

Moat or Mirage? The Arithmetic of AI Lab Valuations

July 2026

Artificial intelligence (AI) is powerful and in high demand, and AI laboratories have built extraordinary products. Indeed, AI is changing the world. But the investment case for AI hinges on different considerations.

Are the future AI profits now being capitalized in private markets protected by durable moats? Or are they the result of temporary competitive advantages and vulnerable to challengers, falling costs, supplier bargaining power, and shifting regulation in the months and years ahead?

We've proposed that much of the AI capital expenditures among such hyperscalers as Alphabet, Amazon, Meta, and Microsoft is better understood as capital churn rather than capital deepening. Gross capex is not capital accumulation; the difference is depreciation. Rapid obsolescence means that apparent additions to capital stock are more like operational expense flows. To AI hyperscalers, the strategic price of underinvestment may exceed the financial cost of overinvestment. As such, their inflated capex may be entirely rational even if it is unprofitable.

If the AI infrastructure buildout among hyperscalers does not capture AI's economic surplus, what will? Private valuations imply that AI labs could be the emerging sector's big winners. Yet the evidence for that potential outcome is thin. To assess the possibility, we apply the same capital stock-versus-expense flow framework we previously used for hyperscalers to AI labs.

"If the AI infrastructure buildout among hyperscalers does not capture AI's economic surplus, what will? Private valuations imply that AI labs could be the emerging sector's big winners."

What We Can and Cannot Observe

Listed companies lead AI's hardware infrastructure and cloud hyperscaler layers. Their margins, capital spending, and segment economics are publicly reported and comparable. ASML specializes in the advanced lithography machines required to print cutting-edge semiconductors and has a secure position in the AI supply chain's upper reaches.

Downstream, NVIDIA designs the chips that TSMC then manufactures. These two companies have an interdependent rather than a competitive relationship. NVIDIA



AUTHORS



Chris Brightman, CFA*
Partner, Senior Advisor

*Corresponding author

Key Points

- Unlike NVIDIA, TSMC, and other artificial intelligence (AI) suppliers or AI hyperscalers like Amazon and Alphabet, AI labs are not audited public companies, but their economics can be gleaned from company statements, press reports, and analyst estimates.
- The leading AI laboratories are worth 20 to 35 times sales, according to private market valuations. Assuming generous margins on current revenue, those valuations imply 60 to more than 100 times earnings and will require years of extraordinary revenue growth and wide defensible moats to maintain.
- AI labs' assumed margins may be contested from three directions both in and outside the AI ecosystem: competitors from below, suppliers from above, and externally from government regulation.
- Given the lack of comparable data and the novelty of the space, rather than focusing on whether AI lab valuation multiples are cheap, realistic, or in bubble territory, investors should consider whether they are buying a durable stock of profits or a transient flow of competitive advantage.

and TSMC have, respectively, reported gross margins near 75% and above 60%.

Among hyperscalers, such dedicated cloud segments as Amazon Web Services (AWS) and Google Cloud have operating margins in the 30% range. Yet their consolidated parent companies – Amazon and Alphabet – have disclosed surging capital expenditures along with the substantial future depreciation burden that entails.

The labs occupy the middle of the AI investment stream. Unlike the hardware developers and hyperscalers, they are privately owned. While they do not file audited financial statements, they do issue a mix of private round valuations and company statements and are covered in press reports and analyst estimates. A key assumption for valuing these labs is the width and robustness of their future profit margins. Today's sales multiples are a bet that the labs have deep moats that will protect their sizeable profit margins for many years to come.

The Case for Durability

Frontier model quality has, so far, retained pricing power. Open-weight models still meaningfully trail their frontier counterparts on the most demanding tasks. Thus, cheaper models exert limited downward pressure on frontier pricing. That's the bull brief for durability. How wide that gap remains and for how long are among the most contested aspects of the durability question.

The strongest case for AI lab durability does not require any single model to take the lead and remain in the top spot forever. All it needs is for the leading labs to be integrated into the enterprise workflows, developer tools, security reviews, and compliance processes. When the model becomes part of knowledge work's broader operating system, customers are no longer buying tokens, they are buying reliability, governance, integration, and trust.

Businesses that invest in new workflows constructed around a particular AI model do not re-platform cheaply. Distribution compounds the competitive advantage: The leading labs sell through the largest cloud channels and enterprise software relationships. Where demand outstrips supply, frontier providers can price to the value customers place on the output rather than the marginal cost of producing it. Each is a legitimate source of durability. If they hold, today's multiples may turn out to be fair or even cheap.

The Case for Obsolescence

AI data center hardware has a steep depreciation rate. Frontier models depreciate too, perhaps faster than the chips they run on. This is economic rather than accounting depreciation: The model still works, but its scarcity value declines.

An AI lab's present competitive advantage depends on its current model, inference cost, product wrapper, and distribution. Each component can be superseded by the next generation. Flagship models routinely reach economic obsolescence within 12 months, often when the same lab that launched it releases an updated version. Published prices for frontier models have also fallen even as their capability has risen. The research and development of new models every cycle, to maintain competitive advantage, is an operating cost flow, not an addition to capital stock.

