ValuAnalysis

28th November 2025

ValuInsight – The AI CAPEX "Bubble"

- **Bubble? No Bubble? An attractive theme for journalists, but not necessarily for investors.** If "bubble" means more supply than demand, we might be at the opposite of a bubble! Let's go back to the root of this unusual capital cycle: the **Scaling Laws**.
- **In less than a decade (2012-2020), three truly remarkable breakthroughs occurred that created AI as we know it today.** In 2012, Google managed to train a 1bn parameter neural network on *unlabelled* pictures and discovered that the network, *unsupervised*, could randomly recognise a cat. This was the birth of modern Machine Learning. In 2017, Google, again, proved that a "parallel structure" (the tokens do not "queue" but create connections in all sorts of directions) is vastly superior in terms of hardware utilisation and training time, and scales exceptionally well with GPUs and TPUs. Finally, in 2020, OpenAI scientists proved beyond doubt what was already evident from the previous breakthroughs: more compute + more data + more model parameters, aka scaling, improves ROI on a power law. **Modern AI was born.**
- **Because of scaling, it is rational for this CAPEX to happen.** More CAPEX means better models, which leads to more usage, more data, more revenues, and more…CAPEX. This creates a complex strategic and financial problem for the players involved. No one wants to leave the **Scaling Loop** too early, risking marginalisation, or fatal extinction.
- **Not all players face the same existential threat.** Microsoft and Google cannot afford to drop out, and have the means to deploy much more capital. Oracle is effectively out of the frontier game – too much debt, no model. Amazon and Meta do not benefit from the strong public cloud + model combination, but have enough of a defensible moat (retail and "attention", respectively) to play a role. Meanwhile, xAI seems to struggle for cheap access to energy.
- **As the loop accelerates, Nvidia and Broadcom are the main outside beneficiaries**. But when the cycle normalises, their valuation will be in great danger, and those who can use AI efficiently (e.g. Enterprise Software) and those who own data will pick up the spoils.

Pascal Costantini (+44) 203 058 2931

The AI Capex “Bubble”

How many specialised AI newsletters and podcasts do you know? Here, at least six or seven. A cacophony of contradicting opinions: AI is in a CAPEX bubble, AI is not in a CAPEX bubble... Getting to the facts, doing bottom-up models was becoming unavoidable. The concept of an AI CAPEX bubble is vague, sometimes misleading and largely based on share price movements rather than economics. If “CAPEX bubble” means an imbalance between too much supply (CAPEX infrastructure) and too little demand, then we are at the opposite of a bubble... Meanwhile, where does scaling come from? Why is this investment cycle so special? How and when will it end? Formulating a tentative response to these real questions is the immodest objective of this paper.

Scaling Laws- the Ariadne’s Thread of AI since 2012

Artificial Intelligence is a real science, based on experiments and published papers. In retrospect, it is relatively easy to discern a clear logical arc of development between three foundational papers, published in 2012, 2017 and 2020, to form “AI” as we know it today. Very few AI scientists during these critical years were likely able to put their developments into this perspective. Perhaps they were just a handful of extremely talented people flapping their butterfly wings at the right moment, but each of these three pivotal years was a stepping stone, and each has unlocked formidable potential with a common thread: **AI improves with more compute (“scaling”)**. The consequences of this inexorable common thread in AI development are immense for equity investors, and justify that we indulge in a brief history of scaling. You already know all the protagonists, but you probably never realised how closely interrelated they all were...

A Brief History of Scaling

Collectively, the “Cat Paper” (2012), the “Transformer Paper” (2017) and the “Scaling Laws Papers” (2020-2021) all reinforce the same conclusion: AI performance reliably improves with more data, more compute and more model capacity. The initial 2012 breakthrough was about large-scale learning. Google scientists demonstrated that a 1 billion connections network using a cluster of 1,000 machines with 16,000 CPU cores could be trained on 10 million unlabelled and randomly selected YouTube standardised images (one standard 200 x 200 pixels random frame per video) to detect “high level concepts”. We burden the text with numerous-digit figures on purpose, to illustrate the large-scale dimension of the project, even by today’s standards. Enter the main character of this story: the cat gave its nickname to the paper, but was an uninvited and unexpected visitor to this experiment. The neural network was not told what to look for, nor was it presented with videos of cats specifically. It just so happens that a neuron (a.k.a. the “Cat Neuron”) started to “respond strongly to cat faces” when presented with some of the randomly selected 10 million video frames. It had taught itself, unsupervised, to recognise cat features, without ever having been told what a cat was, let alone to look for one. This was the foundation of large-scale unlabelled learning, "modern" Deep Learning, demonstrating the power of massive compute and data, where scale is more important than algorithm.

