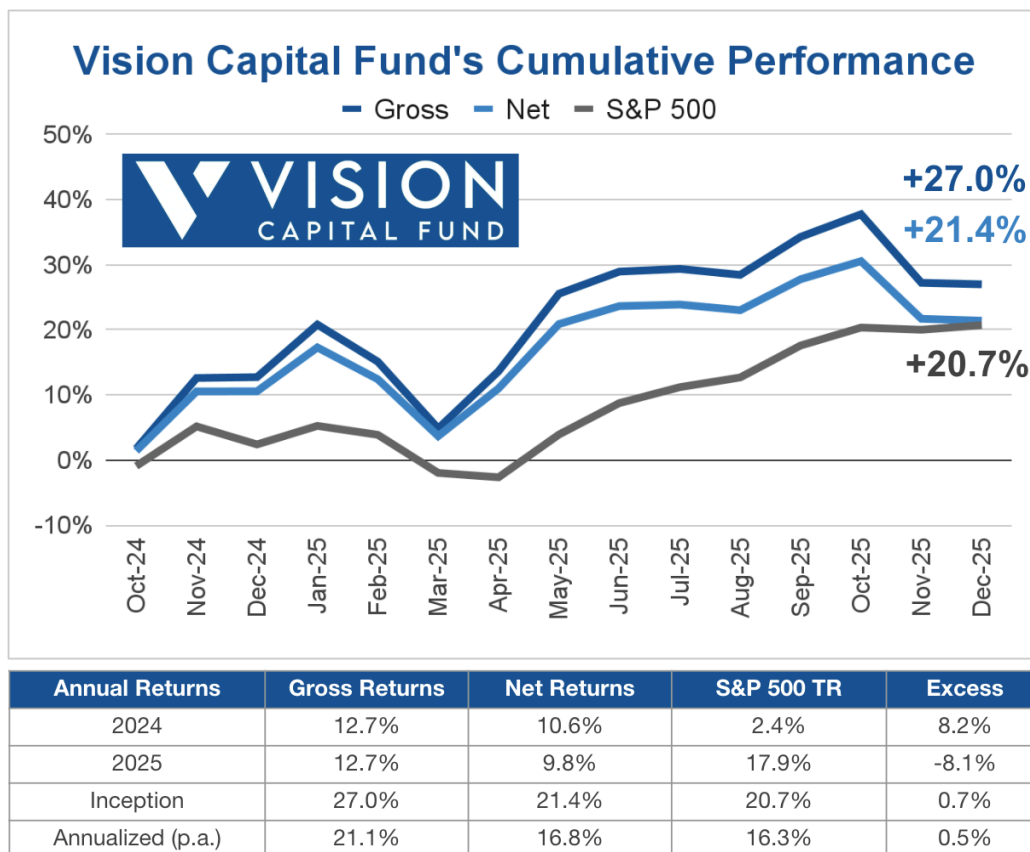January 9, 2026

Dear Partners,

Our ultimate financial measure, which we want to drive most over the long term, is the growth in Vision Capital Fund's net asset value, which will ultimately lead to the growth of your capital.

While we are still in our very early innings, having only started over a year ago, we request that our performance be evaluated over a longer timeframe of three to five years and beyond.

Performance Overview

Vision Capital Fund's unaudited gross and net returns for Q4 2025 were **-5.4%** and **-5.0%**. This puts our gross and net returns for 2025 at **+12.7%** and **+9.8%,** respectively. Our cumulative returns from inception (1 Oct 2024) were +27.0% and +21.4%, respectively[1], versus +20.7% for the S&P 500. Our annualized gross and net returns are **+21.1% p.a.** and **+16.8% p.a.**, respectively, versus +16.3% p.a. for the S&P 500.

Annual Returns	Gross Returns	Net Returns	S&P 500 TR	Excess
2024	12.7%	10.6%	2.4%	8.2%
2025	12.7%	9.8%	17.9%	-8.1%
Inception	27.0%	21.4%	20.7%	0.7%
Annualized (p.a.)	21.1%	16.8%	16.3%	0.5%

Source: Vision Capital Fund, Net (Class A Shares after fees) as of 31 Dec 2025, since inception 01 Oct 2024

Fund Commentary

While we still outperformed from inception, the weaker Q4 2025 resulted in underperformance for 2025. While we strive to outperform over the long term, we are unlikely to consistently outperform every year.

We simply cannot control what the market does in any month, quarter, or year, as most returns will come from changes in valuations rather than from business returns (i.e., the growth of revenues, earnings, and free cash flows). If we keep focusing on owning companies that can keep growing profitably and durably, reinvesting strongly at high returns, business returns will drive the majority of our long-term returns.

Our top 5 largest detractors were Pro Medicus, Zscaler, Wise, Mercado Libre, and The Trade Desk, all of which saw significant price declines. During the quarter, we took advantage of more attractive price levels to add to 9 of our existing positions: Amazon, Meta, Mercado Libre, NVIDIA, Pro Medicus, ServiceNow, Spotify, TSMC, and Wise. Several of which were among our largest decliners.

We initiated a new 2% position in Sea Limited (see more later), Southeast Asia's largest e-commerce platform (52% market share by 2024 GMV), in late December, bringing our total stock holdings to 28. We have long been admirers of management's disciplined approach to turning the business around since 2024 into a much higher-quality business in e-commerce and financial services beyond its original e-gaming business, and the recent 38% drawdown on fears of higher investments and lower profitability presented an attractive entry opportunity.

Our portfolio turnover remains zero. Given our intended long-duration investment holding periods of seeking to own the companies for as long as possible until growth slows rapidly, or if our thesis breaks, we do continue to expect portfolio turnover to remain low (i.e., <10%). This allows us to let our winners run high, rather than trimming our flowers to water our weeds.

Most of the money is made by buying, holding, and waiting, rather than by constantly rotating and trading, which most usually do. In the former, we put a lot of effort into making our single buying decision well. In contrast, in the latter, one has to make two buying-and-selling decisions correctly and repeat them, and the probability of error increases. Some do this well, but we simply don't think we are any good at it.

We wrote three new investment memos on Spotify, Zscaler, and Sea Limited during the quarter. You may find all 22 of our memos in Partner Materials, accessible with your login credentials.

One Question We Constantly Ask Ourselves

Vision Capital Fund is a vehicle for us to own exceptional, growing companies that are top dogs riding secular tailwinds in winner-take-most markets with superior products and services, run by excellent, aligned management. We are long-term investors in businesses run by people, not short-term traders of stocks.

While the public stock market allows us to buy and sell stocks more easily, faster, and cheaper than before, we seek to be highly intentional in our investment decisions, allocating any incremental capital toward acquiring existing additional or new equity ownership stakes. We always prefer to look inward at our existing holdings before looking outside.

The first question we ask ourselves when considering an investment is, "**Can we own and hold on to this investment indefinitely (through the good and bad times)?**"
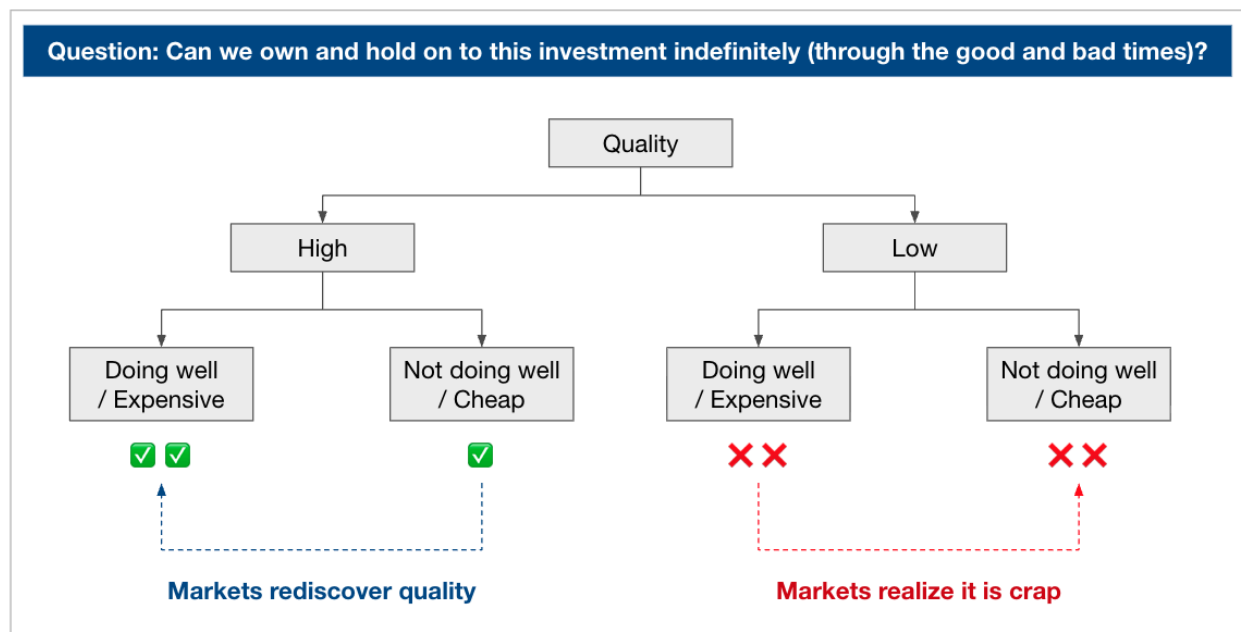
This seemingly simple question helps us avoid (1) statistically cheap lower-quality companies and (2) poorer-quality companies masquerading as higher-quality ones when they are doing well, allowing us to focus much more on higher-quality companies (see framework below).

As investors who allocate capital, we have to say yes and no all the time, to making an investment, or passing on one. Those two actions, in aggregate, will eventually determine our long-term returns and results. Even if we are not doing anything, inactivity has to be an active choice in itself, be it intentional or unintentional.

We recently wrote about how we thought about Type I errors (errors of commission, picking the losers) and Type II errors (errors of omission, missing the winners) here. We realized over time that while avoiding Type I errors is important, avoiding Type II errors is even more important, as the asymmetry of higher upside from winners far outweighs the limited downside of losers.