Pricing pressure also comes from below. Open-weight models have improved and some are approaching proprietary frontier performance at materially lower cost per token. For workloads that customers can self-host or route to a comparable open model, that competitive alternative can cap the pricing power of the premier labs.

The counterweight to high switching costs is expanding software interchangeability. As corporate customers learn to manage multiple AI providers, the cost of changing labs will drop. Rather than being locked into a single provider, enterprises are adopting management tools that let them easily swap models behind the scenes, automatically routing tasks to the provider with the best price and performance. When a business owns the user-facing application layer, the underlying AI model becomes a fungible commodity.

The Margin Suppliers Could Reclaim

One consideration does not depend on models commoditizing at all. Those suppliers with observable margins hold genuinely scarce, irreplaceable resources and prove it with the high gross margins they command. Thanks to its lithography equipment advantage, ASML sits at the top of the AI supply chain, controlling the bottleneck of advanced semiconductor manufacturing. Further down, TSMC dominates the advanced manufacturing infrastructure and business processes that are required to physically build the processing chips; NVIDIA commands the premier frontier AI accelerators and an entrenched software ecosystem; and Broadcom specializes in the networking silicon that links these massive clusters together.

The supplier's margins are the lab's costs: NVIDIA's near-75% reported gross margin is not a precedent for lab profitability but a component of lab cost of revenue. Since the lab's economics depend on input prices that the suppliers could increase, the lab's margin is not owned outright but relies, at least in part, on a supplier's discretion. This bargaining power goes both ways: The largest labs deliver enormous forward demand, financing credibility, and custom-silicon alternatives to their compute partners. Suppliers may choose not to exercise their pricing power for competitive, strategic, or relationship reasons, but the lab's margin is a function not only of the quality of its model and its customers' willingness to pay but also the price of the compute required to serve that demand.

Regulation Cuts Both Ways

The resilience of AI lab margins has both a policy and a competitive dimension. Governments can narrow or widen the moat by fiat, for example, by restricting the use of open-weight alternatives or intervening in the deployment of the frontier models themselves. The proprietary control that protects a lab's pricing power is also what makes it governable. A multiple that capitalizes decades of profits incorporates expectations about competitors, customers, and suppliers as well as policy regimes that have yet to fully take shape.

*“The resilience of AI lab margins has both a policy and a competitive dimension.
Governments can narrow or widen the moat by fiat.”*

What Valuations Assume

The point here is not to estimate fair prices for AI labs but to make the implicit assumptions explicit. Because the inputs are private, unaudited, and uncertain, we cannot make concrete forecasts. We can only present the elementary arithmetic for illustrative purposes.

In late May 2026, the leading AI labs traded at roughly 20 to 35 times sales based on company-stated or press-reported run-rate revenue rather than official, audited GAAP figures, with Anthropic at the low end and OpenAI near the top. Even these bases are not directly comparable since Anthropic and OpenAI define and net their reported metrics differently. A multiple of sales becomes a multiple of earnings only after a net margin is assumed. The first step is simply to translate today's sales multiples into implied earnings multiples under different margin assumptions.

Starting with the calculations at current revenue, even under the extremely generous assumption that one-third of revenue ultimately becomes after-tax profit, a company valued at 20 times sales trades at roughly 61 times earnings. One valued at 35 times sales trades at 106 times earnings. **Exhibit 1** shows the sensitivity: At more conventional margin assumptions, the implied multiples soar.

Exhibit 1. AI Labs: What Today's Sales Multiples ImPLY

Valuation	Implied Earnings Multiple at Assumed Net Margin of:	
	20%	33%
20x sales	100x	61x
35x sales	175x	106x

Notes: Implied P/E equals price-to-sales divided by assumed net margin. A 33% net margin would place an AI lab among the most profitable large companies in history.

Sources: Sales multiples reflect company-stated or press-reported valuations and annualized run-rate revenue as of late May 2026. Supplier and hyperscaler data are derived from the most recent quarterly results and filings of NVIDIA, TSMC, Microsoft, Alphabet, Amazon, Broadcom, and ASML. AI lab revenue, margin, and valuation figures are company-stated, press-reported, or estimated rather than audited, and are from company announcements and private-market sources, such as Reuters, CNBC, The Information, and Sacra, as of late May 2026.



From these levels, growth is the only path to plausible mature earnings multiples. Exceptional growth companies have traded at 100 times earnings and richly rewarded their investors because the growth was achieved and the margin survived competition. AI lab prices today capitalize many years of extraordinary revenue growth while defending wide profit margins despite uncertain moats.

What Price Is Fair?

So, are AI labs in a bubble? We believe the inputs are too soft and the territory too new to render such a verdict – not that we aren't tempted. But a reasonable bull could say that the labs will grow into their sales multiples, while a reasonable bear could counter that such rapid growth rarely survives contact with competition, customer bargaining power, or technological depreciation.

“[A]re AI labs in a bubble? We believe the inputs are too soft and the territory too new to render such a verdict – not that we aren't tempted.”

