

HYPE VS. HIGH CONVICTION

A Quality Approach to Investing in AI

Tom Hancock | March 2026

Three years after the launch of ChatGPT 3.5, positioning around artificial intelligence-related companies has become a (perhaps “the”) critical decision for equity investors. AI investment spending by companies far exceeds current revenue generation from end-market use cases, leaving highly leveraged players facing existential risks.

Still, companies like Microsoft, Alphabet, and Nvidia are aggressively funding innovation from robust cash flows, and, at least superficially, trade at valuation multiples that are moderate relative to the frenzy surrounding the technology. But there is a frenzy.

For investors, the real opportunity lies in distinguishing between companies with durable competitive advantages and those exposed to the boom-and-bust cycle. This nuance is often overlooked in mainstream commentary, but it is essential for portfolio strategy in today’s AI-driven market.

At GMO, we are, of course, excited by the promise of this new technology. With our Quality orientation, what most excites us is that many of the companies driving AI innovation are highly profitable businesses that we believe will remain the winners over time as technology continues to advance. The GMO Quality Strategy focuses its investments on companies with proven cash flows, differentiation, barriers to entry that enable high returns on investment, and strong balance sheets.

As we navigate a path through the current potentially hazardous landscape, our objective is to identify and invest in companies that we have high confidence will ultimately thrive and possess the resilience to withstand and navigate unpredictable storms along the way.

In this piece, we review our approach to different types of AI business models and identify the opportunities and risks we see.

Dissecting the AI Value Chain

We categorize AI businesses into four distinct activities or layers, each with its own set of pros and cons.

1. **Applications:** This is how users, both consumer and corporate, can access AI insights – live examples include subscriptions to Microsoft’s Copilot or the Cursor coding support tools, while future use cases are imagined by thousands of start-ups across the world.
2. **Large Language Models (LLMs):** At the heart of the revolution sit ChatGPT, Claude, Gemini, and the rest, covering text, images, data, and everything in between.
3. **Compute (Hyperscalers & NeoClouds):** This compute layer provides the muscle both for training and inference on AI models
4. **Suppliers:** This is a broad range of upstream companies, including the semiconductor supply chain (e.g., Nvidia, Broadcom, TSMC, and server manufacturers) as well as the Industrial and Utility companies needed to build and power datacenters (e.g., Quanta Resources, Vertiv, and Constellation Energy).

RISKS AND OPPORTUNITIES ACROSS THE AI VALUE CHAIN

	RISK	ADVANTAGE
 APPLICATIONS	<ul style="list-style-type: none"> Picking markets Picking winners 	<ul style="list-style-type: none"> Legacy players have data and customer lock-in
 LLMs	<ul style="list-style-type: none"> Massive training investments precede inference monetization Where's the cash? Model plateauing and commoditization 	<ul style="list-style-type: none"> The heart of a technological revolution
 HYPERSCALERS	<ul style="list-style-type: none"> AI disruption 	<ul style="list-style-type: none"> Strong cash flow today Business diversification Customer lock-in Internal workloads Visibility and flexibility
 SUPPLIERS	<ul style="list-style-type: none"> Volatile revenue, tied to the growth in hyperscaler capacity 	(+/-) Focus on differentiation, diversification, and financial strength

APPLICATIONS

Applications face the classic problem that it is much easier to recognize the tremendous potential of AI broadly than it is to identify the specific winning use case. In the early days of the internet, social media was not an obvious application. During the rise of the smartphone, few anticipated the emergence of Uber.

We are most comfortable participating in the application layer through incumbent software providers, where AI functionality can be added onto an existing product or platform to increase lock-in or raise monetization. That group includes companies like Microsoft with CoPilot or Salesforce with Agentforce. It is worth noting that, of these two examples, we are well within the consensus on Microsoft, whereas with Salesforce, we are taking a more contrarian view that increased productivity and lowered costs will offset any reduction in seats.