Using this single question of "**Can we own and hold on to this investment indefinitely?**", we hope this will give you more insight into how we think.

What did we say no to that turned out to be wrong for us so far?

The largest winners we missed out on investing in 2025 that did really well were the artificial intelligence (AI) memory supply chain players in Dynamic Random Access Memory (DRAM) / High Bandwidth Memory (HBM), with SK Hynix +295%, Micron +227%, Samsung +141%, in Hard Disk Drive (HDD) with Western Digital +268%, Seagate +219%, and in NAND with SanDisk +559%. Our decision to pass on these stemmed from our view of AI's foundational technology, LLMs, and from whether the current memory boom is sustainable.

Name	Ticker	Chart	% 52WH	2025 Returns	Theme
SK Hynix Inc	KRX:000660		-4.1%	295.4%	HBM
Micron Technology Inc	MU		-2.0%	226.8%	HBM
Samsung Electronics Co Ltd	KRX:005930		-3.9%	140.6%	HBM
SanDisk Corp	SNDK		-0.4%	559.4%	NAND
Western Digital Corp	WDC		-9.7%	268.3%	HDD
Seagate Technology Holdings PLC	STX		-7.2%	218.8%	HDD

Source: Vision Capital Fund (via Google Finance data)

While industry consolidation has turned the DRAM memory market into an oligopoly (see chart below), **memory has historically been a more commoditised part of the semiconductor supply chain** (earned essentially zero economic value from 1996-2012).
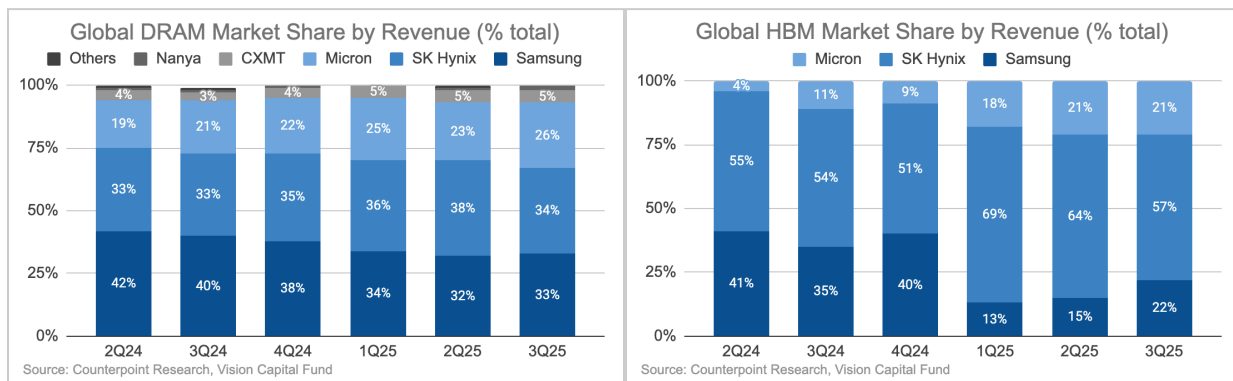
Memory prices and businesses are subject to larger, more volatile cyclical swings, frequently alternating between shortages and gluts. In addition, memory players have growing, positive working capital requirements, which require more capital to operate as they grow, and we tend to dislike companies with such requirements.

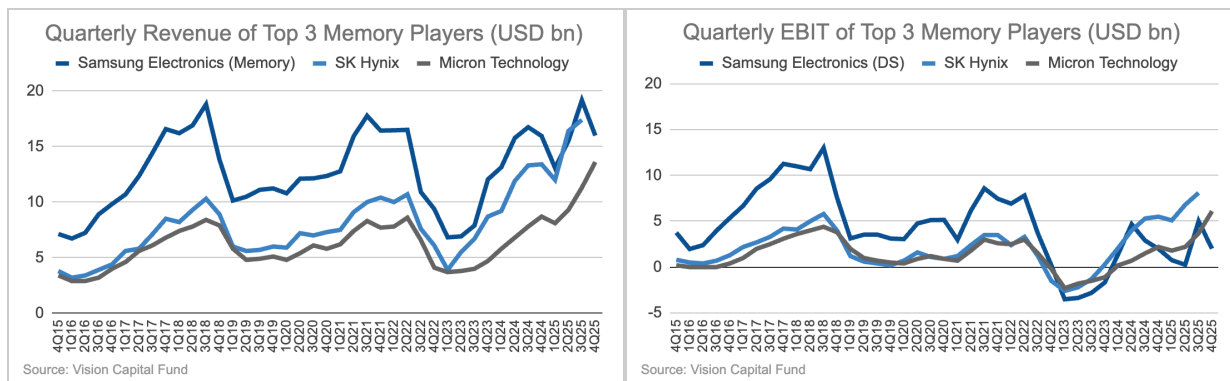**The Consolidation of the DRAM Market over the last 50 Years to 3 Main Players**

DRAM Market Share 1970-2020
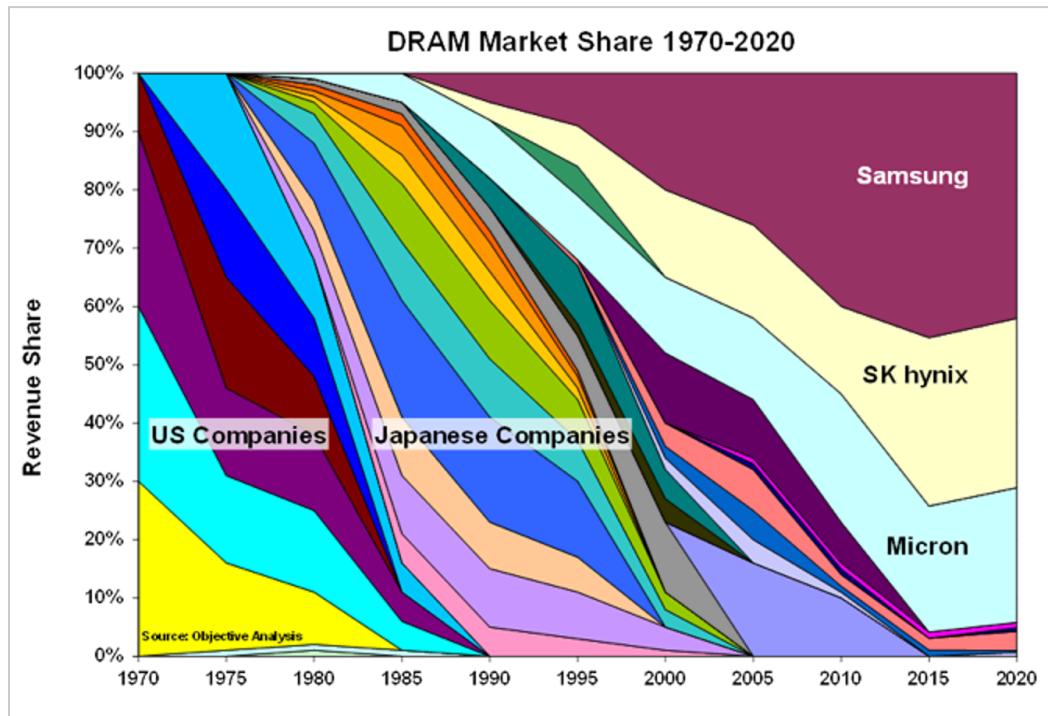
Source: [Potomac Institute, Objective Analysis](#)

The Oligopoly of Samsung, SK Hynix, and Micron in the Global DRAM and HBM Market

Global DRAM Market Share by Revenue (% total)

	2Q24	3Q24	4Q24	1Q25	2Q25	3Q25
Others	4%	3%	4%	5%	5%	5%
Micron	19%	21%	22%	25%	23%	26%
SK Hynix	33%	33%	35%	36%	38%	34%
Samsung	42%	40%	38%	34%	32%	33%

Source: Counterpoint Research, Vision Capital Fund

Global HBM Market Share by Revenue (% total)

	2Q24	3Q24	4Q24	1Q25	2Q25	3Q25
Micron	4%	11%	9%	18%	21%	21%
SK Hynix	55%	54%	51%	69%	64%	57%
Samsung	41%	35%	40%	13%	15%	22%

Source: Counterpoint Research, Vision Capital Fund

The Historical Cyclicality of Samsung, SK Hynix, and Micron's Quarterly Revenues & EBIT

Source: Vision Capital Fund

Source: Vision Capital Fund

While the current bottlenecks have led to higher memory prices and enormous revenues and profits, the question on our mind is whether these are short-term supernormal profits and whether they will normalise towards more sustainable lower levels over time. If we think they would, then investing in such companies involves us selling (or trading) them through these business cycles, which is not what we prefer to do.

