

Artificial Intelligence Investing After the First Wave

Published June 26, 2026

A year ago, in our three-part series Navigating the AI Revolution, the central question for investors was where artificial intelligence's (AI's) disruptive potential would translate into meaningful economic and market change.¹ That question now has a clearer answer. AI is already reshaping parts of the economy and market as capabilities improve rapidly, enterprise adoption broadens, and revenue growth becomes more visible across the ecosystem. Disruption is no longer a distant possibility. It is beginning to show up in software and labor-intensive functions such as customer service.

The investment question has changed with it. The first phase of AI investing was led by hyperscalers, advanced semiconductors, and large language model developers. The current phase has been driven by infrastructure buildout, and markets have already recognized much of that trade. Over the next few years, we expect the most attractive opportunities to center on more persistent bottlenecks, particularly around power, and on companies that control workflow, own the customer relationship, bring domain expertise, and turn AI output into business action. That includes the software infrastructure and applications that will shape how AI is deployed and managed. As companies integrate AI into workflows, adopters across industries should benefit. We expect much of the next phase of value creation to come from emerging AI-native companies and disruptive business models, though at this early stage, even today's disruptors may be displaced.

In this edition of VantagePoint, we focus on three questions:

- Which bottlenecks are durable?
- Can rising revenues justify the capital required to sustain leadership?
- Where can value persist as AI becomes cheaper, more capable, and more widely available?

The broader consequences for labor and society may prove profound, but they remain harder to observe clearly and are progressing slower than the technology itself. Regulation and sovereignty risk considerations clearly matter, as does society's willingness to absorb the pace of change and its consequences. Those forces will shape AI's development, but they are not the focus of this paper. We focus instead on where the evidence is strongest today and where the investment implications are becoming harder to ignore.



Celia Dallas
Chief Investment
Strategist



Katharine Campbell
Partner, Endowment
& Foundation Practice



**Theresa Sorrentino
Hajer**
Head of US Venture
Capital Research

¹ Please see "Navigating the AI Revolution: A Three-Part Series," Cambridge Associates LLC, July 10, 2025.

Tech is still early, moving fast

AI capabilities continue to improve at breakneck speed. The leading labs remain in a tight race, and each major release raises the bar while increasing disruption risk for incumbents and start-ups alike. Leadership among US models continues to oscillate. Chinese models have narrowed the gap on several technical measures despite US export controls on advanced chips. The open-model ecosystem, much of it coming from China, has also become a credible lower-cost option for use cases outside of mission-critical workflows that still demand premium US tools.

Frontier intelligence is becoming more capable, more available, and less exclusive. That should expand adoption, but it also makes raw model access a less reliable source of durable advantage and shifts value toward the assets and capabilities that make intelligence useful, governable, and hard to replace.

AI is shifting from a tool that generates responses to a tool that performs useful work. Recent model releases have improved reasoning, reliability, memory, and the ability to work across different types of data. Anthropic's AI model, Claude, helped accelerate the move from coding assistant to autonomous coding agent through Claude Code, then extended that logic into broader knowledge work. While this is a major disruptive theme and a central focus of this paper, it is only one part of a broader transformation. Specialist models in mathematics, biology, and other fields are proliferating, while researchers continue to experiment with non-transformer architectures such as world models designed for physical AI. Smaller models are improving as well, pushing more inference to run on local systems instead of keeping them entirely in centralized cloud environments.


These advances are beginning to make disruption more visible. In software, faster model improvement is compressing product cycles, narrowing functional differentiation, and pressuring application-level moats that once appeared durable. In labor-intensive workflows such as customer service, the effects are already showing up in shorter handle times, lower staffing needs, and greater pressure to automate routine work. Much of the broader disruption still lies ahead, but the direction is clearer.

Enterprise AI adoption is rising, but scaled deployment remains limited

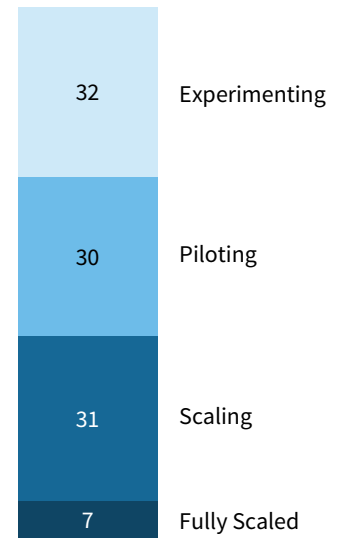
The pace of enterprise adoption is becoming increasingly apparent, even if scaled deployment remains limited and difficult to track in real time. McKinsey survey data captures the key point for investors: adoption is rising, but much of it still reflects experimentation, piloting, and limited deployment rather than full integration across core workflows.

