Montaka
Global Investments

THE AI BUBBLE DEBATE:

WHY INVESTORS SHOULD RE-EXAMINE THE EVIDENCE

2026

By Andrew Macken and the Montaka Investment Team

"Wisdom, itself is often an abstraction associated not with fact or reality but with the man who asserts it and the manner of its assertion.[1]"

These are the words of John Kenneth Galbraith, one of America's foremost economists of the twentieth century, in his famous book, The Great Crash, 1929.

Galbraith was making an important point: we are often persuaded, not by the specific arguments and conclusions, but by the perceived stature of the people putting them forward, and the decisive way they are delivered.

This is exactly what happened in the lead up to 1929. As Galbraith notes, the assurances of sustained stock appreciation appeared convincing but in fact, bore no relation to reality.

As we know, reality eventually caught up. The stock market crashed and was followed by a horrible economic depression.

We now face a similarly challenging period in deciphering rhetoric from reality.

Many persuasive voices – including Michael Burry of the The Big Short fame, and other credible investors – are evoking the ghosts of 1929 (and the later dot com boom) to argue we are now in the midst of a giant AI bubble that will lead to an inevitable bust.

Many market observers now argue that the recent re-rating of major tech companies – including hyperscalers like Alphabet, Meta, Microsoft and Amazon – is tied to concerns about the speculative nature of an AI-driven future. Under this "bubble" case, there is concern that valuations embed aggressive growth in AI workloads, that earnings quality is poor – driven by understated AI chip depreciation, and while funding remains heavily dependent on mature, cyclical business segments like advertising. Recent stock-market signals like sharp pullbacks and concentration of market value among a handful of 'AI darlings' have reinforced fears that we may be experiencing a bubble rather than a durable, structural shift.

But does the bubble hypothesis really reflect reality?

In this whitepaper, we critically evaluate five core premises that underpin prominent arguments that an AI bubble has formed:

1. Hyperscaler cloud revenue growth expectations are too high and will likely decelerate to single digits (SD%) within five years.
2. Barriers to entry in the cloud market are weakening because new players are rapidly entering the space and competing on price.
3. The profitability of new data centre capacity is materially impaired by the rapid depreciation of AI chips.
4. The large-scale capital investments being made in new data centre capacity are funded by lower quality, mature and cyclical advertising businesses.
5. Valuations of today's tech companies are excessive and equity markets are nearing 2000 dot-com levels.

While you could logically conclude from these premises that we are in the midst of an AI bubble, our analysis suggests that many of the underlying assumptions are overstated or unsupported by evidence.

In our view, this leads to an exaggerated, simplistic diagnosis of bubble conditions.

While we also see several instances of excessive exuberance today, we continue to see significant investment opportunity among competitively advantaged players in the AI space.

1: (Galbraith) The Great Crash, 1929 (1954)
2: (Cremers, Pareek) Patient Capital Outperformance, 2016

About Montaka

Montaka's goal is to generate sustainable excess returns by exploiting long-duration market inefficiencies. We observe that advantaged businesses capable of growing reliably over extended time horizons are routinely undervalued by equity markets.

We have built a repeatable process to identify long-duration transformations. But themes alone are not enough. We overlay a rigorous competitive analysis to identify the players within each theme whose competitive advantages are strengthening over time.

These are the opportunities we believe are most likely to be undervalued by the market, a view which we validate through detailed fundamental research and bottom-up analysis.

We construct Montaka's portfolio with high concentration, expressing significant conviction in our largest investments. Montaka's top ten holdings typically represent around 75% of the portfolio.

Empirically, academic research shows that highly-concentrated, long-duration portfolios have demonstrated a persistent ability to generate long-term excess returns[2].

Montaka's investment team brings deep technical expertise and multi-disciplinary insights that differentiate our research. Our backgrounds in areas like engineering, science and financial markets, combined with advanced AI capabilities enable us to understand – not just observe – the rapidly-evolving global trends in which investment opportunities emerge.

Montaka is 100% owned by MFF Capital Investments Limited.

Premise #1 — Hyperscaler cloud revenue growth expectations are too high and will likely decelerate to single-digits (SD%) within five years.