The Historical Cyclicality of DRAM Spot Prices from 2000-2025
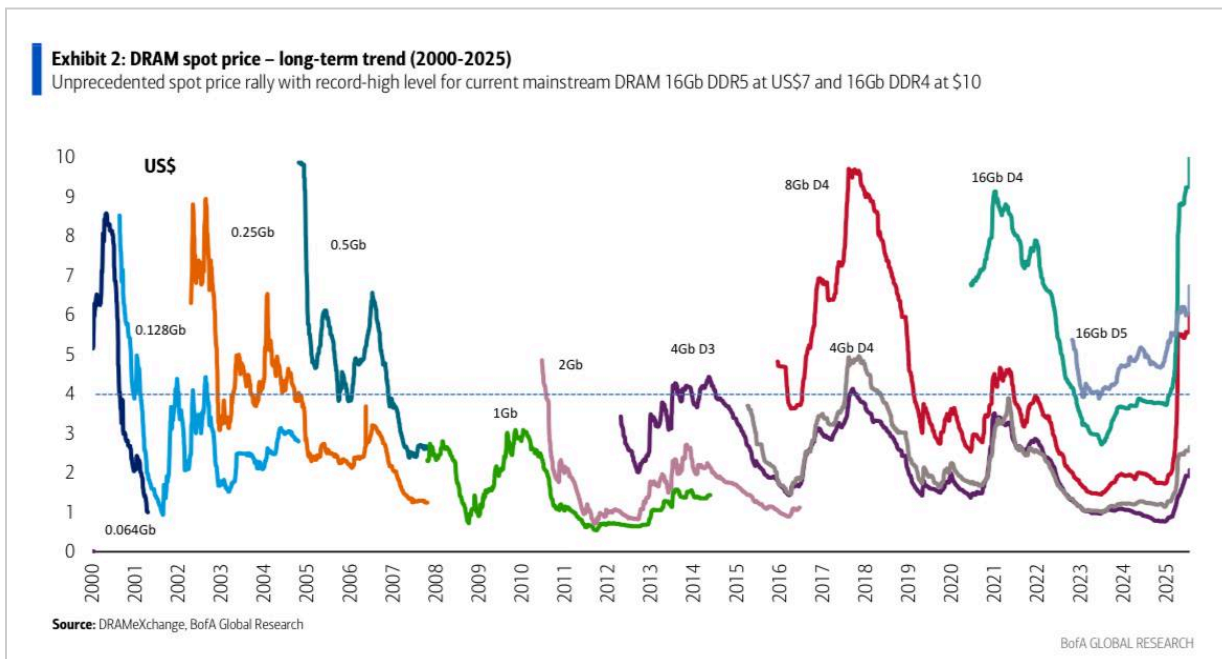
Source: [BOFA Global Research](#)

The Historical Price Chart of DDR5-6000 2x32GB over the last 18 months
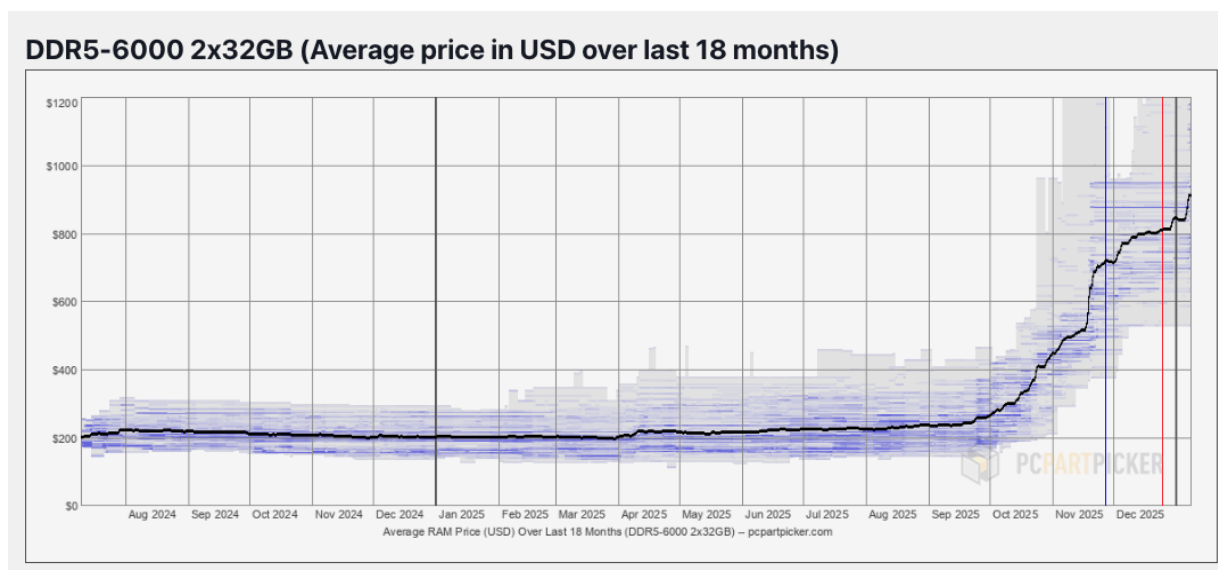
Source: [PC Part Picker](#) (9 Jan 2026)

This brings us back to our prevailing thoughts on large language models (LLMs). LLMs have undeniably revolutionized AI. Scaling LLMs has enabled impressive feats in reasoning, coding, and creative problem-solving. Models can reason step-by-step through complex problems, translate between languages fluently, and invent novel solutions under constraints.

Mimicry vs understanding: However, LLMs, at their core, are autoregressive models trained on pattern recognition on vast datasets of human output to predict the next token in a sequence. This seems to mimic human language use and some cognitive patterns, as humans often rely on probabilistic associations in speech and thought. But this mimicry remains superficial, as LLMs don't possess genuine understanding or consciousness. They sophistically statistically correlate patterns without grasping causality, context, real-world implications, and first principles.

The sweet spot for LLMs for now remains in areas where creativity, fluency, and approximation suffice. Think writing/coding assistance, content generation, or ideation in creative fields. They leverage their strength in mimicking human-like output without needing absolute truth. This has worked well for digital advertising, e-commerce (and e-payments), and cloud computing, where the majority of our holdings benefit from it.

In business contexts, LLMs have proven strong in tasks like summarising reports or automation customer service chats, but only when "grounded" via techniques like retrieval

augmented generation (RAG), which pulls in external data to reduce fabrications. But this mimicry hits hard limits in data-intensive or truth-critical applications, where errors are not tolerable. Hallucinations that confidently output plausible but false information persist as a core flaw of LLMs (a feature, not a bug), stemming from their statistical nature rather than genuine understanding (see chart below, the lower, the better, fewer hallucinations).

AA-Omniscience Hallucination Rate

AA-Omniscience Hallucination Rate (lower is better) measures how often the model answers incorrectly when it should have refused or admitted to not knowing the answer. It is defined as the proportion of incorrect answers out of all non-correct responses, i.e. incorrect / (incorrect + partial answers + not attempted).

Source: [Artificial Analysis](#) (9 Jan 2026)

LLMs are highly inefficient in handling context. In transformer architectures (the backbone of most LLMs), attention mechanisms require recomputing relationships among all tokens in the input sequence at each time step, leading to quadratic computational costs as context windows grow. While techniques or more efficient transformers mitigate this, they don't eliminate the core issue. Regeneration of the entire context for each new token prediction is wasteful and doesn't scale indefinitely. This is an architectural ceiling.

Memory is another bottleneck. The transformer architecture underpinning most LLMs is memory-hungry by design. LLMs are stateless by design. Every LLM conversation starts fresh from zero unless augmented by external systems like RAG, vector databases, or long-term memory stores. Humans don't re-learn every day. LLMs simulate continuity through context windows and memory hacks, but they are fundamentally different from how biological intelligence maintains identity and accumulated understanding.

As context windows balloon, tokens skyrocket, and so do compute and energy demands, regenerating the full cache for each inference step is inefficient, leading to scaling walls. Which is why even advanced models degrade in ultra-long contexts or require hacks like sparse attention. Human cognition, by comparison, relies on hierarchical, sparse memory (e.g., forgetting irrelevancies while retaining core knowledge), which LLMs currently mimic poorly without external aids like vector stores.

Hence, our growing and evolving view is that LLMs could well be a dead end for artificial general intelligence (AGI)/superintelligence because they lack formal understanding, leading to hallucinations or failures on common-sense tasks.

That said, it's unlikely that we will be stuck here forever. We suspect that we need a true, deeper breakthrough in architecture to dissolve this bottleneck, with scalable cures, not surface-level patches. While models like o1 show inference-time improvements, they still rely on patches rather than reinvention. If not, LLMs will not evolve beyond "enhanced autocomplete" for grounded tasks. They will remain as assistants and agents, not as autonomous decision-makers, which we humans continue to hold for now.

This perspective keeps us focused on durable, less cyclical exposures in our holdings in the Vision Capital Fund, and, for now, we prefer to cap our direct AI beneficiaries' exposure to TSMC and NVIDIA, which are our 4th- and 5th-largest positions. Separately, we are actively monitoring the strong progress in custom ASICs and, in particular, Google TPUs for specialized workloads, versus NVIDIA GPUs for general-purpose workloads.

What else did we say no to? Thoughts on OpenAI, Oracle, and Neoclouds.

With ambitious revenue targets of US$100bn and US$983bn by 2027 and 2030, respectively, versus its current US$20bn ARR, OpenAI has committed in advance to over US$1.4T in hardware and cloud infrastructure spend between 2025-2035. Despite our wanting OpenAI to succeed, which it has done thus far by shipping game-changing tech with real demand, OpenAI's massive spending spree and multi-year buildout, against the backdrop of its high cash burn of US$115b to 2029, sparks parallels to WeWork's wild overexpansion.
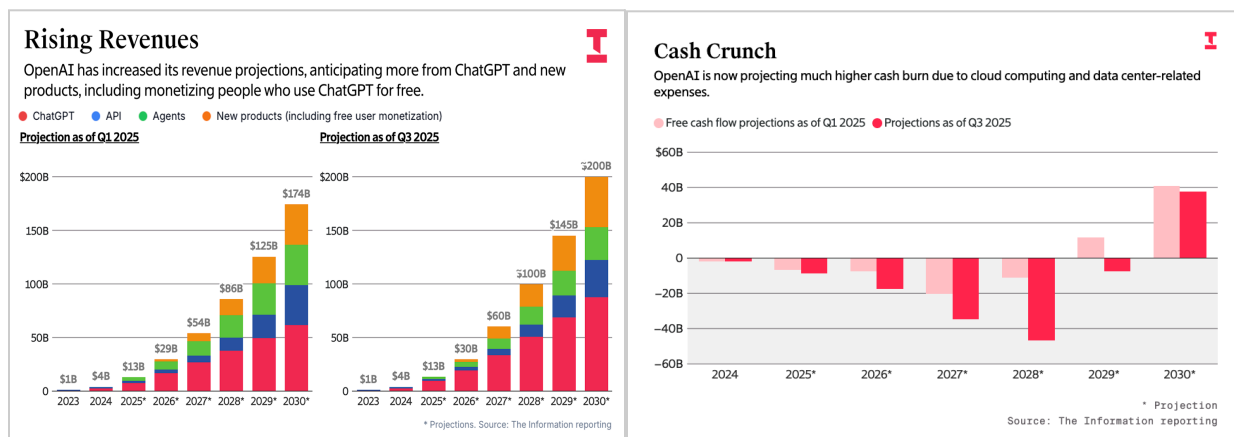
Source: The Information (Sep 2025)

While we think the upside would be enormous if OpenAI can crack artificial general intelligence (AGI) within the next decade, OpenAI's aggressive buildout does echo the WeWork vibe of betting on endless growth when combined with aggressive funding. The next few years could make or break it if OpenAI's losses keep piling up, or if OpenAI loses leadership to Google's Gemini, Anthropic's Claude, xAI's Grok, etc, and and unable to grow as rapidly and monetise.