Enterprise AI adoption is rising, but scaled deployment remains limited

2017–25 • Usage of AI survey responses (%)



2025 • Phase of AI use (%)



Source: McKinsey & Company.

Note: AI adoption definitions evolved over time, from scaled or core-business use to regular use in at least one function.

Coding remains the clearest use case, though productivity gains are still early and uneven. Deployment inside real organizations will be challenging, requiring integration, oversight, and governance as much as good model performance. Usage will evolve as service models change. For example, firms are still enjoying subsidized model pricing, internal controls are weak, and many firms are still learning to use these tools effectively while keeping token costs under control.

As AI advances rapidly, the need for stronger corporate governance is becoming more urgent, even if progress is moving at a more human pace. Some early adopters are pulling ahead in part because they addressed governance, access controls, and oversight sooner, making scaled deployment in sensitive workflows easier. Others may appear to be moving faster precisely because they are deferring those disciplines and accumulating a governance backlog that has not yet surfaced in operating results. Recent events also show how quickly regulatory and security concerns can affect commercialization. Anthropic's temporary withdrawal of Fable 5 and Mythos 5 following a US government directive serves as a reminder that deployment risk may increasingly hinge on security, liability, and policy judgments. Public backlash is also growing, which could make regulation more political over time. As AI use broadens, questions of data provenance, auditability, security, and liability are likely to matter even more.

This shift from experimentation to scaled use is changing the investment question. Technical progress and broader adoption are making AI more commercially relevant, but they do not by themselves determine where durable returns will accrue. That depends increasingly on the economics of deployment and on which firms can turn AI capability into repeatable business action.

How AI economics are evolving

In our last edition of *VantagePoint: The Rearview Mirror Problem*, we argued that investors often mistake recent winners for future return drivers.² That risk is especially acute in AI. The first-wave beneficiaries are well known, and the infrastructure buildout has become the market's central focus. The harder question now is how the economics are evolving beneath that narrative, and which parts of the opportunity set can still deliver durable returns. We think investors should focus on three underwriting questions: Which bottlenecks are durable? Can rising revenues justify the capital required to sustain leadership? Where can value persist as AI becomes cheaper, more capable, and more widely available?


The economics are changing along with the technology. Constraints have shifted from training toward inference, usage, and deployment, and value capture is broadening with them. The relevant opportunity set now extends beyond frontier-training hardware to a broader mix of inference and data infrastructure, deployment software, workflow and permission layers, governance tools, and applications that shape how AI is embedded in business processes.

Earnings have begun to catch up with enthusiasm in parts of the AI ecosystem. Recent gains in AI-linked equities no longer rest on expectations alone. Several leading firms have reported strong revenue growth tied to AI demand, especially in semiconductors, cloud, and selected infrastructure segments. Private model providers, such as Anthropic, appear to be seeing similar momentum, though their economics remain less transparent. Stronger fundamentals validate part of the move. They do not settle the harder question of whether revenue growth will prove durable enough to justify the operating and capital costs required to sustain it.

² Please see Celia Dallas, "VantagePoint: The Rearview Mirror Problem," Cambridge Associates LLC, April 29, 2026.

Fundamentals have strengthened across the AI stack

December 31, 2022 – March 31, 2026



Source: FactSet Research Systems.

Notes: Quarterly revenue and operating income are aggregated at the group level; year-over-year growth rates are based on aggregated quarterly totals; operating margin is calculated as operating income divided by revenue; non-USD results are translated into USD. Groups include Hyperscalers (Microsoft, Meta, Alphabet, Amazon, and Oracle), AI Semiconductors (NVIDIA, AMD, TSMC, and Broadcom), and Memory (Micron and SK Hynix). For Memory, operating income growth is not shown where the prior-year quarter's operating income was negative.

Which bottlenecks are durable?

This first question deals with whether bottlenecks are persistent or simply reflect temporary undersupply. Early in the cycle, scarcity centered on training compute and raw GPU capacity. That is no longer the full story. As AI deployment scales, the tighter constraints are becoming more physical. Data center infrastructure, memory, and advanced packaging remain important, but power is emerging as the clearest hurdle.³ Reliable electricity, cooling, transmission, and the ability to bring new capacity online increasingly matter as much as access to chips.⁴


Falling prices in one layer of the stack do not remove constraints in another. Listed token prices have fallen sharply, which should broaden adoption. But lower prices do not mean lower compute demand. As models have become more capable and AI systems take on more complex tasks, the compute required to complete useful work has continued to rise. More autonomous and always-on systems will reinforce that trend by increasing token consumption and placing greater strain on the physical stack.

3 Please see Celia Dallas, “VantagePoint: Electrifying Returns in the AI Era,” Cambridge Associates LLC, October 31, 2025.