After the 2012 breakthrough, Deep Learning accelerated across products—from vision (e.g., Google Photos) to language (e.g., early predictive text). In parallel, ad platforms benefited from better models trained on massive datasets using GPUs, contributing to strong ROI from scaling compute and data. The equation was simple and immutable: combining as much compute and data as possible makes better models and enables better ad targeting, yielding incredible ROI on the compute investment. For example, Google's advertising business went from ca. $50bn in 2013 to ca. $80bn in 2016. Yet, it took less than five years to hit the limit of this type of models. These models mostly processed tokens sequentially; picture a tunnel vision with tokens queuing up with only adjacent dependencies. The Transformer Paper, subtitled "Attention Is All You Need" solved this issue by introducing "self-attention". In the sentence:

After a day of canoeing down the river, they stopped for the night. Within an hour, the bank was transformed into a fully functioning camp,

it is clear – to a human brain - what "bank" refers to. Models' self-attention means paying attention to long-range dependencies, letting tokens have an interaction with every other token and linking "bank" to clues like "canoeing" and "river", even many words away, to exclude the meaning of "financial institution". The design from the Transformer Paper enabled fully parallel training, where all directions are searched. This dramatically improved hardware utilisation and shortened training time. Most importantly, this infrastructure design scaled exceptionally well across GPUs / TPUs. The "attention first" / transformer blueprint soon became the template of today's big language models (indeed GPT stands for Generative Pre-trained Transformer).

Three years after the Transformer Paper, OpenAI scientists published Scaling Laws for Neural Language Models (2020). They established beyond doubt that increasing (1) model parameters, (2) dataset size and (3) training compute leads to predictable and substantial improvements in performance, following a power-law relationship. We already knew that scaling improved performance, but the predictability aspect was key for the hyperscalers; predictability enhances ROI by enabling targeted investments, leading to reduced costs and higher efficiency for real-world applications.

The Scaling Laws created an extraordinarily complex problem for those companies involved and their capital providers (which are first and foremost their own shareholders who fund CAPEX by giving up some Free Cash Flow). They overlay a technical problem ("how long will "scaling" work") with a game theory problem. Let's start with the former, which we call the Scaling Loop, the loop no one wants to leave.

The Loop No One Wants to Leave

The Scaling Loop (picture, below) makes the AI capital cycle **reflexive**. As long as it "works", there is no incentive to drop out, and a major economic incentive to stay involved. This is clearly the case for the model labs (OpenAI, Anthropic, xAI) and for the hyperscalers, who depend on an LLM for their own products (cloud, retail, advertising, enterprise software etc...).

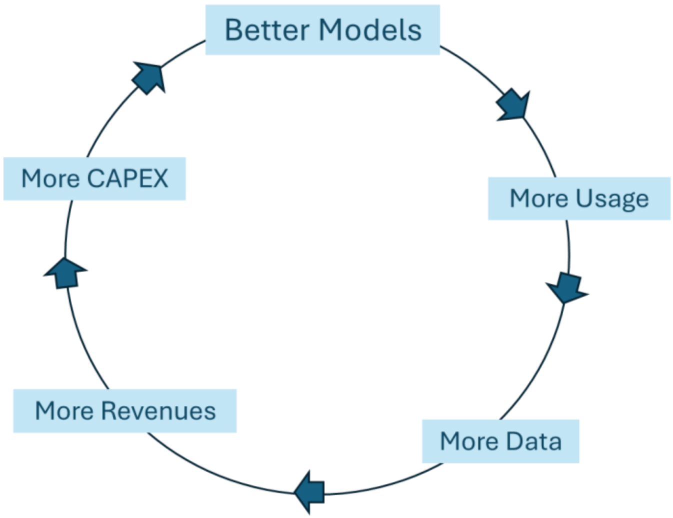

Recent evidence strongly suggests that scaling still works, or, in AI jargon, still "buys real capability".