We declined to invest in Oracle. OCI is the 4th-largest cloud hyperscaler after Amazon's AWS, Microsoft's Azure, and Google's GCP. While we like Oracle's rapidly growing Cloud (IaaS) business, which grew revenues +68% in 2Q26, and Oracle's rapid growth in remaining purchase obligations (RPO) from US$138bn to US$523bn over the last 3 quarters, and its negative working capital, we remain uncomfortable with OpenAI's singular overly high concentration of Oracle's RPOs (~US$300bn 5-year, ~57% share) and its highly leveraged balance sheet (net debt of $105bn) to fund its multi-year high growth capex requirements.

Top 4 Hyperscaler TTM Revenues, TTM Market Share, Revenue Growth Rates
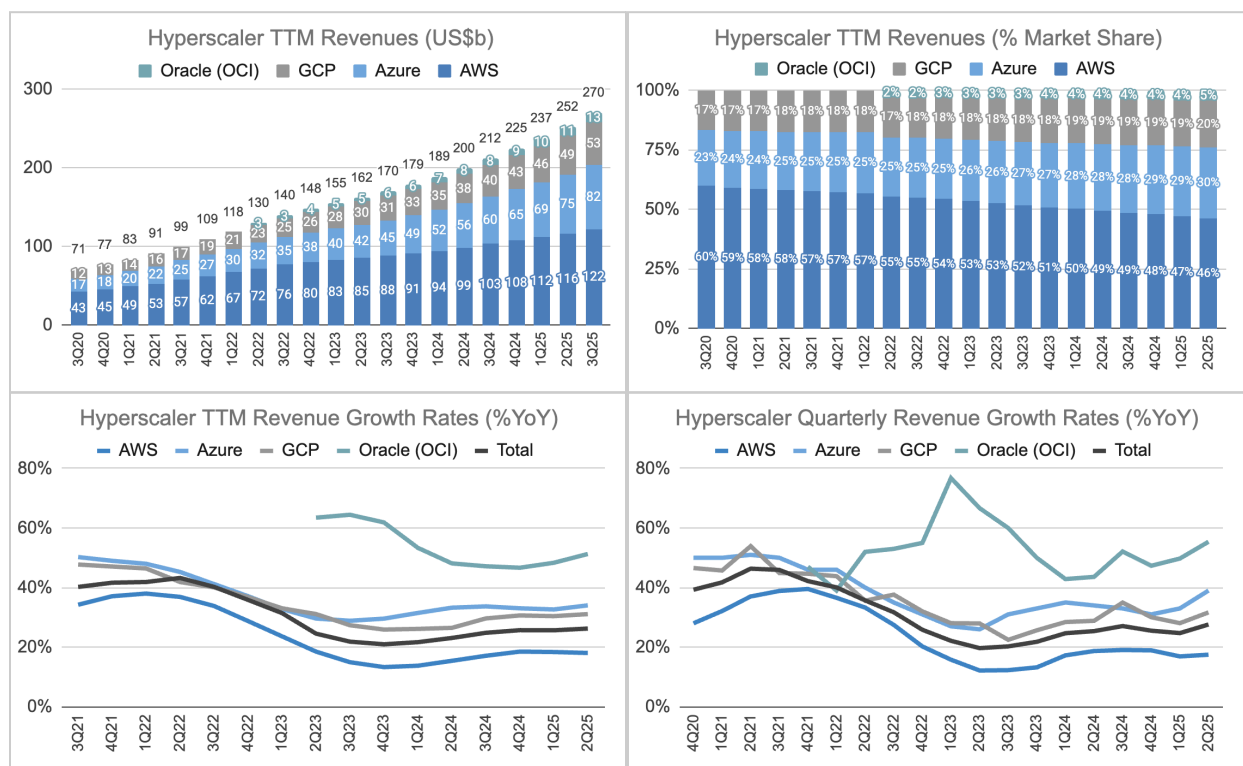
Similarly, we declined to invest in or chase the neoclouds, CoreWeave and Nebius, which posted strong gains of +79% and +174%, respectively, in 2025, on their ability to build and deliver GPU clusters faster at scale than the cloud hyperscalers. We actually thought the less

aggressive pace and more disciplined capacity expansion of the three largest hyperscalers was good and would avoid subsequent industry capacity overbuilding and glut.

By being more conservative, the hyperscalers (Microsoft, OpenAI, Meta) instead turned to the major neoclouds for supplemental compute, signing massive long-term contracts. Effectively, the neoclouds are a buffer for the strong GPU compute demand and are benefiting from it, but if demand slumps, they could well be the first to be hit.

In addition, we remained unconvinced that neoclouds, with their largely more commoditised bare-metal GPU cloud-as-a-service, would provide strong, durable production differentiation versus the more comprehensive full-stack ecosystems offered by the major cloud hyperscalers, particularly given emerging pricing compression, high capex burdens, and reliance on a handful of large customers. While Nebius is arguably a bit more full-stack with managed services than CoreWeave, we have our doubts about their strong, durable competitive advantage.

What did we say yes to? Sea Limited: Southeast Asia's Digital Economy Champion

As mentioned earlier, **we added Sea Limited (SE) in 4Q25** (see memo). Sea Limited represents our newest conviction play on Southeast Asia's (SEA) digital transformation. What began as a gaming company (Garena) has evolved into an integrated ecosystem spanning e-commerce (Shopee) and digital financial services (Monee), each reinforcing the others. Shopee dominates SEA e-commerce with ~52% market share, having outmaneuvered Alibaba's Lazada through relentless hyper-localization and strategic cost discipline.

Rather than chasing premium positioning, Shopee weaponized efficiency to win: filling trucks only when full (slower but cheaper), building networks of small sorting centers instead of mega-warehouses, and focusing on high-density urban routes. This counter-positioning created a logistics moat that competitors are struggling to replicate while Shopee's transaction data feeds Monee's credit scoring, enabling superior lending decisions with NPLs consistently below 2%.

Sea Limited's business has reached an inflection point. E-commerce take-rates have climbed to 13.3%, with advertising contributing 9.6%, with SEA marketplace advertising still sitting at just 2% of GMV versus 7% in the U.S., suggesting a substantial runway ahead. Digital financial services EBIT margins hit 28% as the loan book scales prudently while keeping NPLs low. Gaming, while past its 2021 peak, is showing signs of recovering and continues to generate the cash that funded Shopee's expansion, which now funds Monee's growth.

At current prices, following a 40% drawdown from October 2025 highs, assuming 20% revenue growth into 22% FCF margins, our base case projects attractive ~27% annualized returns over the next five years. Sea Limited deeply aligns with Vision Investing: democratizing commerce and finance for 600+ million people across SEA, Taiwan, and Brazil, while reinvesting and compounding capital at attractive rates.

Shopee continues to grow its e-commerce GMV and revenues strongly.
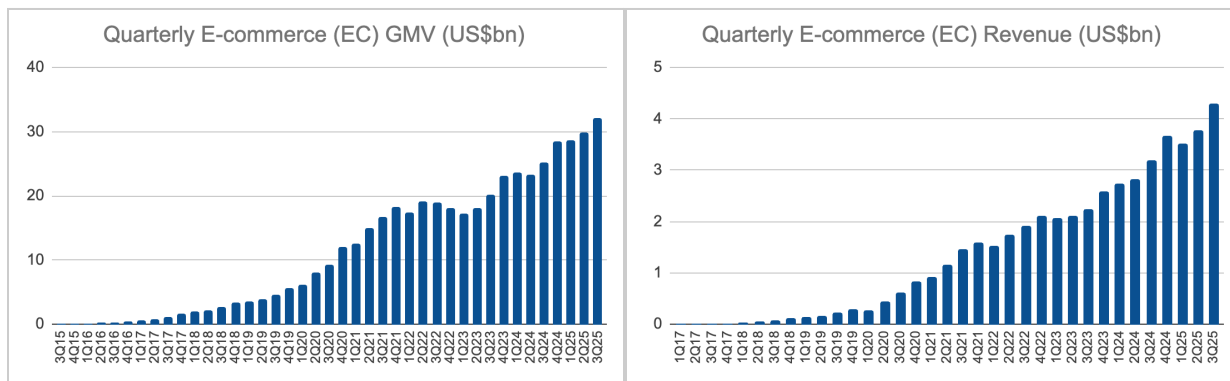
Shopee is the largest e-commerce platform (52% GMV share) in SEA.

Source: Momentum Works Ecommerce in Southeast Asia (various reports)
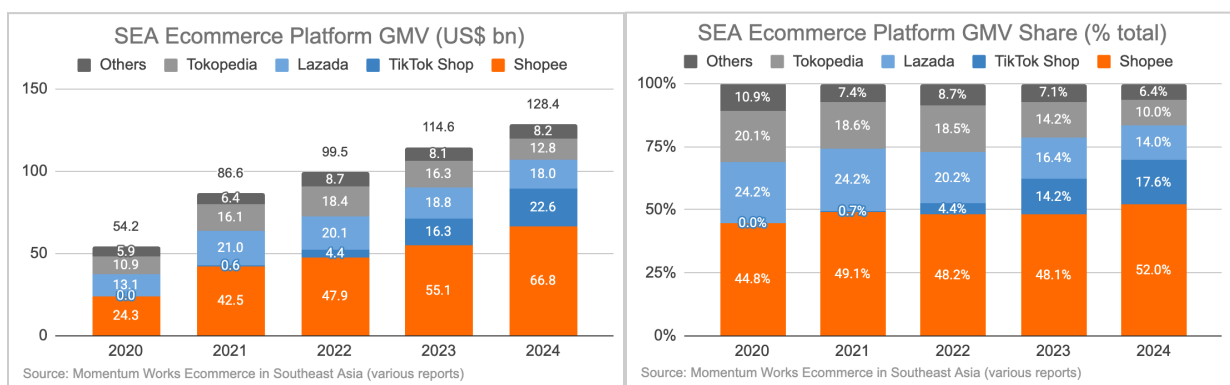
SPX Logistics has become the 2nd-largest SEA logistics provider after J&T Express.