4 The situation is somewhat flipped in China, where access to the most advanced chips remains a constraint, while electricity is more available. Notably, China controls refinement of most critical material, such as rare earths.

Falling listed token prices, but rising effective token expenditure

Mar 2023 – May 2026 • Listed API input prices by model capability



December 1, 2025 – June 10, 2026 • Effective market expenditure per million tokens



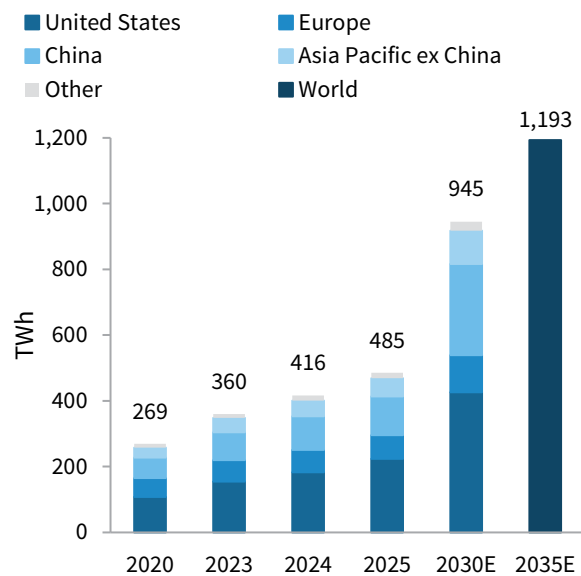
Sources: Bloomberg, L.P., Github, and LLM Price Check.

Notes: LHS shows listed API input token prices for representative models meeting or exceeding GPT-3.5 Turbo, GPT-4 Turbo, and GPT-4.5-like quality thresholds; GPT-5.5 is included to illustrate pricing for OpenAI’s most advanced model. Model quality is approximated using Elo-style benchmark scores derived from public model-comparison data. The chart is intended to show how the market price of access to a given level of model capability has changed over time, rather than to measure providers’ underlying inference costs directly. The chart is shown on a logarithmic scale. RHS shows the Silicon Data LLM Token Expenditure Index, a daily benchmark of effective spending levels in the actively traded LLM market.


Power deserves particular attention because it is increasingly one of the hardest constraints to relieve. Utilities, grid equipment, cooling, and related enabling infrastructure can be difficult to replicate quickly because they depend on permitting, transmission access, engineering capacity, and time to build. Those are more durable barriers than the temporary scarcity that can emerge in parts of the hardware stack early in a buildout.

Power availability is emerging as a key constraint on the next phase of AI deployment

2020–35E • Data center power demand



US installed capacity vs active queues (GW)



Sources: International Energy Agency and Lawrence Berkeley National Laboratory.

Notes: RHS chart data were gathered from the interconnection queues of seven Independent System Operators (ISOs) and Regional Transmission Organizations (RTOs), as well as 50 non-ISO balancing authorities—including utilities and Power Marketing Administrations. Together, these entities account for approximately 98% of the total installed electric generation capacity in the United States.

Memory and advanced packaging also remain important choke points, supporting stronger pricing and earnings across parts of the semiconductor ecosystem. The key investment question, however, is not whether these areas are constrained today, but whether those rents are likely to persist. Some supply bottlenecks may prove temporary as capacity expands. Others may reflect capabilities that are harder to replicate quickly.

In some parts of the market, security, compliance, and regulatory approval may also function as bottlenecks. Where customers need trusted systems for sensitive workflows, firms that can meet higher standards for resilience, auditability, and control may accrue durable advantage.

Not every bottleneck supports durable economics. Some stem from short-lived pricing power. Others are tied to assets, regulation, siting, expertise, or customer relationships that are harder to reproduce. Open and lower-cost models reinforce that point. They may broaden adoption and accelerate experimentation, but they also challenge the idea that frontier capabilities alone guarantee durable pricing power. In areas where customers do not require frontier performance or tightly integrated proprietary systems, improving open models are compressing economics at both the model and software layer.