The next chart plots "capability" versus "capacity". Vertically is a composite index of benchmark results, or how the models are scoring in math / reasoning, Grad-level QA, coding and a few other assessments. Horizontally is on how many runs the models have been trained, in FLOPs of training compute ("Floating Point Operations per second" are a measure of computer performance).

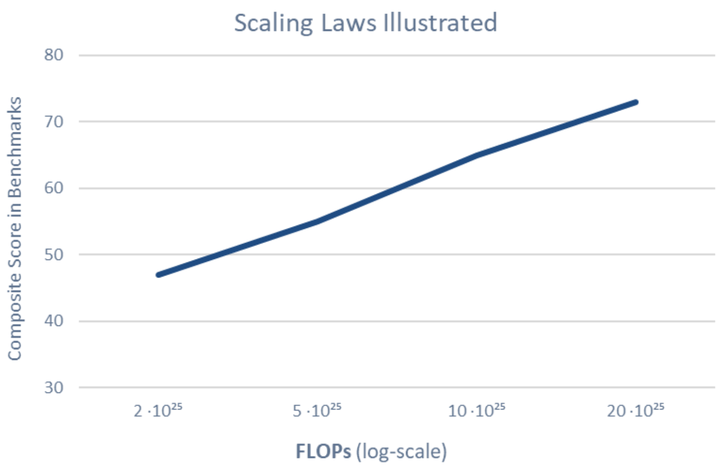

Source: ValuAnalysis ltd

This chart clearly shows that the larger the training compute, the better the results, with no inflection in sight. Scoring well on benchmarks is far from a pure marketing exercise. The scores indicate how the massive energy and capital inputs are actually converting into the intelligence needed to run reliable agents. A high coding score translates directly into products like GitHub Copilot (Microsoft). A high reasoning score allows the AI to handle complex enterprise tasks, for instance legal content. All the major labs (OpenAI, Anthropic, xAI...) keep on releasing improved versions of their models as a result. Google's Gemini 3 is the latest to hit the headlines.

Epoch AI, a research group tracking training compute for frontier models, put the "frontier compute barrier" at 10^{26} FLOPs in 2022. This number of training compute was hypothesised to be close to "AGI-ish" (Artificial General Intelligence) requirements, and was seen as a barrier constrained by costs, GPUs and power. No model had clearly been trained on 10^{26} FLOPs until 2024. By 2025, Grok3, GPT5, Claude4 and Gemini3 are all believed to have been trained at or beyond this threshold. The same research outfit predicts that 6 to 10 such models will have broken the "10^{26} FLOPs barrier" by end 2025 (we are there with OpenAI, Anthropic, Google, xAI, Meta, Microsoft, Amazon and at least four Chinese labs all being able to mount 10^{26} FLOPs runs). They predict 30 models by 2027 and above 200 by 2030, under the continuation of the current trend in investments.

The above matters enormously to investors. There are probably less than 200 people (wild guess) directly involved in these CAPEX decisions world-wide, and these are the type of projections they care about. The 2023-2025 acceleration in training also proves that training compute costs are not yet prohibitive, that training returns are still economically meaningful, and that scaling laws still deliver capability. The 2026-2027 period is likely to see (1) cheaper compute at scale, thanks to Nvidia's GPUs, AMD's, Google's and Amazon's chips, (2) mega datacentre projects: Stargate (recently touted by Satya Nadella himself on the Darwkesh podcast), Colossus (xAI), Oracle's sovereign AI regions and Rainier (Amazon, 1 million Trainiums) and (3) more labs (or at least more models), with at least ten on the 10^{26} FLOPs barrier by 2026. So many labs competing at the frontier, and possibly more to come, amplify the scaling reflexivity; it increases competitive pressure, model release cadence, CAPEX escalation and frontier experimentation.

And this could just be the beginning. Nvidia's Jensen Huang has built a narrative around Scaling Laws, breaking down training into "pre" and "post" training, and, crucially, adding inference (the question/answer part, interacting with users). This may be the new, future twist to this saga: more compute per question gives better reasoning and better answers. We will have to see if some scientific evidence supports Nvidia's marketing. But technology is not the only aspect of what makes the AI CAPEX cycle highly unusual.