LARGE LANGUAGE MODELS (LLMs)

The pure LLM companies are privately held and not accessible to many investors. Perhaps that is fortunate, as we see them as relatively risky given the massive investments they must make to train models substantially in advance of the use cases being crystallized.

We have further concerns about whether the various models will ultimately be differentiated or whether there will be some plateauing of capabilities. Although LLM companies are the disruptors at the heart of a technological revolution, we only invest directly through more diversified, cash-rich companies, such as Alphabet.

COMPUTE (HYPERSCALERS & NEOCLOUDS)

The Hyperscalers are highly profitable incumbents today. As such, they do face at least the possibility of disruption, if only because their status quo is so good and there is only one way to go. But they have the advantage of strong cash flow and business diversification.

It is non-trivial for their customers to switch providers, and the companies have their own internal workloads to offset costs and gather insights. Given their position in the ecosystem, the companies enjoy clearer visibility into any potential slowdown and have the flexibility to scale their capacity investments in response to demand.

We believe that, for the most part, they are exhibiting prudence in their investments. Microsoft invests behind demand, and lets the more specialized NeoCloud providers, like Coreweave, take on its incremental investments. Even Meta has demonstrated, with its recent "year of efficiency," that it is willing to dial back capital expenditures when appropriate.

SUPPLIERS

We see the Supplier category as a mixed bag. Nvidia grew to become the world’s largest company by market value, thanks to strong product differentiation and a wide moat. Not all Suppliers have that.

Some Suppliers have built in resilience through a strong balance sheet and business diversification. Others are more concentrated or have taken on leverage. This is important as Suppliers face more risk to their fundamentals than the upstream Hyperscalers (as we will shortly address).

INTEGRATED PROVIDERS

The mapping of companies to business models is not 1-1. While OpenAI and Nvidia sit relatively cleanly in the “LLM” and “Supplier” layers respectively, Alphabet spans all four with its core AI applications built for search and ad targeting, its own LLM Gemini, its GCP cloud computing platform, and its internally designed TPU chips for internal workload that are now also reported to be sold to external customers.

Cash is King: Tracking Revenue Flow Across AI Layers

As we seek a margin of safety, we examine the relative cash flow position of various AI stocks. We model the risks and opportunities for various companies, considering how varying levels of future AI revenue will flow through the layers as AI investment spend is converted into revenue for the supply chain.

For example, most Hyperscalers generate enough cash from their existing businesses to comfortably fund their AI investments. It would be painful if that investment were to turn out to deliver a low return, but it would not be fatal.

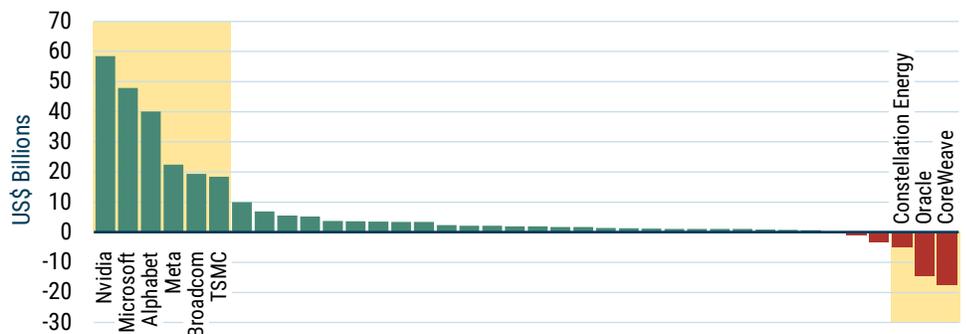
However, within the compute layer, companies like Oracle and CoreWeave are more aggressive and, therefore, are at the mercy of funding markets to support their investment programs. In the case of Oracle, the company has made an impressive transition from a sleepy and unloved (but cash-rich) legacy software company to an aggressive investor in cloud data centers.

We had been long-time holders of the stock and benefited from the multiple expansion that came with this transition. However, we have exited that position as we no longer feel comfortable with the level of leverage. We can see a path where they deliver high returns, but in a world where that plays out, we also see other safer stocks benefiting.