Overall, current AI lab valuations do make implicit assumptions about future profit margins and their durability that the available data can neither confirm nor refute.

Implications for Investors

When presented with an AI lab's assumed margin, an analyst's run-rate estimate, or a pre-IPO valuation, investors should ask two questions before making an investment decision. First, is the number audited, publicly reported, company-stated, press-reported, or estimated? The answer determines how much weight to give it. Second, are they paying for an annuity – a persistent, robust, and moat-protected margin? Or for a transient competitive advantage that model generation may override, customers may navigate around, suppliers may reclaim, or public policy may revoke?

AI is transforming the economy and AI labs could be among the top beneficiaries. But economic transformation is not the same as long-term profit generation. There is no guarantee that the margins today's AI lab valuations imply will be earned over the long term or that they won't be. While the surplus at the supplier and hyperscaler layers is reported and measurable, the surplus at the lab layer is speculative.

Whether that surplus will function more like a capital stock or a cost flow is a question the price asks and answers implicitly. Investors, by contrast, must be explicit about how they address the question and the expectations that underlie their answer.

The material contained in this document is for informational purposes only. It is not intended as an offer or a solicitation for the purchase and/or sale of any security, derivative, commodity, or financial instrument, nor is it advice or a recommendation to enter into any transaction. Research results relate only to a hypothetical model of past performance (i.e., a simulation) and not to actual results or historical data of any asset management product. Hypothetical investor accounts depicted are not representative of actual client accounts. No allowance has been made for trading costs or management fees, which would reduce investment performance. Actual investment results will differ. Simulated data may have under- or over- compensated for the impact, if any, of certain market factors. Simulated returns may not reflect the impact that material economic and market factors might have had on the advisor's decision-making if the advisor were actually managing clients' money. Simulated data is subject to the fact that it is designed with the benefit of hindsight. Simulated returns carry the risk that actual performance is not as depicted due to inaccurate predictive modeling. Simulated returns cannot predict how an investment strategy will perform in the future. Simulated returns should not be considered indicative of the skill of the advisor. Investors may experience loss of all or some of their investment. Index returns represent back tested performance based on rules used in the creation of the index, are not a guarantee of future performance, and are not indicative of any specific investment. Indexes are not managed investment products and cannot be invested in directly. This material is based on information that is considered to be reliable, but Research Affiliates, LLC ("RA") and its related entities (collectively "Research Affiliates") make this information available on an "as is" basis without a duty to update, make warranties, express or implied, regarding the accuracy of the information contained herein. Research Affiliates is not responsible for any errors or omissions or for results obtained from the use of this information.

Nothing contained in this material is intended to constitute legal, tax, securities, financial or investment advice, nor an opinion regarding the appropriateness of any investment. The information contained in this material should not be acted upon without obtaining advice from a registered professional. RA is an investment adviser registered under the Investment Advisers Act of 1940 with the U.S. Securities and Exchange Commission (SEC). Our registration as an investment adviser does not imply a certain level of skill or training. RA is not a broker-dealer and does not effect transactions in securities.

Investors should be aware of the risks associated with data sources and quantitative processes used to create the content contained herein or the investment management process. Errors may exist in data acquired from third party vendors, the construction or coding of indices or model portfolios, and the construction of the spreadsheets, results or information provided. Research Affiliates takes reasonable steps to eliminate or mitigate errors and to identify data and process errors, so as to minimize the potential impact of such errors; however, Research Affiliates cannot guarantee that such errors will not occur. Use of this material is conditioned upon, and evidence of, the user's full release of Research Affiliates from any liability or responsibility for any damages that may result from any errors herein.

The trademarks Fundamental Index™, RAFI™, Research Affiliates Equity™, RAE™, and the Research Affiliates™ trademark and corporate name and all related logos are the exclusive intellectual property of RA and in some cases are registered trademarks in the U.S. and other countries. Various features of the Fundamental Index methodology, including an accounting data-based non-capitalization data processing system and method for creating and weighting an index of securities, are protected by various patents of RA. (See applicable US Patents, Patent Publications and protected trademarks located at <https://www.researchaffiliates.com/legal/disclosures#patent-trademarks-and-copyrights>, which are fully incorporated herein.) Any use of these trademarks, logos, or patented methodologies without the prior written permission of RA is expressly prohibited. RA reserves the right to take any and all necessary action to preserve all of its rights, title, and interest in and to these marks and patents.

The views and opinions expressed are those of the author and not necessarily those of RA. The opinions are subject to change without notice.

©2026 Research Affiliates, LLC. All rights reserved. Duplication or dissemination prohibited without prior written permission.

AMERICAS

Research Affiliates, LLC

660 Newport Center Drive, Suite 300
Newport Beach, California 92660
USA

+1.949.325.8700
info@researchaffiliates.com

EUROPE

Research Affiliates Global Advisors (Europe) Ltd

78-79 Pall Mall
London SW1Y 5ES
United Kingdom

+44 (0) 20 3929 9882
uk@researchaffiliates.com