Psychology, Game Theory and Revealed Preferences

There is a multitude of game theories that we could invoke for this capital race, depending on the lens we wish to use: Game of Chicken (the first to stop scaling loses everything), War of Attrition (victory goes to whoever can absorb the most costs), Arms Race (everyone invests because everyone else invests). On page 7, "Until the End", we review how various exogenous shocks could break the loop. Yet, willingly dropping out of the Scaling Loop early is fatal. The first to blink will lose capability, cost advantage and customers. Its cloud will become more expensive per unit of inference, its eco system will shrink, its talents will leave. All this at the same time as competitors get better models, cheaper inference, more enterprise agent adoption, a larger developer ecosystem… This becomes irreversible. Mark Zuckerberg, in a recent Access podcast, called "misspending a couple hundred billion" (NB: that's more than the GDP of 70% of the countries in the world) on AI infrastructure "very unfortunate", but promptly added that "the risk was higher on the other side", meaning the side of underinvesting. Applying

Paul Samuelson's revealed preferences theory - the analysis of people's actual choices when faced with trade-offs - to Mark Zuckerberg's statement leaves investors under no doubt: barring an external shock, CAPEX will stop after scaling stops working. Not a second earlier.

In detail, not all the players are subject to the same constraints. The general AI labs are fighting an existential battle, and the one stopping too early will disappear. If we judge by comparable instances, it seems that only a handful of general models, 2 to 4, are likely to survive. This is what happened in operating systems for computers (Windows and Mac), and again for mobile phones, this time with Google (Android) vs. Apple. In cloud, there is a small oligopoly of three, complemented by niche players.

Outside of the AI labs, the other players are under more differentiated threats, and we see two distinct groups. In the first one, Microsoft and Alphabet. These are now mortal enemies, because the hyperscaler that hosts the best model becomes the centre of gravity. Indeed, the cloud that becomes home to the leading model automatically gets enterprise migration, developer platform share, ecosystem gravity (third party apps) and enhanced cost and scalability optimisation. AI is now a core differentiator for Microsoft, for Azure (AI native cloud), but also for Windows, M365, Copilot, GitHub, security, Dynamics, LinkedIn and probably Activision. The AI stack is now Microsoft's growth engine, and stopping early would mean a huge strategic climbdown, and would stop in its tracks the attempt to make Copilot the new OS. They probably have to ride the scaling curve to its economic limit, then pivot to efficiency. Can they cope with being number 2? Possibly, as long as the gap to number 1 is narrow to protect their cloud market share; if the gap is small, Azure could still win distribution, enterprise relationships and integration.

The same applies to Alphabet with respect to its cloud business. On the models, Google DeepMind was perceived to lag OpenAI, but the new release of Gemini 3 appears to have somewhat closed the gap. Search and You Tube must be AI native to defend monetisation and user experience. Google Cloud needs competitive models to compete with Azure and AWS. Surrendering the frontier leadership risks the core cash machine (search ads). Could they settle for number 2? Yes, as long as they remain number 1 in search and ads. Finally, Alphabet has the added benefit of TPUs, which is substantial.

We see Amazon, Meta, Oracle and xAI belonging to a different, and somewhat less intense group, as all have some weaknesses that the first two have not, but also live under a less existential threat. Of these four, Amazon and Oracle don't really have a state-of-the-art model, and Meta and xAI don't really have a (public) cloud. Losing the race would not have the same catastrophic consequences as for Microsoft or Google, in our view. Meta in particular seems to be willing to fight a war that it may not be required to fight. Its ad empire is based on a closed-end capture of attention that cannot migrate anywhere. It is often wrongly compared to Google, whose ad empire depends largely on its gatekeeping function with search, which could migrate to AI agents if it loses the race. Meta could rent a better model than its own Llama and still function. But underestimate

Meta at your peril. Despite Llama being positioned as an open-source foundation and not actively seeking monetisation, some early indications suggest that Meta is experimenting with revenue sharing agreements with hosting clouds.

“Until the End”

There is no finer way to put it: “they” will likely fight it until the end. Our premise is that AI is the biggest revolution in computer technology, and we rule out the idea that AI “won’t work”, or will not find enough applications. This does not mean that we will reach AGI any time soon. If we do, then that will be the natural ending. If we don’t, within say, the next five years, the end of the AI CAPEX cycle will likely come from hitting one or two roadblocks: energy and capital availability, with a possible third, regulation.