The first claim that bubble proponents make is that revenue growth of hyperscale cloud computing businesses – like Amazon, Microsoft and Alphabet – is about to decelerate rapidly. Therefore, they argue that the long-term revenue and earnings projections of these businesses, which underpin current valuations, are too optimistic.

Yet, when we reconstruct the total addressable markets (TAMs) for cloud computing under different scenarios, we find that for such an extreme deceleration in revenue growth for hyperscalers to materialise, which the bubble proponents expect, it would imply that AI workloads contribute minimal incremental demand beyond the traditional digital transformation (i.e. the structural migration of on-premise workloads to the public cloud – which has been underway since well before the launch of ChatGPT).

That the implied 'AI TAM' is so small appears implausible – given current adoption trends, the economics of model training and inference, the non-linear improvements in AI models, and the structural investment cycle underway .

As such, we regard the revenue deceleration scenario as excessively conservative and therefore low probability.

The growth rates of the cloud businesses within three major hyperscalers are shown below. Also included are recent cloud backlog growth rates – the future commitments already made by customers for services not yet delivered – which are very high and indicate strong future revenues to come. We've also shown how each cloud business is growing as a share of total group gross profits (GP), highlighting increasing significance over time.

Company	Metric	2023	2024	2025	2026*
Amazon	AWS Revenue Growth	13%	19%	19%	24%
	Backlog YoY Growth (as at 3Q25)			22%	
	AWS GP Share of Total Group GP	28%	30%	31%	32%
Microsoft	Azure Revenue Growth	30%	33%	34%	37%
	Backlog YoY Growth (as at 3Q25)			51%	
	Azure GP Share of Total Group GP	17%	19%	21%	25%
Alphabet	GCP Revenue Growth	28%	33%	37%	50%+
	Backlog YoY Growth (as at 3Q25)			82%	
	GCP GP Share of Total Group GP	2%	6%	9%	15%

**internal estimates*
Source: Company Filings; Montaka analysis

For these cloud businesses to decelerate to SD% growth over the next five years, it would imply that global cloud TAM (ex-China) will be not much greater than the original cloud TAM (made prior to the ChatGPT and associated generative AI breakthroughs) from traditional digital transformation alone, with minimal contribution from AI.

Cloud TAM Forecasts (Various Scenarios)

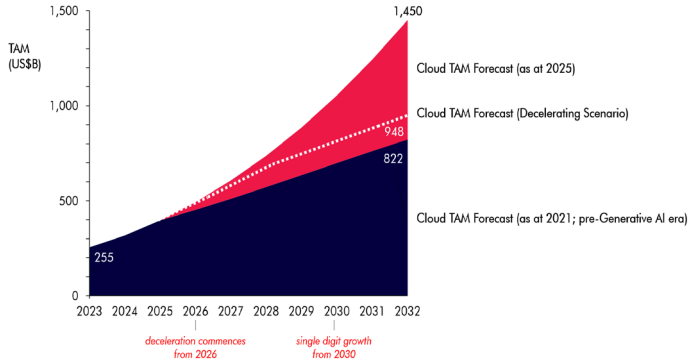

Source: Montaka analysis, IDC

Yet according to IDC's forecasts, AI could lift the cloud TAM from less than $0.5 trillion per annum today to nearly $2 trillion per annum over the next decade.

In this scenario, as global (ex-China) GDP expands to around $150 trillion per annum, cloud spending rises from well below 0.5% of GDP today to roughly 1.3% by 2034.

These implied levels of cloud spending and penetration do not appear excessive in a world where AI diffuses broadly across the economy. Viewed through this lens, a rapid deceleration to single-digit cloud growth is difficult to reconcile with observable indicators of AI adoption and advancement today.

Premise #2 — Barriers to entry in the cloud market are weakening because new players are rapidly entering the space and competing on price.

The second argument for bubble trouble is that hyperscalers are overvalued because a new generation of cloud providers – such as CoreWeave and Oracle – are entering the market, leading to a deteriorating competitive landscape.

But these new cloud providers are not comparable to the hyperscalers in several important ways.