Source: Momentum Works Ecommerce in Southeast Asia (various reports)
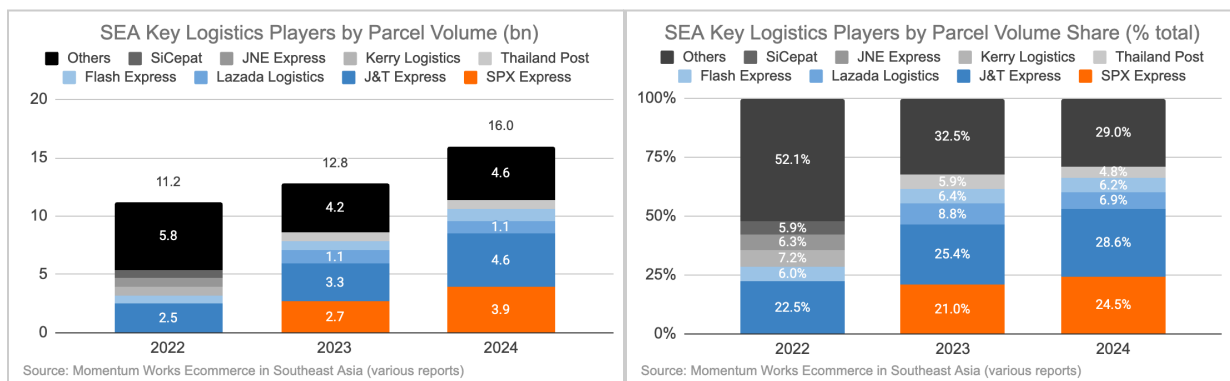
ShopeeFood is growing rapidly, particularly in Vietnam, Indonesia, Malaysia, and Thailand, and is the 3rd-largest in SEA. Food delivery riders could provide opportunities for instant (<2h) commerce, leveraging Shopee (similar to China's Meituan, Taobao, JD).

Southeast Asia Food Delivery Platform GMV (US$bn)

Source: Momentum Works Food delivery platforms in Southeast Asia (various reports)

Southeast Asia Food Delivery Platform GMV (% total)

Source: Momentum Works Food delivery platforms in Southeast Asia (various reports)
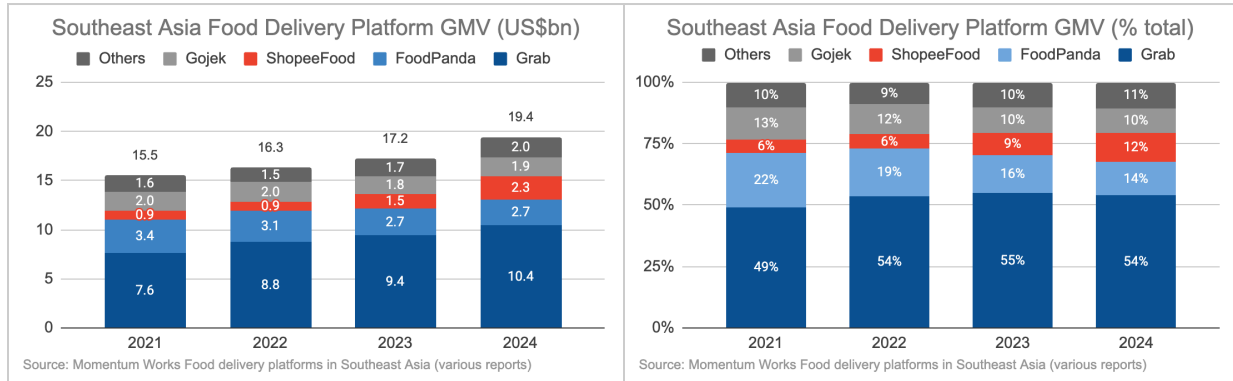
Shopee's overall e-commerce take-rates continue to grow, supported by faster-rising core marketplace (3P)/advertising take rates sufficiently offsetting the slower-declining VAS marketplace (3P)/logistics take rates.

Annual EC Take-Rate Breakdown by Type (%)

Quarterly EC Take-Rate Breakdown by Type (%)
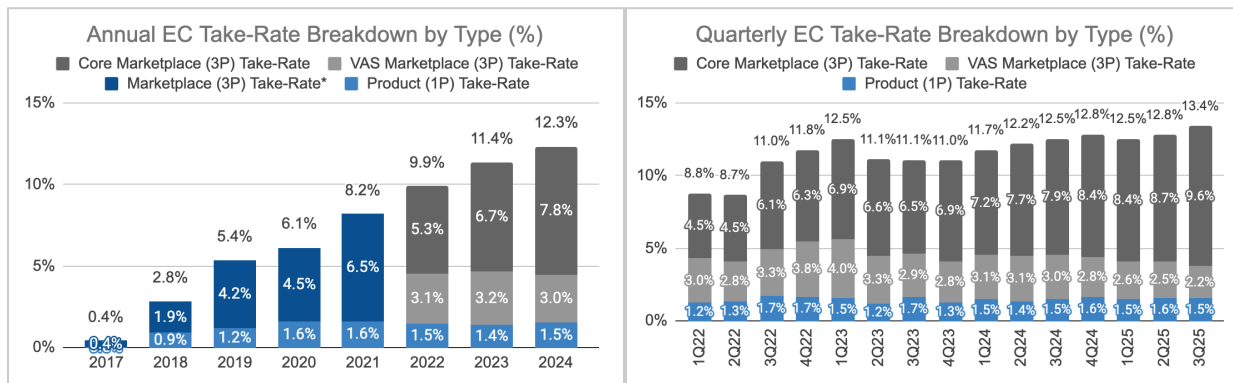
Shopee's profitability continues to improve, and is now Adj EBITDA/EBIT profitable.

Annual E-commerce (EC) Adj EBITDA Margin (%)

Annual E-commerce (EC) EBIT Margin (%)
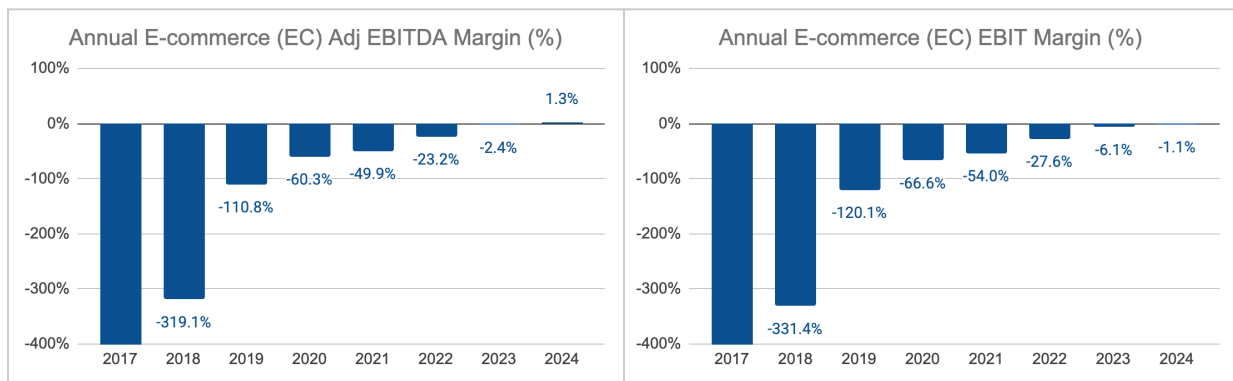
Monee's Loan Book continues to grow, largely driven by on-book (87%) versus off-book (13%), and most (94%) of the loan book is to consumers (>90%) and commercial (SMEs).

Quarterly On-Book and Off-Books Loans by Type (US$bn)

Sea Limited On-Book Loans by Type (% total gross)
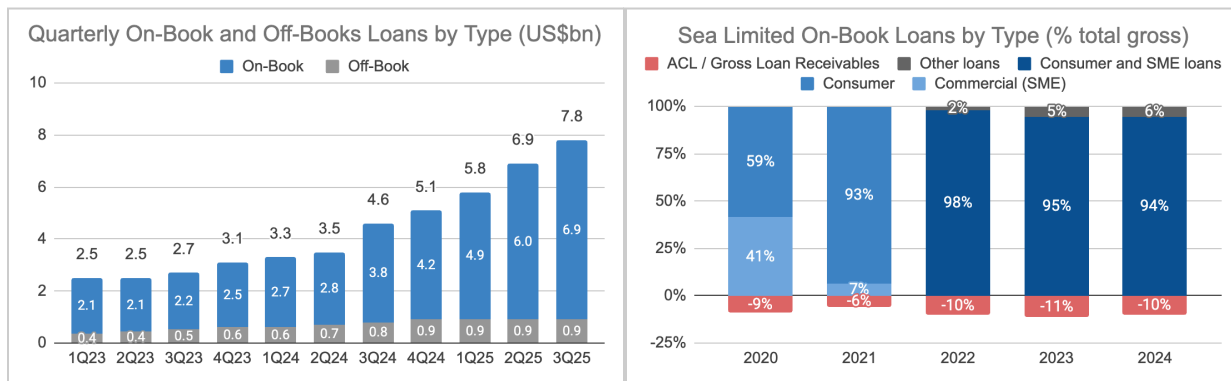
Monee's loan 90+ NPLs continue to remain low and stable, below 2%.

