



19 December 2025

Why We Worry – Part I

Dear Investors

The global economy is entering a fragile and unusually complex moment. Capital is flooding into artificial intelligence at a pace without historical precedent, even as the economics of that investment look increasingly stretched. The US labour market is showing increasing signs of strain. And the political narrative that has underpinned market confidence - the “Trump Bubble” – appears to be unravelling.

To give these issues the depth they warrant, we are releasing this analysis as a three-part series.

- Part I explores what we believe is the central flaw in the artificial intelligence boom and why the economics underpinning today’s infrastructure buildout may be far weaker than consensus suggests.
- Part II examines the growing cracks in the US labour market and the rising probability of a layoff-driven recession.
- Part III evaluates the unwinding of the “Trump Bubble” – the political-economic feedback loop that has quietly powered sentiment, spending, and risk-taking over the past year.

Each of these forces is important in isolation. But taken together, they form a set of risks with the potential to shift the trajectory of markets abruptly. Our goal in this series is not to forecast precise outcomes, but to highlight the fault lines now most relevant for investors and to explain how we are positioning in response.

Searching for the Flaw in the AI Narrative

In *Soros on Soros: Staying Ahead of the Curve*, the famed money manager describes how his search for flaws in ideas provides him with a competitive advantage:

I look for the flaw in every investment thesis. When I find it, I am reassured. As long as I can see only the positive side I worry...My sense of insecurity is satisfied when I know what the flaw is. It doesn't make me discard the thesis. Rather, I can play it with greater confidence because I know what is wrong with it while the market does not. I am ahead of the curve...I'm only rich because I know when I'm wrong.

It’s an illuminating framework, and it raises the obvious question: what is the flaw in today’s dominant AI narrative?

In our judgment, we have become increasingly concerned about the underlying economics of AI. The scale of the infrastructure build-out is unprecedented, but the incremental returns remain small. The industry leader, Google, provides a clear illustration of this imbalance.

The Economics of AI Are Becoming Harder to Justify

This year alone, Google will spend roughly \$60 billion more in annualised capex than it did before ChatGPT launched. Since late 2022, the company has deployed an additional \$85 billion in cumulative capex on AI-related development. With similar spending levels expected next year, Google’s capex now exceeds its net profit – a sharp departure from pre-AI years, when capex represented only about 25% of profit.

In effect, a business once celebrated as a capital-light software platform is being pulled toward becoming a capital-intensive infrastructure and construction company. And it is doing so without a clear line-of-sight to commensurate financial returns.

What is Google receiving in return for this extraordinary level of investment? At present, Google processes roughly 1.4 quadrillion AI tokens per month. If we make a simplifying assumption and apply Google’s API input pricing across all of those tokens, the result is an additional \$21 billion of annualised revenue.



For context, this is not an especially compelling trade-off: \$85 billion of incremental capex for \$21 billion of low-margin revenue. In effect, Google is deploying vast sums of capital for what amounts to a modest 5% uplift in annual revenue, and materially lower returns than its core search and advertising franchise generates. A major outlay for just a 5% uplift in annual revenues doesn't sound like a great use of capital to us.

And if this is the underlying economic reality for the industry leader, it is difficult to see how outcomes will be more favourable for its competitors. Over time, we doubt that the return on capital employed (ROCE) from datacentres will meaningfully improve from today's levels.

The Coming Commoditisation of Large Language Models

The more we study the space, the more convinced we are that large language models (LLMs) are on a path toward commoditisation, especially for non-market leading models. Once models reach a threshold of "good enough" performance for their intended tasks – whether automating tasks, powering chat interfaces or search requests – competition will naturally shift toward cost. There doesn't seem to be the same benefits from scale as in the cloud business. Most customers will not care which model is marginally better, in much the same way travellers care more about airfare than the nuanced engineering differences between aircraft. When functionality converges, providers compete on price.

In such an environment, it is hard to see how the economics of AI improve meaningfully from here.

At some point, absent a radical technological breakthrough, the rate of AI usage growth will begin to taper. We increasingly believe that there will be one major wave of corporate AI automation, particularly across back-office and administrative functions, before growth normalises. Historical base rates support this view: whether electricity adoption or internet usage, early-stage "hockey stick" narratives tend to give way to more modest, predictable curves.

