valuations require, the dashboard will document the bubble thesis failing, filing by filing. If instead it stalls while the leverage migrates outward and the circular deals multiply, that record will be there too: checkable, doubted, improved. Either way, the economist's contribution to the defining capital question of the decade should be what it has always been at its best: not prophecy, but an honest ledger.

A note on method. The full construction (the user-cost denominator, the quasi-rent and social-value numerators, the assumption bands, the fragility decomposition and the deal-ledger schema) is set out in a companion working paper and, with live figures, on the dashboard's methodology page. Every published number traces to a public primary source; the contestable assumptions live in one versioned file, and changing any restates the entire back-series, so a movement is never a quiet change of method.

References: Arrow (1962); Brunnermeier & Nagel (2004); Brynjolfsson, Li & Raymond (2025); Cui et al. (2026); Dell'Acqua et al. (2023/2026); Greenwood, Shleifer & You (2019); Jordà, Schularick & Taylor (2015); Jorgenson (1963); Minsky (1986); Nordhaus (2004); Noy & Zhang (2023); Pástor & Veronesi (2009); Shleifer & Vishny (1997); company filings via SEC EDGAR, retrieved June 2026.