On-Book Consumer & SME Loans (% total gross)

Quarterly NPL 90+ of Total On-Book & Off-Book Loans
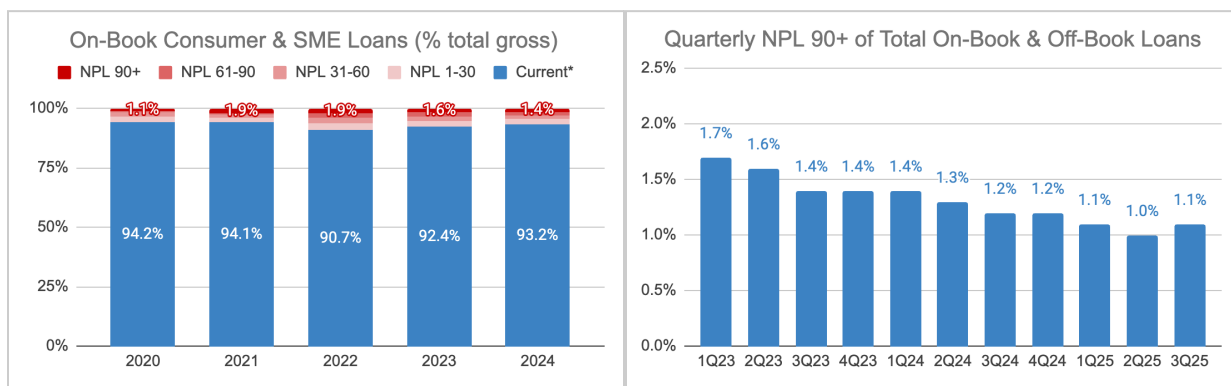
In digital entertainment, Garena's active, paying users have begun to return.

Quarterly DE Quarterly Active Users (QAU) (mil)

Quarterly DE Quarterly Paying Users (QPU) (mil)
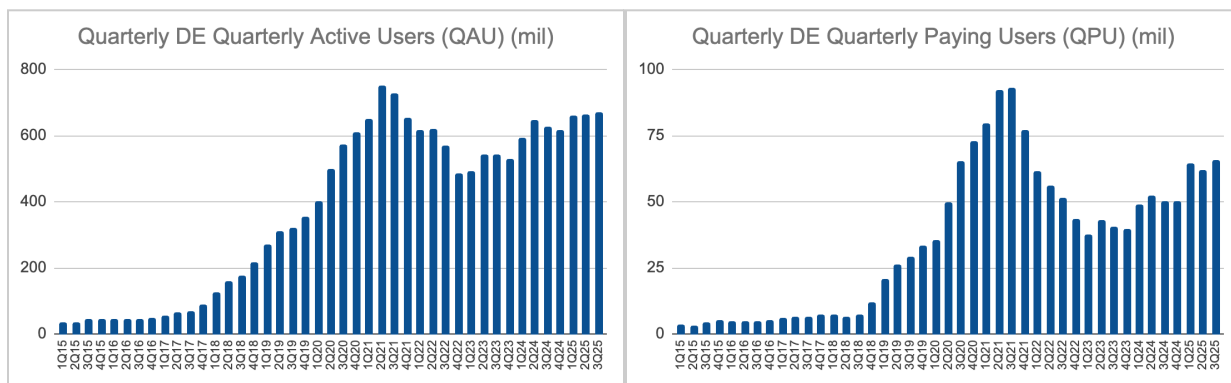
Garena's bookings and revenues are recovering from their 2023/2024 lows.

Quarterly Digital Entertainment (DE) Bookings (US$bn)

Quarterly Digital Entertainment (DE) Revenue (US$bn)
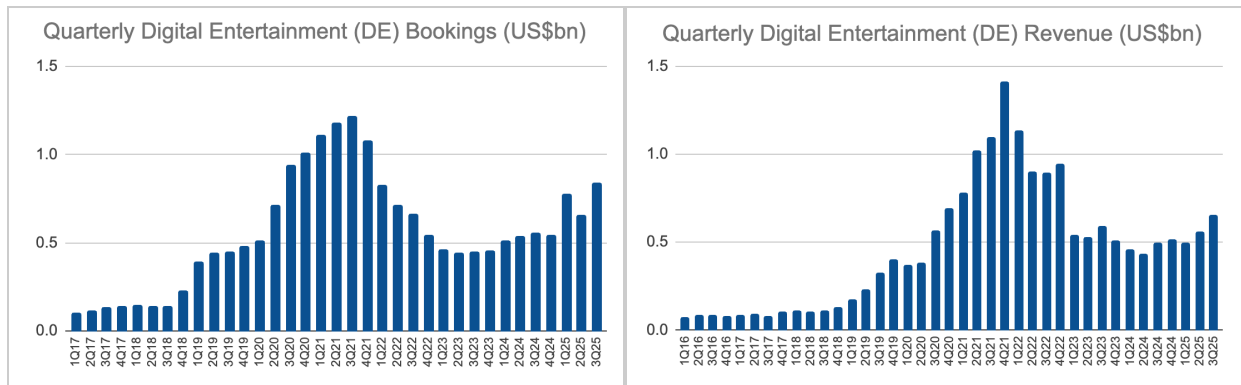
Sea Limited started with digital gaming (Garena), subsequently added e-commerce (Shopee), and is increasingly expanding into digital financial services (lending via Monee).

Quarterly Revenue by Core Segments (US$bn)
- Other Services
- Digital Entertainment (DE)
- Digital Financial Services (DFS)
- E-Commerce (EC)

Quarterly Revenue by Core Segments (% total)
- Other Services
- Digital Entertainment (DE)
- Digital Financial Services (DFS)
- E-Commerce (EC)
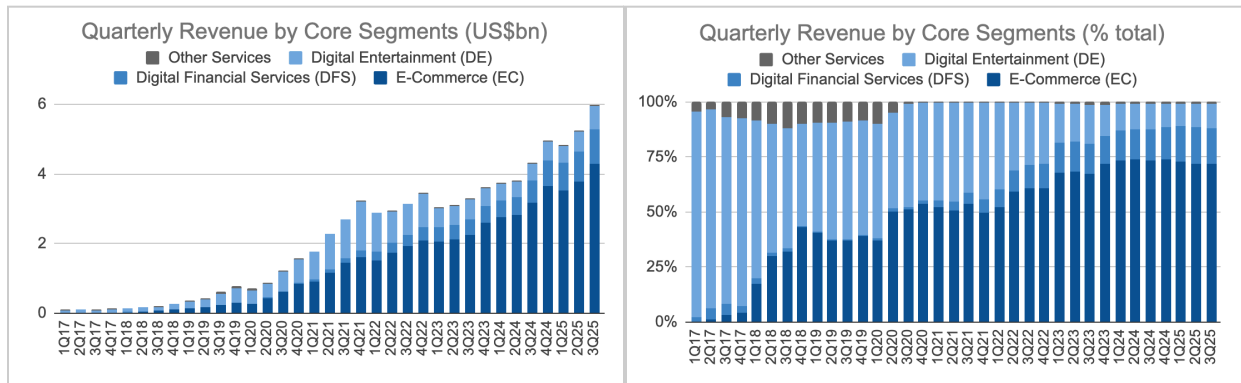
Sea Limited has been reaccelerating revenue and gross profit growth since 2024.

Quarterly Revenue by Segment/Type (US$bn)
- Sales of Goods
- Services (DE)
- Services (DFS)
- Other Services
- Services (EC)
- Services (EC & Other Services)

Quarterly Gross Profit by Segment/Type (US$bn)
- Sales of Goods
- Services (DE)
- Services (DFS)
- Other Services
- Services (EC)
- Services (EC & Other Services)
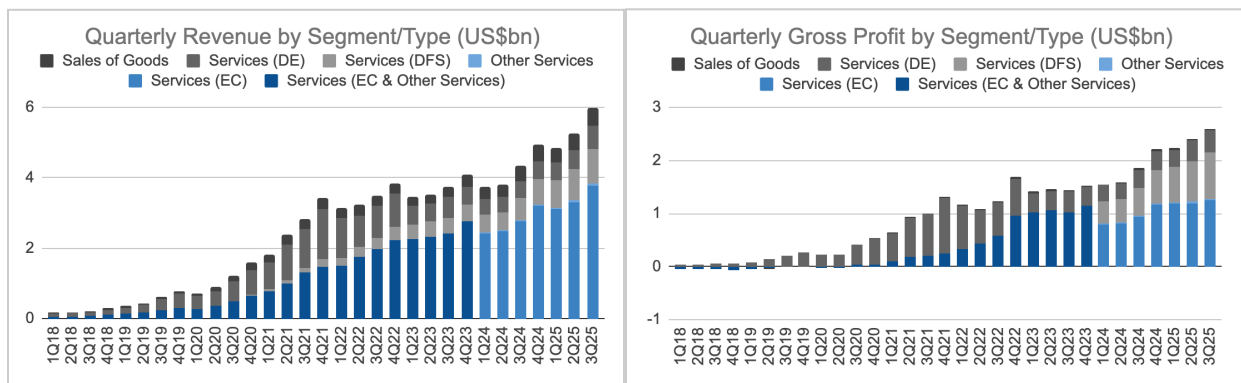
Sea Limited's profits continue to grow, driven by improving EC and DFS profitability, while DE profitability continues to hold strong.