As noted in the book, *Capital Account: A Money Manager's Reports from a Turbulent Decade (1993-2002)*:

"The mantra of the late 1990s was that demand for telecoms bandwidth was doubling every three months... demand for data-transmission did expand at about this rate for a couple of years when the Internet first took off. By 1997, however, real growth had slowed, though it continued to double every year. At the time, these ludicrous projections were justified by the prospect of imminent 'killer apps'... which always beckoned but never quite arrived."

The parallels to today's AI exuberance are difficult to ignore.

While NVIDIA's CEO, Jensen Huang, has pointed to robotics and autonomous vehicles as the next "killer apps," we think these will take far longer to materialise than markets expect. The bottlenecks are not model quality, but the underlying physical technology.

We're also not overly concerned about the shareholder-unfriendly accounting choices that Dr Michael Burry has highlighted in his recent publications. Whether companies take generous depreciation assumptions or shift more equity compensation to employees, these practices are unlikely to determine the timing of the cycle. They may make the optics of profitability look better, but they don't change the underlying economics.

The reality is far simpler: being early in markets is indistinguishable from being wrong. Bernie Madoff's Ponzi scheme persisted for years until the inflows supporting the outflows could no longer keep pace. Similarly, we expect the AI bubble will continue as long as capital keeps flowing in. Company executives compensated in equity, and investment bankers underwriting the next wave of datacentre and chip deals, all benefit from the continuation of the narrative. Incentives remain firmly aligned to "keep the music playing."

And critically, at present it is not profit-seeking that is driving Big Tech's capex behaviour – it is fear. When a technology carries the possibility of existential disruption, it can be rational to invest aggressively to protect profits so as not to lose them. The dynamic resembles a prisoner's dilemma: each firm keeps spending until all firms



simultaneously stop believing that AI threatens their dominant position. The threat ends only when everyone agrees it has ended.

Unsustainable Capital Cycles and Intensifying Competition

But there are more troubling developments. If Big Tech and data centre operators collectively spend around \$400 billion on AI infrastructure in 2026, then, by our estimate, at least \$80 billion in annual net income would need to be generated to justify that investment. The hurdle is high because processors, which make up the bulk of capex, have a useful life of only about five years. Back-solving this requirement implies that something like 333 million paying users of ChatGPT – roughly the entire US population – would be needed to support such economics.

Today, the numbers fall drastically short. Only around 5% of users (about 20 million people) pay for ChatGPT, and both paid and non-paid user growth has begun to stall in recent months. OpenAI's attempt to introduce advertising as a revenue stream has met fierce consumer resistance. And unlike Google, directing users to websites does not generate economic value for OpenAI. This raises the critical question: how will OpenAI, or any non-advertising-based AI provider, monetise its service at the scale needed?

To put this in broader perspective, an estimated \$500 billion in AI capex between 2022 and 2026 equates to roughly \$1,500 per American – or \$4,000 per household. What other industry has spent this much capital in such a short span in an attempt to extract incremental revenue from the same underlying customer base? And what are the realistic odds of recouping it? At the moment, real wages for the average US consumer are falling by about 2% year over year. We struggle to see how someone earning roughly \$40,000 after tax will allocate \$500 of discretionary income toward AI products. Perhaps some consumers will reduce discretionary spending elsewhere, but we remain sceptical.

One possible avenue for AI to generate meaningful economic value is through labour displacement – reducing the need for firms to hire or retain workers. But based on our conversations with practitioners and industry experts, the technology remains far from achieving this at scale. By most credible assessments, we are still years and several breakthroughs away from this scenario. For now, the mismatch between capital outlays and monetisable economic return remains stark.

A key premise of our AI bubble thesis is that the technology does not make a sudden leap to general intelligence in the near term. As far as we understand, there are structural limitations that make AGI improbable without breakthroughs well beyond current architectures. But we are not dogmatic. If the facts change, we will update our priors accordingly.

In our view, two developments could flip the psychology of corporate decision-makers:

1. *AI adoption growth slows materially*, exposing the weak economics of LLMs, or
2. *A recession emerges*, forcing companies to prioritise cash preservation over technological “defence spending.”

These events are not mutually exclusive – each could trigger the other.