MOST OF THE PARTICIPANTS HAVE A LOT OF CASH

Cash Flow from operations in excess of estimated CapEx

CASH FLOW FROM OPERATIONS MINUS CAPEX



An example of a company with a relatively controlled range of outcomes is Taiwan Semiconductor, which manufactures GPUs for Nvidia and AMD, as well as custom ASICs for Broadcom and other clients. TSMC generally operates at a base level capex-to-sales ratio of around 30%, which it ramps up during growth cycles, peaking at around 50% historically. They have seen this rodeo before as the smartphone revolution was one such cycle, and this is another. With their gross margins of 50% or more, their investments are well-funded. And should it turn out that TSMC built capacity before demand, the risk is hardly existential.

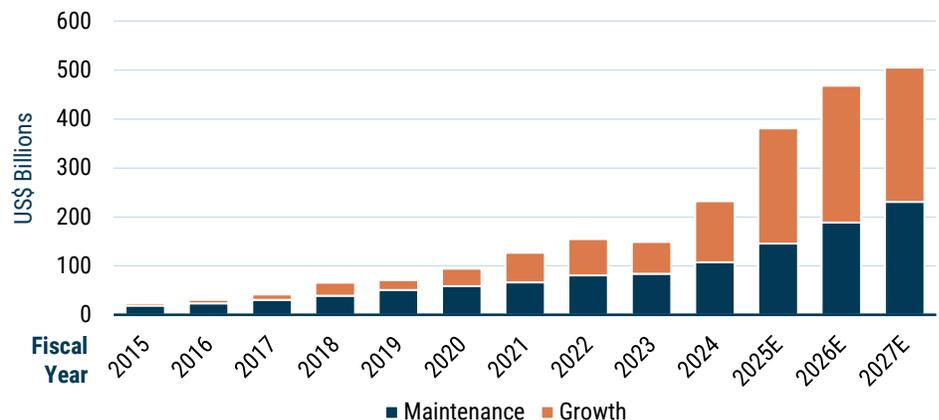
In general, it is helpful to split the investment spend at the Compute layer into two buckets. One is the “maintenance capex” that Hyperscalers need to maintain their current compute capacity. This is proxied by their depreciation schedules (and we think the controversy about whether 4, 5, or 6 years is the right number for a GPU is not of the highest consequence.)

The remainder is “growth capex”. This is investment to extend compute capacity. While the growth capex supports a fraction of the Hyperscalers’ revenues, it accounts for over half of the Suppliers’. And so, if the Hyperscalers were to take their foot off the gas and say, “Hmm, capacity is good for the next year or so”, there would be a bloodbath in the supply chain. It is this risk that makes us think Nvidia is not as cheap as Microsoft, even though the difference in their current P/E’s is not large.

HYPERSCALER CAPEX: GROWTH VS. MAINTENANCE

When the buildout stops, the companies revert to maintenance CapEx

HYPERSCALER CAPEX



Source: Visible Alpha

The “Hyperscalers” include Microsoft, Alphabet, Amazon, Meta, Oracle, and CoreWeave. Maintenance CapEx is estimated based on depreciation schedules.

Conclusions

Summing this up, we look to invest where we find companies with:

- Durable competitive advantages that support profitable growth
- The opportunity to prosper across a wide range of scenarios
- Resilience to downturns via less volatile AI revenue, reliable non-AI revenue, and balance sheet strength

We find those features across most established Hyperscalers and selectively within the Suppliers. Among the Suppliers, we focus on the diversified and strong businesses of TSMC



Tom Hancock

Dr. Hancock is the Head of GMO's Focused Equity team and a portfolio manager for GMO's Quality Strategies. Dr. Hancock

is a partner of the firm. Previously at GMO, he was Co-Head of the Global Equity team. Prior to joining GMO in 1995, he was a research scientist at Siemens and a software engineer at IBM. Dr. Hancock holds BS and MS degrees from Rensselaer Polytechnic Institute and a PhD in Computer Science from Harvard University.