Quarterly Adj EBITDA by Core Segments (US$bn)

Quarterly EBIT by Core Segments (US$bn)
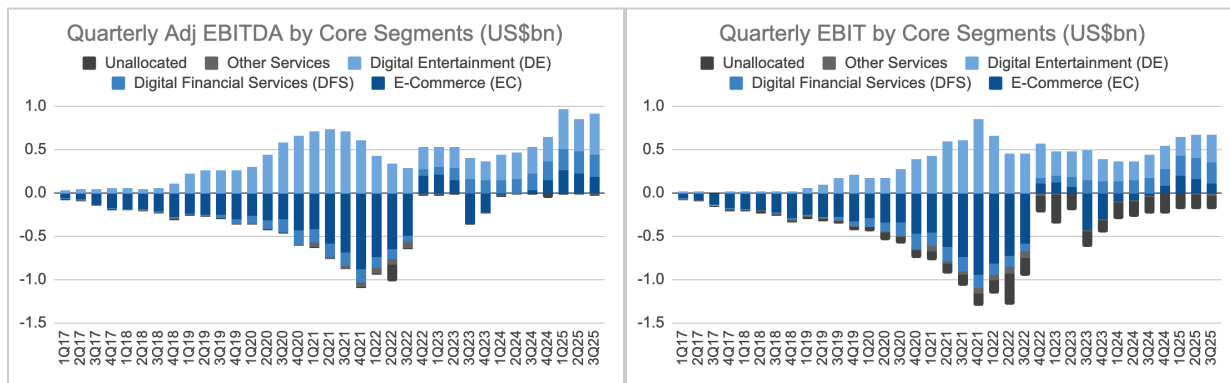
Profitability for Garena remains strong, while Monee has become profitable rapidly. Key to Shopee becoming more profitable in the coming years.

Quarterly Adj EBITDA Margins by Core Segments (%)

Quarterly EBIT Margins by Core Segments (%)
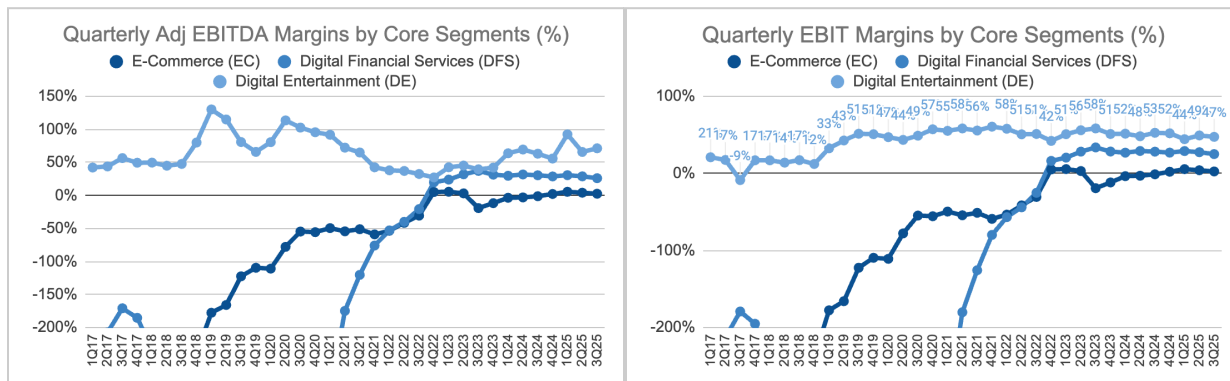
Unknowable versus Knowable

It is the time of the year again when many eloquent experts, from economists to strategists, op-ed columnists, speculators, and analysts, will have an opinion about what will happen in 2026. You might be tempted to ask us how we think 2026 will do. Those who know us well will know we will make no attempt to predict what will happen in any given year, for it is simply wasted energy and futility, for it is truly unknowable. We simply have no idea what will play out.

It is far better to have and manage our expectations about what could happen, covering a wide range of possibilities, especially the worst, rather than the best. Being directionally and approximately correct rather than being precisely wrong is a far better, longer-term approach. The chart of S&P 500 1Y earnings squiggles over the last 30 years below says it best. Many will forecast incorrectly, but it is the long-term direction that matters.

S&P Operating EPS Forecast Squiggles versus Actuals over the last 30 Years
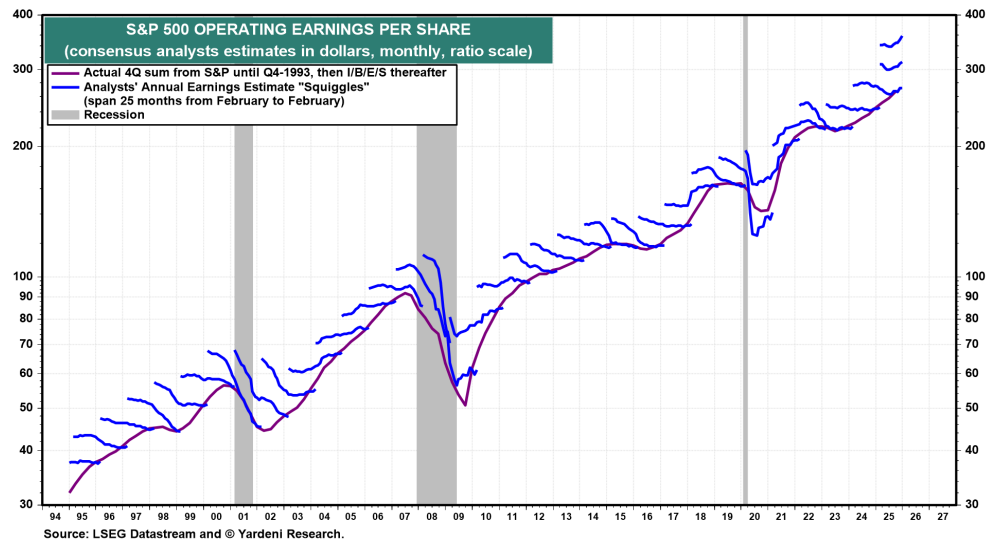
Source: Yardeni Research

Observations about the World - US versus China in Electricity and AI

US electricity production has barely grown over the last 15-20+ years. In contrast, China has been growing rapidly in renewables (hydro, wind, solar) and nuclear power, and is >2X the US.
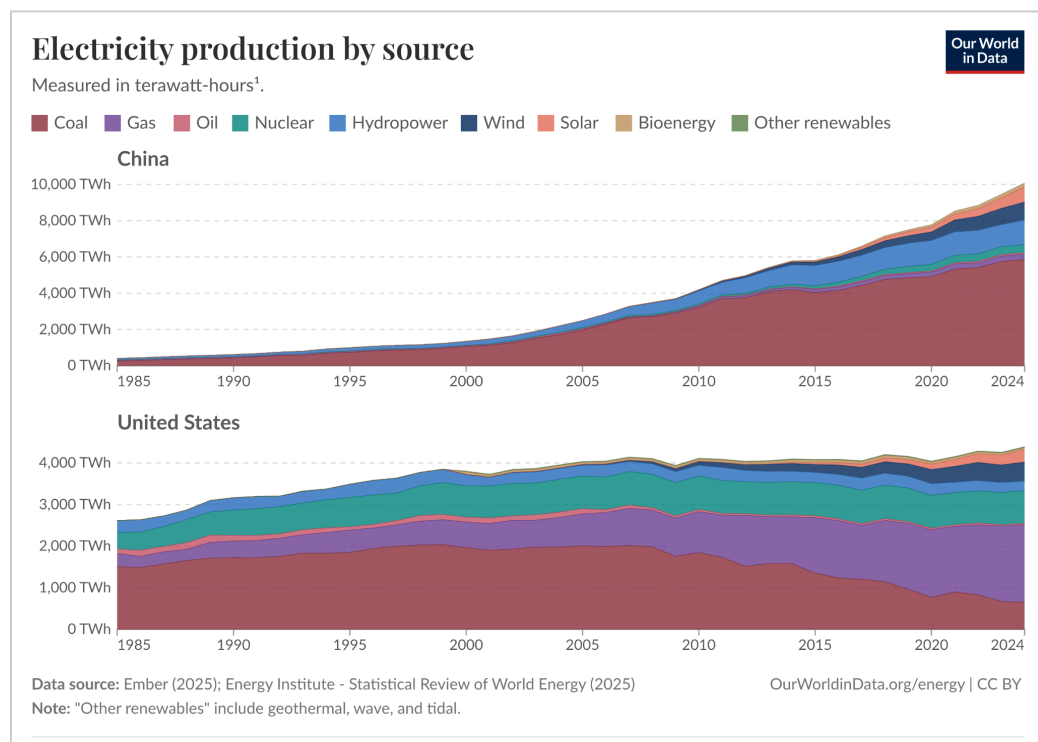
Source: Our World in Data

On the hardware side, US AI chips continue to be much more powerful and energy-efficient than the Chinese chips, while on the software side, Chinese AI LLMs are more cost- and compute-efficient to train/run than US LLMs due to optimizations like MoE architectures and inference scaling, delivering strong performance at lower resource use. Efficiency gains have come from architectural innovations and training on constrained hardware than having more superior hardware. Overall, China's open-source LLMs have been rapidly catching up with US closed-sourced LLMs.

NVIDIA's versus Huawei's AI Chips in terms of processing power & memory bandwidth
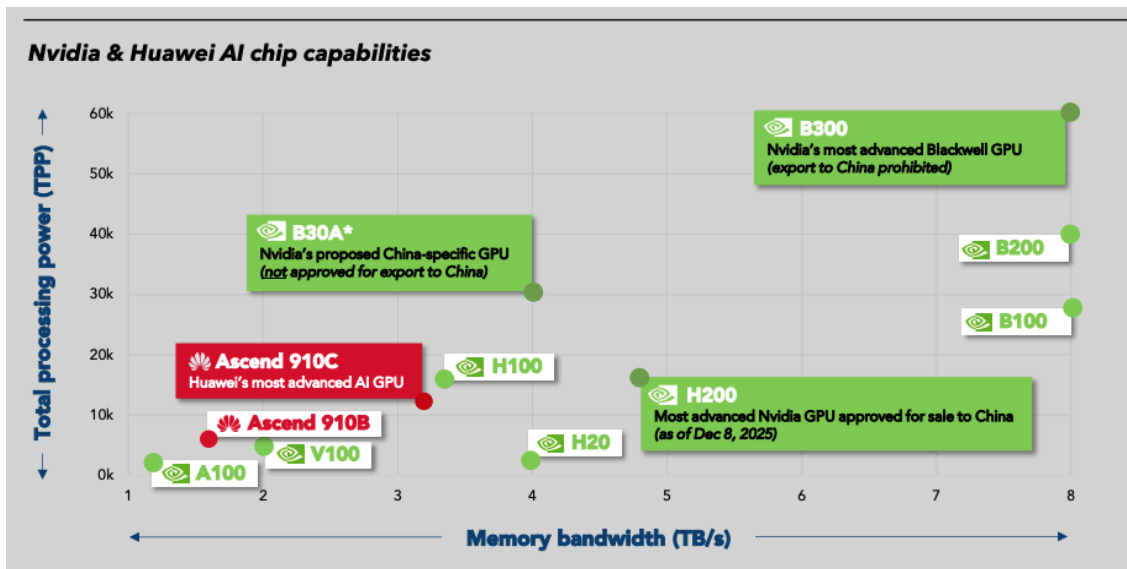
Source: MUFG Americas Research