Energy access is becoming a problem. Driven by their AI component which is growing at ca. 30% p.a., datacentres’ power demand is growing at 12 to 15% per annum, and will have doubled in early 2030s (source: IEA). Hyperscalers are in a race to secure long-term energy supply, and only two companies have a substantial portfolio of PPA (Power Purchase Agreements), or long-term (decades) supply agreements of (often clean) energy: Microsoft and Google. Microsoft has contracted above 34 gigawatts of carbon-free electricity and is at the forefront of securing PPA deals, seen as a strategic resource. Google is equally active, with PPAs totalling above 22 gigawatts of clean energy. Amazon is a distant third, and no one else is able to match remotely this access to energy. If we hit an energy wall, these three – and their associated models – will be hit last.

But the bigger, all-encompassing issue is capital access. Despite the size of their balance sheets, the hyperscalers had to move quite a bit down the financing stack in the past 12 months, from equity (e.g. retained earnings) to debt issuance, after having pulverised free cash flow (but, incidentally, retaining outsized Stock-Based Compensation). Alphabet just did $25bn in USD and EUR, with maturities out to 2075, Meta raised up to $30bn in six tranches (5 to 40 years), Amazon just returned to the US bond market for the first time in three years raising $12 to 15bn, and Oracle did $15 to 18bn in bonds with similar terms, 7 tranches out to 40 years. Overall, nearly $100bn was raised in debt, with the notable exception of Microsoft. Investment-grade bond investors seem to be very happy to fund AI CAPEX in this way; the yield range is attractive (5 to 6% for 10 to 30 years paper), and deals are heavily oversubscribed (Meta reportedly saw $100bn of demand, and Amazon $80bn). In this respect, there are no signs of investor fatigue.

Yet this represents the cleanest source of capital, but an entire web of circular deals haunts the lower habitats of the financing stack. The vendor-financing deals that we have seen emerging in the past few months are clearly reminiscent of the worst aspects of the 2000/2003 bubble, and it is no surprise that investors are focusing on the next one, the so-called “AI bubble”. Below, we verify how the macro projections stack up “bottom up”.

A Bottom-Up Model of AI Funding

A word of warning: these are, inevitably, rough numbers that will require adjusting over time; but they are directionally right. The framework is the next five years (2026-2030). Our model starts with the latest figures of the key players, the hyperscalers plus Meta. On the following table, we assume that SBC costs and an additional 10% of Cash from Operations (CFO) are untouchable. We estimate an amount of legacy, pre-AI baseline CAPEX based on pre 2023 data and calculate what is in theory available for AI spending ("AI available"), to be compared with actual AI CAPEX. The total AI spent is $172.5bn, a figure consistent with other calculations.

Three points come out of this table:

- The capital wall has not been hit. In aggregate, there is a self-funded headroom of ca. $107bn as of today ("AI available" minus "AI CAPEX", negative funding of Amazon and Oracle ignored).
- Two companies (Amazon and Oracle) have no headroom left
- Two companies (Alphabet and Microsoft) are in a league of their own, and have massive headroom left relative to their group.

We think that the trajectory of the current (2025) $172bn CAPEX spend will be affected by the following during the next five years:

- The future growth of the hyperscalers' CFO (likely low double digit)
- The Scaling Laws, of course, with larger training runs, more inference, partly compensated by efficiency gains
- Falling unit costs for compute, which may reduce costs per TeraFLOPs substantially over five years
- Physical constraints (land and construction availability and power)

Table 1: The Five Key Players – Last Twelve Months Figures, in $ bn

	Alphabet	Microsoft	Amazon	Meta	Oracle	TOTAL
Owner CFO (post SBC)	128	135	111	89	17	
Retained (10%)	-13	-13	-11	-9	-1.5	
Legacy (pre-AI) CAPEX	-38	-29	-76	-38	-3.5	
"AI available"	77	93	24	42	12	
AI CAPEX	-40	-40	-44	-25	-23.5	-172.5
Self-fund. headroom	37	53	-	17	-	107

Source: ValuAnalysis Research

We estimate that growth in CAPEX over the next five years ought to be around 15% per annum, which means that AI CAPEX will double, and reach a cumulative spending of 1.34tn for the hyperscalers. To this figure we need to add the second-tier capital requirements from the “neo clouds” (CoreWeave, Crusoe, Lambda, Voltage Park, the list goes on…) and the labs. The numbers are difficult to pin down but, in any case, significantly smaller; we estimate that these capital requirements may be around $30bn per annum, with limited opportunity to grow due to constraints linked to their status. For instance, private credit will not fund their expansion without signed contracts or commitments. Overall, we think that their AI bill might amount to $200bn over five years.