First, businesses like CoreWeave only offer specialised AI model-training infrastructure. These services are effectively high-performance Graphics Processing Unit (GPU) compute delivered on a usage-based model for customers training large AI models – in essence, GPU-for-rent. This is a very narrow, infrastructure-only segment of the market – in contrast to hyperscalers' ecosystem of products and services.

Secondly, while AI model training currently represents a significant share of AI workloads today, the much larger long-term opportunity is inference – the phase after training that represents the models actually being used to solve problems.

Inference workloads will likely dominate over time as AI is deployed across enterprises, governments and consumer use cases as part of broader applications (that require substantial higher-margin services beyond just training-compute infrastructure).

Cloud TAM Forecast & Penetration

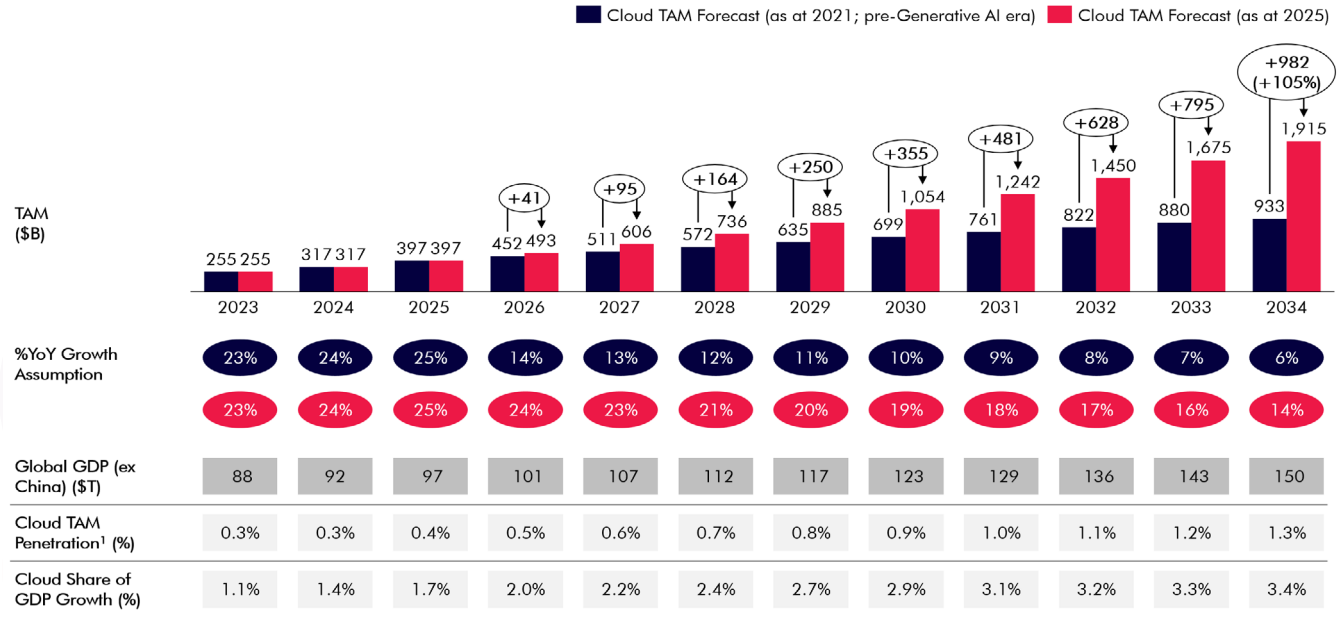

Source: Montaka analysis, IDC | 1. Cloud TAM as a proportion of Global GDP (ex China)

While the new cloud entrants are focused on AI training, hyperscalers, on the other hand, have a much broader mix of businesses, as you can see in the table below:

Hyperscaler Revenue Across the Cloud Stack

	Cloud Workload Types		
Cloud Stack	Traditional Cloud	AI Training	AI Inference
Application Layer	✓		✓
Platform Layer	✓		✓
Infrastructure Layer	✓	✓	✓

Source: Montaka analysis

There are important distinctions to be understood between each of these three lines of business:

- AI Training is the 'GPU rental' business, described above and is the lowest margin, least differentiated, and most commoditised.
- Traditional Cloud represents the broad range of cloud computing services required to deploy applications from the public cloud, including compute, storage, networking, databases, security, authentication and many others. This is a much higher margin line of business that is dominated by the hyperscalers.
- As mentioned, AI Inference refers to running trained AI models in 'production mode' for broader software applications. This is where AI actually creates end-user value. In recent years, there have been concerns that expensive inference costs could erode software margins. However, several factors point in the opposite direction:
    - Inference pricing is often consumption-based. This means that variable revenue (what the user pays) is directly aligned with variable infrastructure costs, eliminating the risk of high-cost users eroding profitability.
    - Many inference use cases generate exceptionally high value relative to cost (e.g., automating workflows or improving decision quality). This means it makes economic sense for customers to continue spending, and in turn, hyperscalers can continue to price these services attractively.
    - Inference costs have fallen by a factor of 10 each year, on average, for an LLM of equivalent performance[3]. To the extent inference costs deflate faster than inference pricing – which is very likely – then inference profit margins will expand over time, all else equal.

Together, these factors suggest that inference is likely to generate high margins for hyperscalers over time.

Fantastic flywheel effects

Importantly, there is clear evidence of positive 'flywheel' dynamics between Traditional Cloud and AI Inference, from which hyperscalers benefit.

Firstly, enterprises typically need to migrate on-premise workloads to public clouds to deploy AI applications, which drives demand for Traditional Cloud services.

Secondly, as enterprises expand their use of AI applications, demand for AI inference rises. This in turn accelerates further cloud migration, a dynamic which reinforces hyperscaler advantages over time.

Indeed, based on Montaka's recent analysis of interviews with industry experts, for every $1 spent on the third-party AI models, an additional $1-4 is typically spent with the hyperscaler on broader cloud services to make the workload functional, reliable, secure and compliant.

So while the cloud market has seen newer entrants such as Oracle and CoreWeave, they are not threatening the higher-margin businesses of the hyperscalers, and therefore do not necessarily undermine the valuations of the hyperscalers ascribed by the market today.

Premise #3 — The profitability of new data centre capacity is materially impaired by the rapid depreciation of AI chips.

The third argument of bubble proponents is that the unprecedented capital investment in data centres will fail to generate acceptable returns and ultimately lead to enormous value destruction.

Combined Capex of Major Tech Companies ($B)

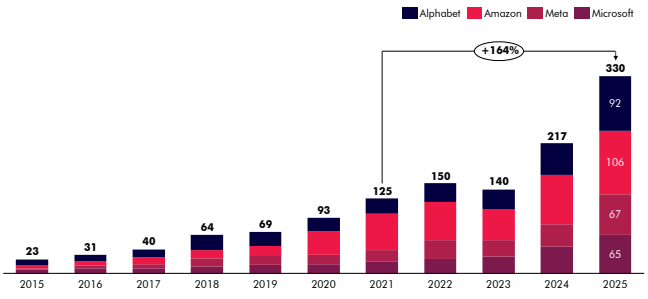

Source: Canalyst Models

There is no question that if cloud computing demand, either first-party or third-party, is not there to support these investments, or the profitability of future revenues is impaired, then these investments will be value-destructive.

We showed in the section above that there is strong demand for cloud computing.

But we agree that investors need to carefully consider how profitably this new capacity will be monetised, including (but not limited to) the impact of depreciation.

When a new capital investment is made, such as a data centre, its value depreciates over time based on its useful life.

Many prominent figures in the investment management industry have raised concerns about the assumed depreciation schedules employed by hyperscalers for compute infrastructure.

The hyperscalers have lengthened depreciation for compute infrastructure to approximately 5-6 years, up from 3-4 years several years ago.

3: (Andreesen Horowitz) Welcome to LLMflation – LLM inference cost is going down fast, November 2024

At the heart of this concern is the fact that a large share of hyperscale data centre capex is concentrated in specialised AI chips that power training and inference. If customers insist on using leading-edge chips for AI workloads, and if new generations of chips are released annually, then the effective economic life of each chip could be far shorter than 5-6 years.

In that scenario, a significant portion of data centres would depreciate much faster than reported, reducing the true profitability of these investments and calling into question the returns on the entire build-out cycle.

But we can resolve this apparent dilemma by making an important distinction between AI model training workloads and AI model inference workloads.