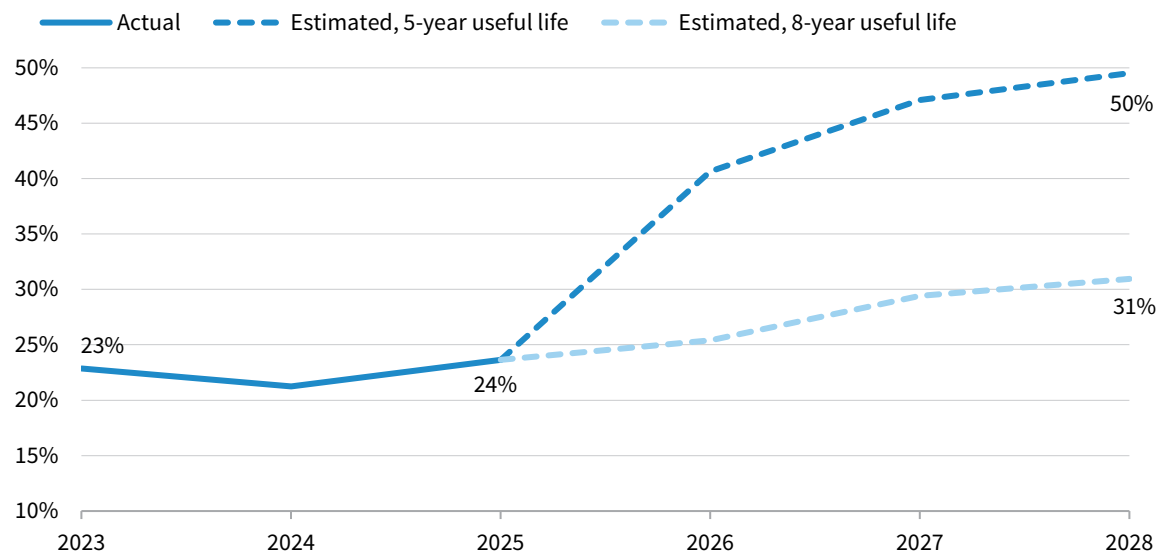
Can rising revenues justify the capital required to sustain leadership?

The second question asks whether rising revenues and earnings can justify the capital required to sustain AI leadership. That issue is now most visible in the capex cycle. For current spending to earn attractive returns, revenue growth must continue to catch up with investment, usage must remain high, enterprise monetization must deepen, and margins must hold up despite a much larger capital base.

Consensus expectations for the five major hyperscalers call for combined revenue to increase 54% while EBITDA is expected to increase about 111% from year-end 2025 through 2028. Over the same period, depreciation is expected to rise much faster. Depending on assumed asset lives, it could increase by roughly 175% to more than 340%. That drag is large enough to matter. At the low end of those estimates, 2028 depreciation would come close to the group's 2025 net income of \$405 billion.

Depreciation is absorbing a growing share of hyperscaler EBITDA

2023–28E • Calendar years • Depreciation / EBITDA ratio (%)



Source: FactSet Research Systems.

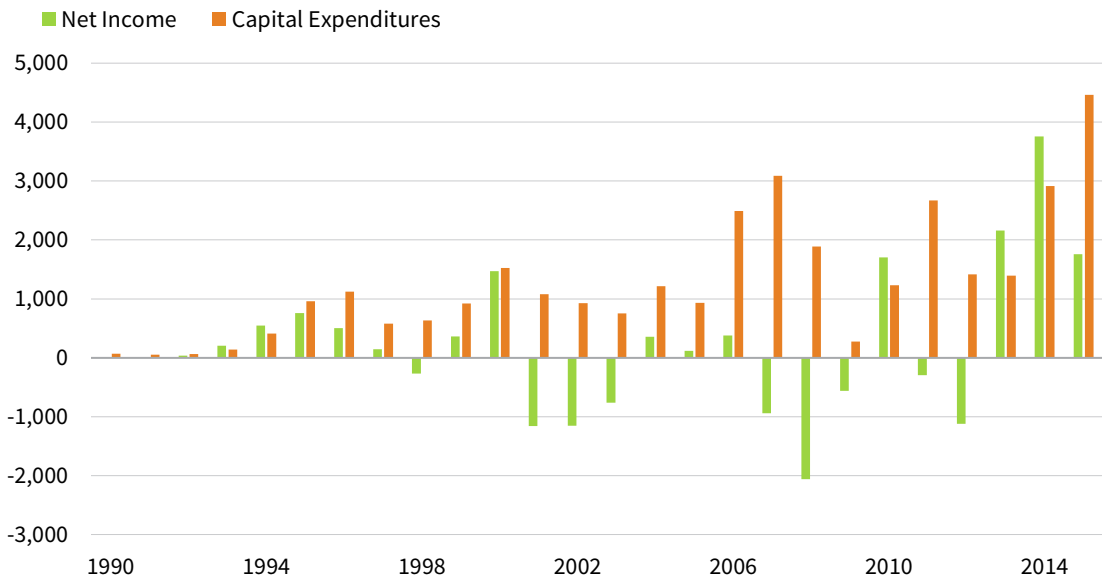
Notes: Hyperscalers include Microsoft, Meta, Alphabet, Amazon, and Oracle. Actual reported figures are used for 2023–25; FactSet mean estimates are used for 2026–28 capital expenditures and EBITDA. Forward depreciation is estimated for each company using a five- or eight-year straight-line capex-vintage methodology, under which each year's capital expenditures are depreciated evenly over five or eight years and roll off thereafter. The ratio is calculated as aggregate depreciation divided by aggregate EBITDA.