Finally, power builds will also be a substantial CAPEX bill that will need to be financed, traditionally or creatively (e.g. leasing). In this respect, the energy problem becomes a subset of the capital problem. With an extra 550 TWh/year datacentre demand implied by the IEA, the world might need 100 to 150 gigawatt of additional capacity, depending on the energy source. Since one GW costs about $2.5bn to build, the final bill over 5 years ought to be around $300 to $400bn, or $60 to $80bn per annum. **Overall, we estimate that a reasonable and realistic total cumulative AI bill in 2026-2030 might hit the 2tn mark** (1.34tn + 0.2tn + 0.35tn = $1.9tn). This amounts to a full buildout from the main hyperscalers, including power.

So much for headline-grabbing deal announcements; there are today **only four capital providers** to the AI eco-system: Microsoft, Google, Amazon and private credit distributors. As was clear from Table 1, the bulk of hyperscalers’ CAPEX can be self-funded, including the ca. $107bn of “self-funded headroom”, which we can safely assume will be engulfed into the overall future CAPEX. Hyperscalers’ CAPEX is about half of the funding requirement, straight out of CFO. The other half will need to be a call on the world’s saving pool, and will likely breakdown into a “big three” (Microsoft, Google and Amazon) bond issuance and private credit (including sovereigns) funding.

Microsoft and Google, and, to some extent, Amazon, have got ample room to leverage their balance sheet. Meta does not participate directly in the ecosystem (it does not have a public

Table 2: Debt Level Compatible with AA Rating

	Alphabet	Microsoft	Amazon	Total
EBITDA	145	166	140	
Net Debt (cash)	(54)	18	66	
Net Debt to EBITDA	-0.4x	0.1x	0.5x	
Additional Debt	145	150	70	365

Source : ValuAnalysis Research

cloud, uses all its infrastructure internally and does not fund anybody else), and Oracle already has a debt to EBITDA ratio of 4, which is BBB and limits additional leverage. Most rating agencies will give an AA rating with a debt to EBITDA ratio of 1x or less. On that basis, we estimate that Alphabet and Microsoft, in particular, but also Amazon, could raise approximately an additional $365bn, probably a little bit more (see Table 2).

Private credit is the fourth source of AI funding, and would need to step up to fund the rest, or about $0.6tn. Blue Owl, Blackstone, KKR, Brookfield, to name a few, are already quite active. Sovereigns and public programmes (Gulf funds, national AI/DC plans) have a plausible role to play in this order of magnitude.

Final Thoughts

The purpose of this paper is not to defend a pro or anti AI position. Rather, it attempts to analyse, from as many angles as possible, this enormously complex investment theme. This is where we stand:

- **Scaling Laws are real**. There is no evidence that this loop is about to stop, as the recent Gemini 3 release seems to indicate. Yet, the day will come. Some researchers (in MIT 2025 AI Scaling study) warn that data scarcity could create an inflection point by 2028. The day the loop breaks will be a formidable event.
- **Funding is not about to stop**. There is no current "capital wall", as Table 1 makes it clear. The hyperscalers still have plenty of ammunition to drive their ambitions further. Expect aggressive leveraging from the big three, but don't anticipate an Oracle-like deterioration of their prudential ratios.
- **The pace is accelerating**. The battle between Microsoft and Alphabet will be epic. Behind, xAI is privately funded and sits at a disadvantage on the energy front. Its Colossus project does not have access to enough power from the grid, and runs on mobile gas turbines with no deep PPA book. Meta pretends to be up for a fight it has no real business fighting unless it plans to pivot to full monetisation of Llama. Oracle has already lost any hope of competing at the frontier. Only three contenders mean business, and only two, Microsoft and Google, are really up for it. Amazon has the benefit of the biggest cloud business, but its heart is on retail, not LLMs.
- **The situation will evolve further**. We have limited information about the frontier models that China may own. Meta's Llama monetisation potential is intriguing and could be a belated twist. Whilst we don't expect either Microsoft or Alphabet to make a fatal strategic mistake, plenty will also happen around them. It is in their interest to talk valuations down. If/when some neo cloud fails, they will be interested buyers. If funding gets tough among the labs, they will be willing to help, on their terms.
- **Overall, this is a dangerous area**. The status quo depends on the goodwill of investors to keep their cool, and maintain a very low risk premium on these players. This will become even more critical as the big three issue debt, which will add a dependency on the level of interest rates.