It is true that leading-edge AI chips are typically the most economical for AI model training workloads. Their superior energy efficiency minimises the effective 'capital cost' of training a new model, while their processing speed accelerates time-to-market.

If training were the only AI workload, then it would make sense for chip depreciation schedules to be tied to the rapid product cycles of leading-edge chips – implying much faster depreciation than the 5-6 year schedules used today.

But as AI diffuses more broadly into the economy, it is likely that AI inference workloads will significantly outnumber training workloads.

And unlike training, a much wider range of chips can be (and are) used for AI inference. For example, eight-year-old Google TPUs are still operating at 100% utilization in the Google Cloud Platform (GCP)[4].

Other ways to monetise compute

Additionally, there are several techniques for hyperscalers to monetise new compute more effectively – without changing depreciation schedules – including:

- **AI model distillation** — Where a large foundation AI model transfers much of its knowledge and capability to a smaller, less compute-intensive, model. Distilled models are cheaper to run and can perform inference on older or less powerful chips. This allows hyperscalers to serve growing inference demand using a much broader portion of their existing and planned future compute, thereby improving utilization.
- **Customised chips** — Where the hardware design is explicitly optimised for a specific range of AI models to boost efficiency. These chips allow hyperscalers to extract more value from each unit of compute.
- **High-density data centres** — Where more compute is packed more densely into the physical footprint of the data centre, reducing both fixed overhead costs and variable operating costs (such as cooling and maintenance) per unit of compute, improving the overall economics of the data centre.

The future benefits of R&D spend underestimated

Finally, we believe that bubble proponents are failing to appreciate the optically depressing impact of research and development (R&D) expense on the true economic earnings power of the hyperscalers.

Just as overstating the useful life of chips would inflate earnings (by understating depreciation), understating the useful life of R&D would have the opposite effect: it depresses reported earnings by expensing long-lived investments entirely in the current period.

Most would surely agree that the more than US$200 billion spent annually by Amazon, Microsoft, Alphabet, and Meta – on R&D activities like AI model development, new products, and chip design – will likely bring some future economic benefits.

Economically, these investments behave more like long-lived capital assets. But the accounting rules require all R&D to be expensed in the current period, irrespective of the future benefits that stem from today's spend. This is akin to expensing all capital investments in the current period instead of capitalizing and depreciating over a future useful life.

In fairness to those who make the accounting rules, they are simply siding with conservatism on the basis that future benefits from R&D spend are difficult to know with a high degree of certainty.

And, importantly, the magnitude of this distortion is just as significant as any possible distortion from capital investment depreciation – at least for the hyperscalers – as illustrated by the chart below.

R&D and Capex Share of Gross Profit (GP) (%)

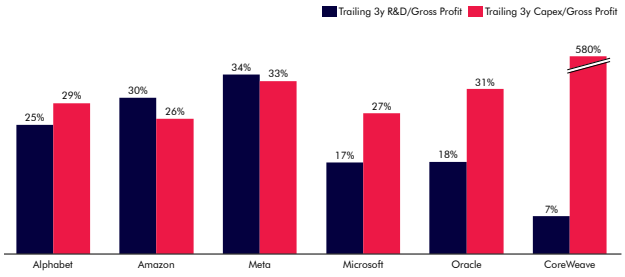

Source: Company Filings, Montaka analysis.
Note: CoreWeave figures are estimates based on three quarters of reported data.

When we consider the evidence, it's clear that bubble proponents' concerns that hyperscaler earnings are being overstated by under-depreciation are overblown. What's more, if we shine the spotlight on their also-huge R&D spend, we can see that their current and future earnings potential is actually understated, something the market is failing to appreciate.

Premise #4 — The large-scale capital investments being made in new data centre capacity are funded by lower quality, mature and cyclical advertising businesses.

The next argument of those espousing a belief that we are in the midst of a bubble is a concern around the funding sustainability of the large advertising businesses attached to hyperscalers like Alphabet, Amazon and Meta.

4: (DC Dynamics) Google says TPU demand is outstripping supply, claims 8yr old hardware iterations have "100% utilization", October 2025

The critics note that digital advertising is cyclical, maturing, and becoming more competitive, creating headwinds in these businesses that will erode the cash flow of the ad-centric hyperscalers in the near term.