AI is making important parts of technology more capital intensive. Some parts of the market may still be valued as if AI were reinforcing capital-light software economics, when in fact it is making important parts of the stack more asset-heavy and operationally demanding. Investors should place more weight on depreciation, reinvestment needs, financing conditions, and the durability of pricing power on these more asset-heavy companies.

Memory is not a direct analogue for hyperscalers, and the current AI cycle has different drivers. Still, its history is a useful reminder that periods of tight supply, strong pricing, and high margins can look more durable than they prove to be once capacity expands. Shortages have repeatedly lifted margins and encouraged new investment, only to erode those same margins as supply caught up. AI may not follow that path exactly, but the lesson is familiar. Strong demand does not by itself protect returns when supply can respond and pricing power is not well defended.

Micron illustrates memory's cyclical past

1990–2015 • Annual • USD millions



Source: FactSet Research Systems.

Note: The chart ends in 2015 to preserve visibility into earlier cycles and capex responses; net income and capital expenditure figures are based on calendar years rather than fiscal years.

Indeed, Micron and SK Hynix—two of the memory companies most directly exposed to advanced AI demand—have increased capex by a combined 70% in each of the last two years, and consensus expects another roughly 55% increase in 2026. Investors should be careful not to assume that today’s strong pricing and profitability will persist unchanged as capital spending rises and supply responds.

The quality of demand matters as well. Investors should distinguish between durable end demand and demand supported by ecosystem-linked commercial arrangements, strategic subsidy, or circular deal structures that make near-term economics look stronger than they are. As the system matures, leverage and structured financing also deserve more attention. Risk rises when capital assumptions become aggressive ahead of proven cash flows. After rising by roughly \$900 billion since the start of this year, gross supply of investment-grade credit is expected to increase by about 17% over the full year to a record \$2.1 trillion, with much of the increase coming from hyperscalers and related infrastructure. Structured credit markets are expected to see data center securitizations rise by nearly 50% in 2026 to \$30 billion.⁵

⁵ For further discussion, see Wade O’Brien, Kate Thomes, and Abby Varney, “Should Credit Investors Be Concerned About Rising AI-Related Debt Issuance?,” Cambridge Associates LLC, March 3, 2026.

Stronger fundamentals have made the buildout more credible, but not necessarily more durable. High depreciation expense creates a demanding hurdle for hyperscalers that are increasingly competing with one another for business. Capital-intensive businesses facing rising competition may struggle even if the addressable market continues to grow. Not all participants will fare well. Semiconductors and advanced memory should benefit from tight supply and, in some cases, multi-year contracts, but supply is likely to catch up over time as capacity expands and technology becomes more efficient. Investors should be cautious about extrapolating today's pricing and profitability too far into the future.

Where can value persist as AI becomes cheaper, more capable, and more widely available?

The final question considers where durable value can persist. Access to models alone will not remain enough. As intelligence diffuses, we expect more defensible positions to belong to firms that control how it is used: who owns the workflow, governs permissions, controls distribution, and connects output to execution. Hyperscalers and other large platforms are trying to capture value across multiple layers of the stack through vertical integration, from compute and cloud infrastructure to model access, routing, deployment, and enterprise tooling.

The most important shift is in what software and adjacent systems actually do. As agentic systems begin to perform economically meaningful work rather than simply assist users, part of the addressable market shifts from software budgets to labor budgets, which are much larger. That may expand revenue pools, deepen integration, and create stronger business models. It may also intensify disruption across software, services, and selected consumer sectors.

Software and adjacent control layers may still be where much of the value ultimately accrues, but they also pose the hardest underwriting questions. AI may expand revenue pools even as it weakens traditional moats. Customers may expect broader functionality without proportional price increases, while model, compute, orchestration, and support costs remain material. We expect the stronger positions belong to firms that are deeply embedded in a workflow, possess privileged task-specific context, and can convert AI output into completed work rather than simply sell access to a feature.

That distinction matters because control over the workflow is different from access to the model. A firm that helps generate an answer may be easy to displace. In an agent-driven environment, durable advantage should rest increasingly on control over permissions, approvals, and execution rather than on data alone. Systems that determine what autonomous software can access, trigger, and complete could have a competitive upper hand over systems of record alone.

This logic extends beyond enterprise software. In consumer markets, such as commerce, education, and travel, AI is likely to reshape discovery, service, and execution. New products and business models should emerge. But these same layers may also face the greatest pressure from improving models and larger platforms, especially where functionality is easy to replicate, customer relationships are weak, or distribution is controlled by someone else.