- **Two adjacent areas offer an attractive risk / return profile**. As long as Scaling Laws work and are funded, the chips and electricity equipment suppliers will likely do well. But they are a leveraged play on the hyperscaler trade. When

  Scaling Laws reach their point of inflection, Microsoft might not be in such a bad position, controlling the best AI cloud and one of the frontier models, whilst reducing precipitously its investment plans. Now think about Nvidia in this context…

- **Those who can use AI will do disproportionately well** as the AI investment cycle normalises. The market stubbornly ignores this investment thesis, and lets enterprise software companies, for instance, trade on sub-par FCF multiples. An ever-increasing list of stocks get derated, sometimes in spectacular fashion: Wolters Kluwer is down 50% this year, on suspicion of AI threats for its clinical solution software. We believe that the best risk-adjusted expected returns are among these, including some of the data owners.

DISCLAIMER

This document is provided by ValuAnalysis Limited, which is authorised and regulated by the Financial Conduct Authority (firm reference number 710908). This document is only permitted for individuals or firms who would fall within the definition of a professional client as defined by the Financial Conduct Authority's rules.

The information in this document is for informational purposes only and does not provide personal recommendations based on your individual circumstances. By making this information available to you, ValuAnalysis is not advising you or making any recommendation. Investments carry risk, including the risk that you will not recover the sum that you invested.

The views expressed in this document are as of the published date and based on information available at the time. ValuAnalysis does not assume any duty to update any of the information contained in this document.

The information herein is believed to be reliable and has been obtained from sources believed to be reliable, but we make no representation or warranty, express or implied, with respect to the correctness, accuracy, or completeness of such information.

The information contained in this document is strictly confidential and may not be reproduced or redistributed in whole or in part, nor may the contents of the document be disclosed to any other person without the prior consent of ValuAnalysis. By viewing this document, you confirm that you have read and accepted this disclaimer.

CONFLICTS OF INTEREST

To mitigate the possibility of conflicts of interest, ValuAnalysis' employees are subject to internal organisational and administrative arrangements in relation to the management of inside information, handling of unpublished research material, gifts and hospitality, external business interests, remuneration and personal transactions. These internal organisational and administrative arrangements have been designed in accordance with applicable legislation and relevant industry standards. These internal organisational and administrative arrangements are considered appropriate and proportion in light of the nature, scale and complexity of ValuAnalysis' business.

As at the time of writing, ValuAnalysis does not perform services for any issuer mentioned in this report. Notwithstanding, ValuAnalysis may, to the extent permitted by law, perform services for, solicit business from, or otherwise be interested in the investments, directly or indirectly, of any issuer mentioned in this report.

ValuAnalysis prohibits its analysts, professionals reporting to analysts and members of their households from making personal transactions in any issuer in the analyst's area of coverage for a period of 10 days before and after the publication of research pertaining to the issuer. ValuAnalysis prohibits its analysts, persons reporting to analysts or members of their households from serving as an officer, director, advisory board member or employee of any company in the analyst's area of coverage.

ValuAnalysis has no agreements with issuers with respect to dissemination of recommendations. Analysts do not, nor will they, receive direct or indirect compensation in exchange for expressing any of the views or the specific recommendation contained in this report. Analysts are paid in part based on the overall profitability of ValuAnalysis.

ValuAnalysis
The Global Intrinsic Value Specialist

For more information:

www.valuanalysis.com

Pascal Costantini
Founding Partner
The Clubhouse, 8 St. James's Square
London SW1Y 4JU
(+44) 20 3058 2931
pascal.costantini@valuanalysis.com

Joakim Darras
Founding Partner
The Clubhouse, 8 St. James's Square
London SW1Y 4JU
(+44) 20 3058 2933
joakim.darras@valuanalysis.com

Janet Lear
Partner
The Clubhouse, 8 St. James's Square
London SW1Y 4JU
(+44) 203 058 2935
janet.lear@valuanalysis.com