Disclaimer

The views expressed are the views of Tom Hancock through March 2026 and are subject to change at any time based on market and other conditions. This is not an offer or solicitation for the purchase or sale of any security and should not be construed as such. References to specific securities and issuers are for illustrative purposes only and are not intended to be and should not be interpreted as recommendations to purchase or sell such securities.

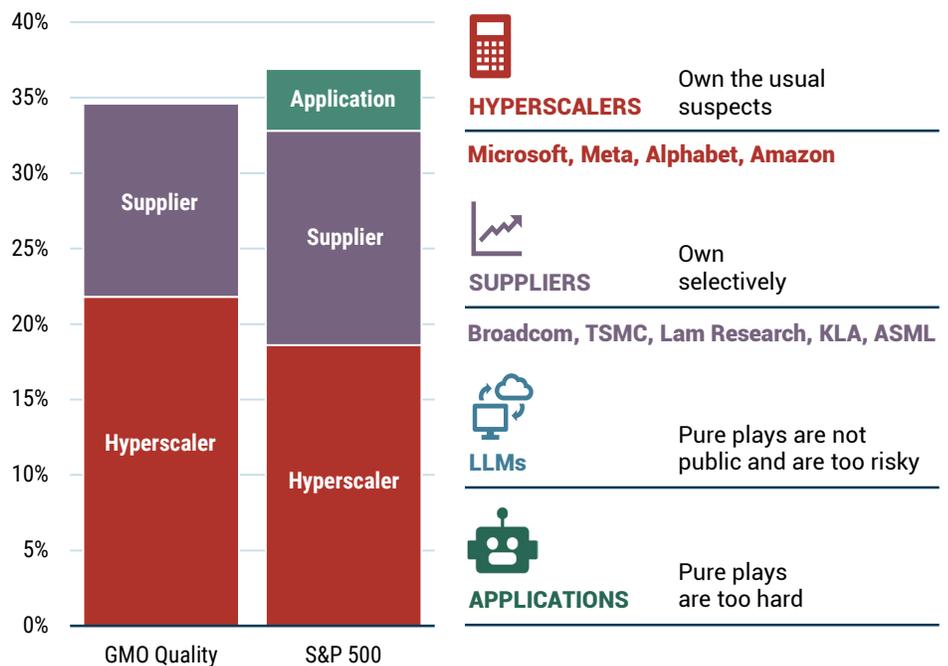
Copyright © 2026 by GMO LLC. All rights reserved.

and some of their semiconductor equipment supply chain. We also believe that Broadcom's focus on a few large customers with more established use cases (primarily centered on content targeting) makes them a safer bet.

Outside of our "direct AI" exposures, our stock selection reflects our contrarian position that AI is likely to be a benefit rather than a threat for providers that tightly control critical enterprise data, such as in software and services, or in sectors like healthcare, where the benefits of AI use cases appear not to be fully priced in.

In aggregate, we believe our positioning leaves us substantially exposed to AI upside; yet, even so, we remain a bit below the market, both on a notional level (adding up stock weights) and as measured by observed volatility. This gives us the dry powder to top up during sell-offs, as we did post DeepSeek and Liberation Day, while leaving us not too far behind on the way up if those sell-offs occur from much higher levels than where we currently sit.

GMO QUALITY POSITIONING



Source: S&P, GMO

The above information is based on a representative account in the Strategy selected because it has the fewest restrictions and best represents the implementation of the Strategy. S&P does not guarantee the accuracy, adequacy, completeness or availability of any data or information and is not responsible for any errors or omissions from the use of such data or information. Reproduction of the data or information in any form is prohibited except with the prior written permission of S&P or its third-party licensors. Please visit <https://www.gmo.com/americas/benchmark-disclaimers/> to review the complete benchmark disclaimer notice.

For more on this topic, watch our recent presentation from [GMO's 2025 Fall Conference](#).