Where value may persist across the AI stack

| Layer of the AI stack | What matters now | Why it may matter for returns | Main risk |
|---|---|--|---|
| Power and physical infrastructure (utilities, grid equipment, data-center developers) | Power availability, cooling, grid access, data-center capacity | Physical constraints may support durable returns where assets are scarce, regulated, or slow to build | Capital intensity, policy risk, local opposition, and eventual supply response |
| Semiconductors, memory, and hardware (NVIDIA, Broadcom, TSMC, SK Hynix) | GPUs, custom silicon, memory, packaging, and networking | Bottlenecks can create strong economics where supply is constrained and difficult to expand | Temporary scarcity may be mistaken for durable advantage as architectures and supply evolve |
| Cloud and compute infrastructure (Microsoft, Meta, Amazon, Alphabet, Oracle) | Inference demand, utilization, orchestration, and enterprise deployment | Rising usage can support durable demand for compute, storage, and AI tooling | Overbuilding, utilization pressure, and weaker pricing power as competition broadens |
| Frontier models (OpenAI, Anthropic, Google) | Model performance, reasoning ability, multimodal capability, enterprise reliability | Leading models may still support premium pricing and strong adoption, especially for high-value use cases | Overbuilding, utilization pressure, and weaker pricing power as competition broadens |
| Inference, data, and deployment infrastructure (Databricks, Snowflake, Clickhouse, and other companies providing hyperscaler tools, model-routing and optimization platforms, and emerging private companies) | Efficient inference, model routing, latency, cost control, reliability, and deployment across models | As usage scales, this layer may capture value by improving cost, performance, and flexibility for enterprises deploying AI in production | Overbuilding, utilization pressure, and weaker pricing power as competition broadens |
| Software and agentic applications (Salesforce, ServiceNow, Anthropic Claude Cowork, and emerging AI-native private companies) | Workflow integration, execution of economically meaningful tasks, customer ownership, and monetization of AI-enabled work | Applications may capture more value if AI moves from assistance to execution and displaces selected labor | Overbuilding, utilization pressure, and weaker pricing power as competition broadens |
| Governance and control layers (CrowdStrike, Palo Alto Networks, private tooling vendors) | Security, compliance, monitoring, testing, and operational reliability | These layers become more important as enterprises move from pilots to scaled deployment and seek trusted control over AI activity | Overbuilding, utilization pressure, and weaker pricing power as competition broadens |

Source: Cambridge Associates LLC.

Notes: Illustrative framework comparing where value may accrue across the AI stack as adoption broadens. Categories are simplified and not mutually exclusive.

The same logic also shapes investment underwriting. Strong adoption and fast top-line growth may not be enough if a company with limited bargaining powers depends heavily on a single model provider, hyperscaler, or distribution platform. Downstream growth may prove real without translating into durable economics. Strategic acquisition may become a common end state for promising firms, supporting investment outcomes alongside a select group of independent long-duration compounders. For private equity and venture investors, underwriting should place more weight on customer ownership, monetization after model and compute costs, governance quality, likely end states, and the durability of economics if acquisition interest fades.