Hyperscaler Ad Revenues

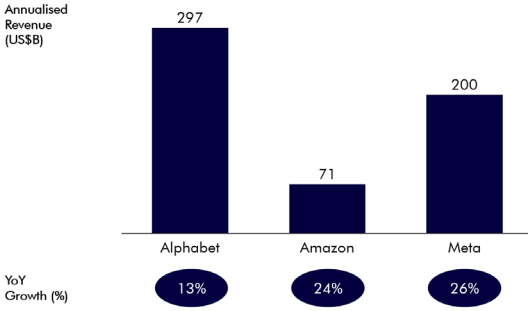

Source: Company Filings (3Q25)

To be sure, many of the observations are valid:

- Historical advertising demand has been cyclical.
- New competitors have entered the digital advertising space, including Walmart, Netflix and Uber, increasing supply.
- Given there are only 24 hours in a day, ad impressions have a natural ceiling.

Yet recent observations show that when the world's most advantaged digital advertising platforms – including Google, Meta, Amazon and even Tencent – have deployed large AI models internally, it has substantially and tangibly improved content-recommendation, ad-targeting, and conversion metrics. These gains have been broad based and, importantly, show no signs today of reaching a ceiling.

Such improvements matter. They directly reduce advertising waste and increase return on ad spend (ROAS). And that means these platforms are inching closer to the theoretical 'holy grail' of personalised performance-based marketing: presenting the right ad to the right person at the right time with a high probability of conversion.

In this world, ad budgets behave less like fixed line items and more like self-funding growth investments. When a merchant can essentially buy revenue through high-performance marketing platforms, the incremental revenue funds its own marketing expense. As a result, ad budgets expand to the point at which the required ROAS is not being met.

From an aggregated economic perspective, we view these platforms as increasingly enabling transactions – connecting the right buyer with the right seller at the right time – that otherwise could not have been made. Said another way, they are effectively eating into the natural (and enormous) dead weight loss in the economy from buyers and sellers not having a way to meet at exactly the right moment.

And this is how it is possible for hyperscaler ad revenues to grow – even in a theoretical world in which, given there are only 24 hours in a day, ad impressions have reached a ceiling. As ad performance increases with AI, the 'price' of ads can continue to grow as the value delivered by these ads grows.

Consider a performance marketing ad conversion rate (CVR) of 5%, for example, which is typical today. This means advertisers need to buy 20 ads for one conversion, on average. At a 10% CVR, the number of ads needed for one conversion reduces to 10 ads, on average – and the advertiser would be willing to pay double the price per ad, all else being equal.

Two Types of Market Power: Parasitic vs Symbiotic Monopolists

At Montaka, we make the explicit distinction between two categories of monopolist (or dominant industry player):

	PARASITIC MONOPOLIST	SYMBIOTIC MONOPOLIST
	The **Parasitic Monopolist** uses its dominant competitive position to extract value at the expense of its customers.	A **Symbiotic Monopolist** operates under the principle of mutual benefit, where the monopolist's success is tied to the success of its customers or partners.
Core Dynamic	The relationship is fundamentally zero-sum. The Parasitic Monopolist's gain is a direct loss for its customers.	The relationship is a positive-sum or 'win-win'. The Symbiotic Monopolist creates value for its customers and participates in that value creation, meaning both parties gain.
Customer Experience	Customers are forced to use the monopolist's service (due to its size or market necessity). Customer profitability can easily be transferred to the Parasitic Monopolist through pricing & other mechanisms.	Customers benefit from using the monopolist's service, despite few alternatives. As the Symbiotic Monopolist creates more value for a customer, it participates and creates more value for itself.
Key Example	Rightmove in the UK. Property agents have little choice but to list vendor properties on the Rightmove platform. Agent profitability has been structurally declining and is approaching zero as Rightmove continues to extract economics.	High-efficacy performance marketing platforms, such as Google, Meta, Amazon and Tencent. Customers benefit from the increasingly better performance of ads, while the platforms benefit from structurally increasing ad prices & new adjacent properties over time.

The distinction between these two categories is important to assessing the durability of growth of the hyperscalers, in our view.