We're also increasingly worried about OpenAI's competitive positioning relative to Google, particularly its inability to monetise its product. Without offering material advantages over Google's Gemini, yet operating at a far higher cost base, OpenAI's business model may ultimately prove unsustainable. The trillions of dollars in future capex commitments across the industry risk becoming little more than a performative show of technological bravado – impressive on paper, but economically hollow. If OpenAI cannot find a way to monetise its technology or alter entrenched consumer habits, then Google, with its integrated ecosystem and superior economics, will simply outlast the capital available to its rivals. Should that scenario unfold, we are confident it would trigger a recession.

There is historical precedent. A similar pattern emerged during the late-1990s dot-com bubble. Telecom operators, despite enormous capital outlays, found their services rapidly commoditised. Usage growth slowed, pricing power collapsed, and the industry could not extract the household revenues required to justify the capex binge. High returns initially attracted more competition, which eventually eroded margins for even the leading “pick-and-shovel”



equipment suppliers. Investor belief in unassailable competitive advantages proved illusory. Once reality set in, the bubble burst and triggered a shallow recession.

Another concern is the intensifying competition in graphics processing units (GPUs). *Capital Account* offers a useful explanation of why this is such a big problem:

As things turned out, the history of the telecoms world in this period provides a near perfect case study of the operation of the capital cycle. To repeat, the theory tells us that when companies trade at a premium in the stock market to their replacement cost, new investment is stimulated; that demand forecasts are inherently unreliable; that when competition is increasing, there is a great danger of supply exceeding demand; that investment bankers will promote excesses and that investors will capitulate to these developments.

Essentially, when companies earn supernormal returns, capital floods in, demand forecasts prove unreliable, and investment bankers amplify the excesses. Ultimately, supply overwhelms demand, and margins collapse.

This is precisely the risk unfolding today. SemiAnalysis notes that Google's TPU infrastructure now rivals NVIDIA's latest commercially available GPUs – at significantly lower cost. Sensing the threat, NVIDIA has resorted to taking equity stakes in companies that depend on its support (OpenAI among them), effectively subsidising its customer base to stave off competition and preserve margins. This is not a sustainable strategy. Amazon's upcoming Trainium 3 chip has also narrowed the performance gap and is likely to be cost-competitive upon release.

With credible alternatives emerging, NVIDIA's 75% gross margins – the foundation of its current valuation – will not hold indefinitely. When investors fully appreciate this, and when that realisation intersects with the economic unsustainability of OpenAI's model, the conditions for a sharp correction may be in place.

Overall, we remain concerned about the economic sustainability of today's AI capex boom. Big Tech can fund datacentre expansion through cash flow and balance sheet strength, but that doesn't make it a good use of capital. Spending \$85 billion for \$20 billion of low-margin, commoditising revenue is not a viable long-run equilibrium. Ultimately, these dynamics risk diluting the very profitability that has made Big Tech the cornerstone of global equity markets.

For now, markets view AI as a tailwind for free cash flow. We are far less convinced. The market also isn't incapable of coming to this conclusion – for instance, it worked out that Oracle's revenue growth was largely profitless. If the market senses that the music will stop, then it could very likely drag us into a recession on the scale of 2000-01.

Kind Regards,
Fawkes Capital Management

Please click [here](#) to subscribe to receive future Fund or market updates.



Fawkes Capital Management Disclaimer

One Funds Services Limited ("OFSL"), ACN 615 523 003, AFSL 493421, is the issuer and trustee of the Fawkes Capital Fund (the "Fund"). The material contained in this communication is general information only and was not prepared by OFSL but has been prepared by Fawkes Capital Management Pty Ltd ("Fawkes"), a Corporate Authorised Representative of One Wholesale Fund Services Ltd ("OWFS"), ACN 159 624 585, AFSL 426503, CAR number 1308574. Fawkes has made every effort to ensure the accuracy and currency of the information contained in this document. However, no warranty is made as to the accuracy or reliability of the information. Investors should consider the Information Memorandum ("IM") dated 24 May 2024 issued by OFSL before making any decision regarding the Fund. The IM contains important information about investing in the Fund and it is important investors obtain and read a copy of the IM before deciding about whether to acquire, continue to hold or dispose of units in the Fund. You should also consult a licensed financial adviser before making an investment decision in relation to the Fund. Past performance is no guarantee of future performance. This report does not take into account a reader's investment objectives, particular needs or financial situation and is general information only to wholesale investors and should not be considered as investment advice and should not be relied on as an investment recommendation.