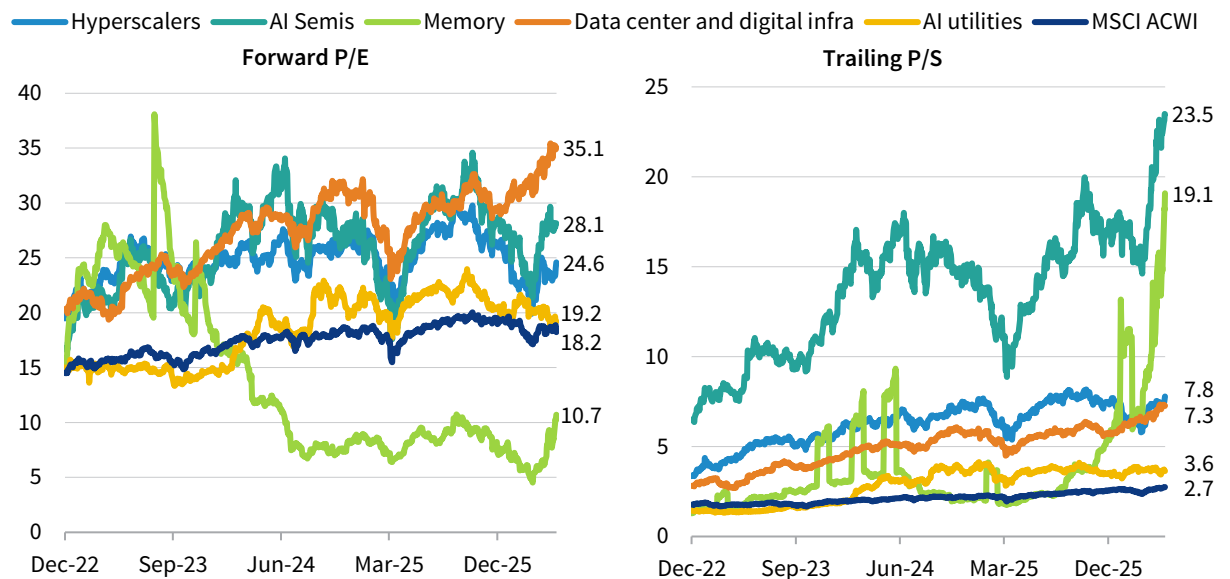
Investment implications

AI should be treated as a system-wide set of exposures rather than a narrow thematic trade. We see the stronger opportunities ahead in harder-to-relieve bottlenecks—especially power and related infrastructure—alongside the software and application layers that govern deployment in real workflows and emerging AI-native businesses that can reshape industry economics. The same shift should benefit adopters that use AI to improve their own economics while increasing disruption risk for incumbents that fail to adapt or are displaced by new business models.

This shift argues for more caution toward parts of the AI ecosystem where expectations, capital spending, and competition have all risen sharply at once. Large hyperscalers remain central to the buildout and may continue to benefit from scale, distribution, and enterprise integration. But they are also engaged in an increasingly costly race to secure compute, power, and physical

Valuation dispersion across AI-linked groups remains wide

December 31, 2022 – May 31, 2026 • Daily



Sources: FactSet Research Systems and MSCI Inc. MSCI data provided "as is" without any express or implied warranties.

Notes: Calculations use harmonic average index methodology and are based on daily price data from Factset. Baskets include Hyperscalers: Microsoft, Meta, Alphabet, Amazon, and Oracle; AI semiconductors: NVIDIA, AMD, TSMC, and Broadcom; Memory: Micron and SK Hynix; Data center and digital infrastructure: Equinix, Digital Realty, Vertiv, Eaton, and Schneider Electric; AI utilities: Vistra, Constellation Energy, NextEra Energy, Duke Energy, and Southern Company.

infrastructure, and the associated depreciation burden is becoming harder to ignore. We therefore lean away from the most crowded first-wave winners, particularly where valuations still leave limited room for disappointment. The same caution applies to parts of semiconductors and memory, where recent earnings strength has been real, but history suggests investors should be careful not to mistake tight supply and current pricing power for durable advantage.

By contrast, we are more constructive on select infrastructure and real assets tied to harder-to-relieve constraints, particularly electricity infrastructure, grid access, and related enabling assets. These areas appear better positioned to benefit as AI deployment scales and physical bottlenecks become more binding.

The widening opportunity set also creates room for emerging disruptors, many of which were inconceivable before recent AI advances. As AI becomes cheaper, more capable, and more widely available, value will migrate toward firms that control workflows, permissions, customer relationships, and operational integration rather than those relying on thin wrappers or temporary model arbitrage. In software, the more durable positions are likely to belong to companies that can embed AI into economically meaningful tasks, govern it effectively, and monetize completed work rather than simple access. Private equity may benefit where businesses need capital, operational support, and technology investment to adapt successfully to AI-driven changes in cost structure and competition.

Venture capital remains the key channel for accessing emerging AI-native companies and disruptive business models, and investors seeking that upside likely need some participation in private markets. But this technology cycle is still early and, as in past cycles, a few winners are likely to emerge alongside many losers as innovation advances faster than commercial adoption. In that environment, exposure to AI is not the same as access to strong investment returns. Because private commitments are long-lived, pacing matters as much as manager selection. Investors should continue to allocate selectively, with discipline on timing and valuation rather than rushing to add exposure simply because the theme is compelling. As a new wave of highly anticipated technology IPOs comes to market this year, investors should be thoughtful about re-deploying capital to venture capital, balancing those opportunities against other market segments with more attractive valuations and differentiated return potential.

The same framework should also be applied defensively. AI is not only a source of new opportunity. It is also a source of disruption risk in existing holdings. Businesses with weak differentiation, labor-intensive models, or information-heavy processes may be more vulnerable than they appear, even if they sit outside any obvious AI category. Equity long/short hedge funds may be well positioned to benefit from rising dispersion as AI creates clearer winners and losers across software, services, and other information-intensive industries. As disrupted companies—especially ones that took on private debt during the period of zero interest rates and high valuation from 2021 to early 2022—struggle to refinance over the next few years, stressed and distressed opportunities may emerge.