- Accelerated growth of a Parasitic Monopolist, for example, may simply reduce its duration of growth, as customer profitability approaches zero more quickly.
- Accelerated growth of a Symbiotic Monopolist, on the other hand, can be much more sustainable over longer durations, as more value is created for customers over time.

As these platforms increase their mix of high-efficacy performance marketing (and move away from lower performing and cyclical brand advertising), growth in their ad businesses becomes much more structural (and less cyclical), and more tied to improvements in AI – which continue at a rapid pace.

It is also true that much of the capital investment being made by hyperscalers today is in service of third-party enterprises and governments. Through this lens, the underlying demand drivers supporting these investments are highly diversified and tied to global economic activity, rather than the specific non-cloud business verticals of the hyperscalers.

Finally, far from seeing hyperscaler non-cloud businesses as a liability, we view them as important sources of competitive advantage.

Amazon, Microsoft, Alphabet, (as well as Meta and Tencent) have large, valuable core businesses that are clear beneficiaries of larger and more powerful AI models.

Therefore, any 'excess' capacity that might materialise from the data centre buildout over the coming years will more rapidly be absorbed by their internal needs.

So overall, we see the existence of large, tech/AI-enabled non-cloud businesses attached to the hyperscalers, not as a risk, but as a major strategic advantage.

Premise #5 — Valuations of today's tech companies are already excessive and equity markets are nearing 2000 dot-com levels.

Perhaps the bubble proponents' biggest argument is that the market is wildly overvalued.

There is no question that the price-to-earnings (PE) ratio of the S&P 500 index is elevated by historical standards. While we cannot say definitively whether today's PE ratio is fair or not – only time will tell – we can make some definitive observations with respect to historical comparisons:

- Aggregate earnings growth expectations today are above their average. This will help push the PE ratio above average levels, all else being equal.
- Interest rate expectations today are below their average. And this will also help push the PE ratio above average levels. Indeed, it's subtle but important to note that the 2000 dot-com bubble occurred in the midst of a materially higher (and rising) interest rate environment than today. This likely means that if the dot-com bubble was translated into today's interest rate environment, the market PE ratio at the time of the dot-com bubble would be even higher than what it actually was 25 years ago.
- Credit spreads are below their average levels today, which will also help to push PE ratio above average levels.

Putting these together, it is probably fair for today's market PE ratio to be above its long-term average.

S&P500 Index: Price-to-Earnings (PE) Ratio

Source: Bloomberg

But what about individual companies? Aren't they overvalued?

We are not interested in defending every individual case of potential overvaluation, such as Palantir or OpenAI, nor do we defend extreme cases like Mira Murati's Thinking Machines venture, which is reportedly in talks to raise money at a US$50 billion valuation[5] despite not yet having a product, let alone revenues, let alone earnings.

Even market-darling, Nvidia, is facing intensifying competition from custom hyperscaler chips, especially Google's TPU[6], the effects of which remain to be seen, and which could arguably impact its fair value.

Again, time will tell if these valuations are excessive. All that we can say definitively today is that the growth expectations built into these valuations are:

(i) Very high,

(ii) Long-duration, and

(iii) Profitable over time.

And that's why many investors fairly conclude that there is a significant probability that these expectations may ultimately not be met.

We also share concerns about the 'circular' financing arrangements (where one company extends credit to, or invests in, another company, and the recipient company then turns around and buys the creditor's products and services) that have made headlines recently, with OpenAI and Nvidia at the heart of many of these.

Undoubtedly, vendor-financed sources of revenue are of a lower quality and potentially less persistent. Oracle, in particular, appears to be taking on financial risks by making large capital investments primarily on the basis of a single customer's (OpenAI) very large medium-term revenue projections[7].

5: (Bloomberg) Murati's Thinking Machines in Funding Talks at $50 Billion Value, November 2025
6: (Bloomberg) How Google's TPUs Are Giving Nvidia a Run for Its Money, November 2025
7: (WSJ) Oracle, OpenAI Sign $300 Billion Cloud Deal, September 2025

How Nvidia and OpenAI Fuel the AI Money Machine

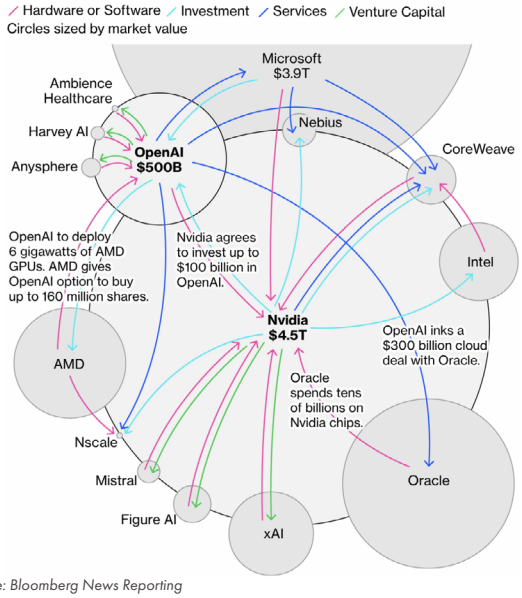

Source: Bloomberg News Reporting

But despite these concerns, we believe that many tech businesses – including the hyperscalers – remain undervalued. Relative to our expectations about the longer-term cloud computing TAM, as well as the durability of growth we see in their non-cloud businesses, today's market-implied expectations embedded in today's stock prices are too conservative. And this represents significant investment opportunity.

It's a mixed bag. But it's always a mixed bag – with some stocks appearing to us as overvalued and some undervalued. (See Montaka's top 10 investee companies for hints on which stocks we assess as materially undervalued).

If one believes that growth will decelerate rapidly and that the profitability of incremental growth is poor, then a bubble conclusion seems logically valid.

If, on the other hand, one sees durable, profitable growth ahead as AI diffuses into large swathes of the economy – especially when adjusting carefully for things like depreciation and R&D – then investment opportunities in selected businesses is a more likely conclusion.

The conclusion from our AI Bubble Audit: No bubble trouble

This whitepaper has sought to fairly assess five premises laid out by many prominent industry figures who believe we are in the midst of another market bubble.

If each of the five premises were true, it would be valid to conclude that we are in fact in a bubble.

However, we currently believe that the premises are unlikely to hold true.

Instead, we see a high probability that cloud computing growth will be sustained over a long period, driven by demand for both traditional digital transformation and AI (with these two driving each other under a positive flywheel effect).

We make the important distinction between full-service hyperscalers and those new industry entrants who are seeking to compete only on the commoditised rental of bare-metal GPU compute for AI model training.

We believe earnings power of hyperscalers is probably not overstated due to depreciation schedules and, indeed, highlight the point that significant R&D expenses (which very likely provide future benefits but are fully expensed in each period) almost certainly understate earnings power.

Additionally, we expect structural gains in digital advertising due to AI-enabled content recommendation and ad targeting. Even as growth in impressions slow, structural pricing gains will likely stem from ongoing AI advancements.

And, more broadly, we identify the significant non-cloud businesses attached to the hyperscalers as a source of major strategic advantage that derisks new data centre capacity.

Finally, as a result of the above assessments, we believe the market-implied expectations that are embedded into the stocks of select companies (such as the hyperscalers) as undemanding. Therefore, we continue to view them as undervalued investment opportunities.

This is our assessment, even against a backdrop of many individual examples of head-scratching valuations – especially in private markets.

Of course, we acknowledge, there are many possible scenarios that could make us revise these assessments.

If inflation and interest rates were to spike, it would surely curtail demand and make financial conditions much tighter. Today, that does not appear to be an issue – and deflation being exported by China is actually quite helpful.

But things can change. Politics are volatile in today's world. Inequality is high; affordability is low. Many believe major armed conflicts are just around the corner[8].

We remain open-minded.

But today, on balance, after examining the arguments of bubble proponents using first principles, we disagree with the AI bubble hypothesis.

8: (FT) France isn't ready to send its children into battle, November 2025

CONTACT DETAILS

Montaka Global Investments
W: montaka.com
E: office@montaka.com
P: 02 7202 0100

Zarina Kalapesi
Sr. Adviser - Distribution, MFF Group
E: zarina.kalapesi@mffcapital.com.au
P: 0488 136 635