AI remains an important area of exposure. From here, we expect the best opportunities to come from identifying durable bottlenecks, defensible control points, and the businesses most likely to benefit from disruption rather than suffer from it. Active management across public and private markets will be central to success. ■

Graham Landrith and Justin Hopper also contributed to this publication.

INDEX DISCLOSURE

MSCI All Country World Index (ACWI)

The MSCI ACWI captures large- and mid-cap representation across 23 developed markets (DM) and 24 emerging markets (EM) countries. With 2,558 constituents, the index covers approximately 85% of the global investable equity opportunity set. DM countries include Australia, Austria, Belgium, Canada, Denmark, Finland, France, Germany, Hong Kong, Ireland, Israel, Italy, Japan, the Netherlands, New Zealand, Norway, Portugal, Singapore, Spain, Sweden, Switzerland, the United Kingdom, and the United States. EM countries include Brazil, Chile, China, Colombia, Czech Republic, Egypt, Greece, Hungary, India, Indonesia, Korea, Kuwait, Malaysia, Mexico, Peru, the Philippines, Poland, Qatar, Saudi Arabia, South Africa, Taiwan, Thailand, Turkey, and the United Arab Emirates.

Copyright © 2026 by Cambridge Associates. All rights reserved.

This document, including but not limited to text, graphics, images, and logos, is the property of Cambridge Associates and is protected under applicable copyright, trademark, and intellectual property laws. You may not copy, modify, or further distribute copies of this document without written permission from Cambridge Associates ("CA"). You may not remove, alter, or obscure any copyright, trademark, or other proprietary notices contained within this document. This document is confidential and not for further distribution, unless and except to the extent such use or distribution is in accordance with an agreement with CA or otherwise authorized in writing by CA.

This report is provided for informational purposes only. The information does not represent investment advice or recommendations, nor does it constitute an offer to sell or a solicitation of an offer to buy any securities. Any references to specific investments are for illustrative purposes only. The information herein does not constitute a personal recommendation or take into account the particular investment objectives, financial situations, or needs of individual clients. Information in this report or on which the information is based may be based on publicly available data. CA considers such data reliable but does not represent it as accurate, complete, or independently verified, and it should not be relied on as such. Nothing contained in this report should be construed as the provision of tax, accounting, or legal advice.

Past performance is not a reliable indicator of future results. All financial investments involve risk. Depending on the type of investment, losses can be unlimited.

Any information or opinions provided in this report are as of the date of the report, and CA is under no obligation to update the information or communicate that any updates have been made. Information contained herein may have been provided by third parties, including investment firms providing information on returns and assets under management, and may not have been independently verified.

Cambridge Associates is a global group of companies that provide investment management, investment advisory, research, and performance reporting services. For the purposes of this document "us", "the Firm", "our", "we", "CA", "Cambridge Associates", and similar terms refer collectively to the following list of companies. Similarly, unless otherwise stated the figures provided are the combined total for the following list of companies: Cambridge Associates, LLC (a registered investment adviser with the US Securities and Exchange Commission, a Commodity Trading Adviser registered with the US Commodity Futures Trading Commission and National Futures Association, and a Massachusetts limited liability company with offices in Arlington, VA; Boston, MA; Dallas, TX; New York, NY; and San Francisco, CA), Cambridge Associates Limited (a registered limited company in England and Wales, No. 06135829, that is authorized and regulated by the UK Financial Conduct Authority in the conduct of Investment Business, reference number: 474331); Cambridge Associates GmbH (authorized and regulated by the Bundesanstalt für Finanzdienstleistungsaufsicht ("BaFin"), Identification Number: 155510), Cambridge Associates Asia Pte Ltd (a Singapore corporation, registration No. 200101063G, which holds a Capital Market Services License to conduct Fund Management for Accredited and/or Institutional Investors only by the Monetary Authority of Singapore), Cambridge Associates Limited, LLC (a Massachusetts limited liability company with a branch office in Sydney, Australia, a registered investment adviser with the US Securities and Exchange Commission and registered in several Canadian provinces ARBN 109 366 654), Cambridge Associates Investment Consultancy (Beijing) Ltd (a wholly owned subsidiary of Cambridge Associates, LLC which is registered with the Beijing Administration for Industry and Commerce, registration No. 110000450174972), Cambridge Associates (Hong Kong) Private Limited (a Hong Kong Private Limited Company licensed by the Securities and Futures Commission of Hong Kong to conduct the regulated activity of advising on securities to professional investors), Cambridge Associates AG (a Swiss Limited Company, registration number CHE-115.905.353, that is authorized and Regulated by the Swiss Financial Market Supervisory Authority (FINMA), and Cambridge Associates (DIFC) Limited (incorporated as a Private Company and regulated by the Dubai Financial Services Authority, License Number: FO11237).