US vs Chinese vs Others LLMs across different modalities

Source: Artificial Analysis (9 Jan 2026)

US closed-source LLMs (OpenAI, Anthropic, and Google) continue to lead, but the Chinese open-source LLMs (Z AI, DeepSeek, MiniMax, and Xiaomi) are rapidly catching up, closing the gap
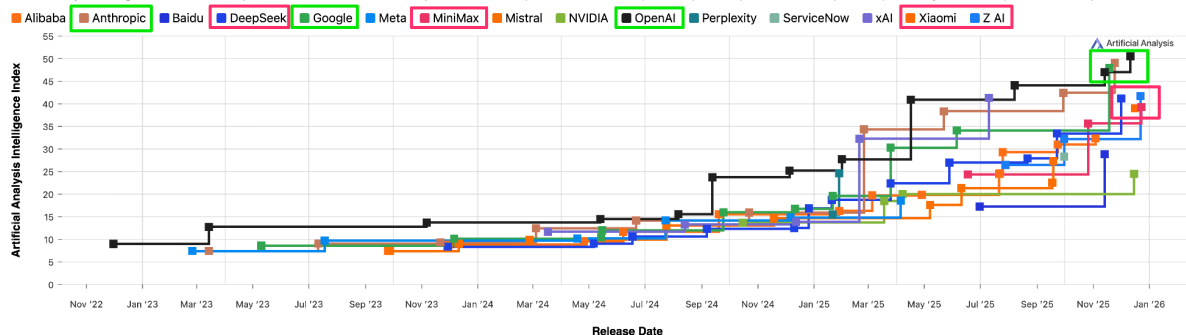
Source: Artificial Analysis (9 Jan 2026)

Chinese LLMs (e.g., Z AI's GLM, DeepSeek, and Xiaomi's MiMo) are less intelligent, slower, but are also much cheaper to operate.
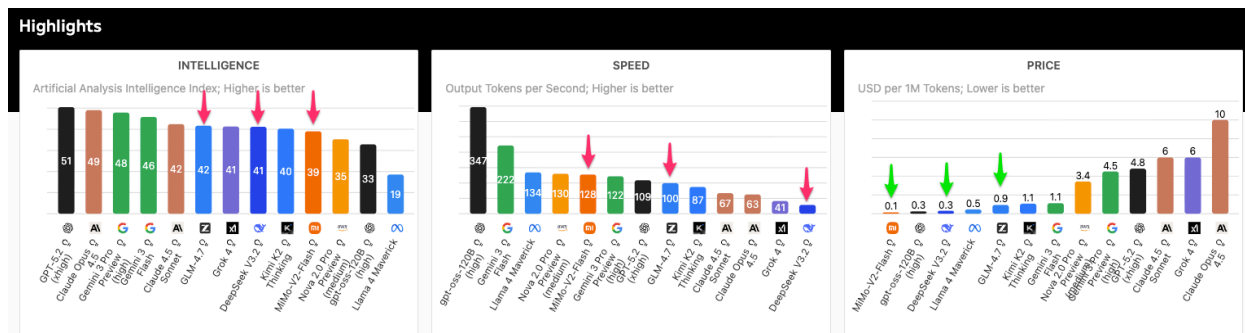
Source: Artificial Analysis (9 Jan 2026)

Observations about the World - Manufacturing becomes critical again

The US leads in capital-light innovation. US dominance in goods is limited to aerospace (Boeing/Lockheed), reusable rockets (SpaceX), defense, medical devices, and advanced industrial machinery. But the US is much stronger in services: fabless semiconductor design (NVIDIA, AMD, Broadcom, Qualcomm), software/cloud (AWS, Azure, GCP), biotechnology, and financing.

China leads in capital-intensive production, with ~64% global share of global Electric Vehicles (EVs) sales, ~70% share of EV batteries, and ~70% share of global drones. China is the global manufacturer of the most important products that require integrated supply chains at

scale. While everyone else largely supplies China with commodities and critical inputs, and then consumes China's output. BYD and Xiaomi are great examples.

China's focus on manufacturing excellence is exceptional, making everything from electronics and EVs to solar, drones, and robotics. China's industrial success is due to its strong infrastructure (ports, rail, data connectivity, electrification, process knowledge) and robust manufacturing ecosystem of self-reinforcing parts enabling unmatched scale and speed.

America invents, China builds, everyone else supplies, and Europe regulates. It will become clearer that whoever establishes the manufacturing strength will come to dominate.

Strong convictions, loosely held. Above are our current observations/opinions of selected trends, and they continue to evolve as things change. As things change, we will change our minds and not be stubborn about it.

Final Words

One can tell a great deal about a person by the books they read. Similarly, one can learn a great deal about an investor from the stocks they own and how long they have owned them. **Books reveal your interests, words distill your thoughts, stocks reveal your actions, and holding periods reveal the investor.**

We share the holdings of Vision Capital Fund and their breakdown because we want you to see that we are consistent in our approach, both in how we think, communicate, and act, and that we continue to select the best companies that we believe are likely to do well the right way.

Overall, we remain optimistic about the long-term returns of stocks and our holdings in Vision Capital Fund. The long-term business returns of the underlying businesses (i.e., growth of revenues, profits, and free cash flows), in which we are part owners, are expected to continue to drive the majority of our long-term returns, than changes in valuation multiples.

Our holdings are faster-growing and more profitable, with **25.9% TTM revenue growth** and **25.8% TTM FCF margins** on a weighted-average basis (see table on next page).

No.	Holding	Percentage Allocation	1Y TTM Rev Growth	1Y Fwd Rev Growth	TTM FCF Margin	Secular Theme
1	Palantir Technologies Inc	12.4%	47.2%	53.8%	46.0%	Software (Data Integration)
2	Shopify Inc	6.9%	30.3%	29.1%	17.8%	Ecommerce/Payments
3	Nu Holdings Ltd	6.0%	20.4%	35.9%	17.1%	Payments/Digital Banking
4	Taiwan Semiconductor Manu	5.9%	37.0%	36.5%	24.6%	Semiconductors, AI
5	NVIDIA Corp	4.9%	65.2%	63.5%	41.3%	Accelerated Computing/AI
6	Meta Platforms Inc	4.9%	21.3%	21.3%	23.7%	Digital Advertising
7	Wise Plc	4.4%	9.9%	17.4%	21.1%	Cross-Border Payments
8	MercadoLibre Inc	4.3%	36.9%	37.6%	32.9%	Ecommerce/Payments
9	Crowdstrike Holdings Inc	4.1%	22.1%	21.5%	25.7%	Software (Cybersecurity)
10	Tesla Inc	3.8%	-1.6%	-3.3%	7.2%	Electric Vehicles/Solar/Robotics
11	Amazon.Com Inc	3.6%	11.5%	12.0%	1.5%	Ecommerce, Cloud, Advertising
12	Pro Medicus Ltd	3.3%	31.9%	36.4%	52.1%	Software (Medical Imaging)
13	ServiceNow Inc	3.2%	21.1%	20.6%	31.3%	Software (Workflow Automation)
14	Adyen Nv	3.2%	21.0%	34.7%	43.1%	Payments
15	Lululemon Athletica Inc	2.8%	8.8%	4.2%	10.2%	Athleisure
16	Meituan	2.5%	12.3%	13.6%	1.2%	Ecommerce (Services)
17	Spotify Technology Sa	2.4%	11.9%	23.7%	17.3%	Streaming (Audio)
18	Trade Desk Inc	2.2%	20.8%	18.2%	24.7%	Digital Advertising
19	Jd.com Inc	2.1%	16.6%	18.8%	3.8%	Ecommerce
20	Tencent Holdings Ltd	2.0%	13.5%	19.0%	28.5%	Digital Advertising/Gaming/Payments
21	Netflix Inc	2.0%	15.4%	15.6%	20.7%	Streaming (Videos)
22	Zscaler Inc	2.0%	23.2%	23.3%	32.6%	Software (Cloud Security)
23	Sea Ltd	2.0%	35.8%	32.8%	18.5%	Ecommerce/Payments/Gaming
24	Cloudflare Inc	1.9%	28.1%	28.5%	11.7%	Software (Connectivity & Security)
25	Microsoft Corp	1.8%	15.6%	16.0%	26.6%	Software, Cloud
26	MasterCard Inc	1.8%	15.6%	16.1%	54.0%	Payments
27	Paypal Holdings Inc	1.1%	4.5%	4.7%	16.9%	Payments
28	Paycom Software Inc	1.0%	9.7%	8.9%	19.7%	Software (Human Capital Mgmt)
	Cash	1.6%				
	Weighted-Average		**25.9%**	**28.6%**	**25.8%**	

Source: Koyfin and Vision Capital Fund (as of 7 Jan 2026)

Thank you for your continued trust and support in allowing us to invest for the long term. I look forward to navigating the coming years and decades ahead together.

Please feel free to reach out if you have any questions about Vision Capital Fund, would like to add more capital, or refer another investor you think might be interested and suitable. Your recommendations are always well appreciated.

Excelsior (Latin for "ever upward" or "higher"),

Eugene Ng | eugene.ng@visioncapitalfund.co

Managing Partner, Vision Capital Fund

Disclaimer

The Information published herein is intended for "Accredited Investors" and/or "Institutional Investors" only as defined in the Singapore Securities & Futures Act (Cap. 289) of Singapore ("SFA"). This Information is provided for informational and discussion purposes only and is not, and may not be relied on in any manner (legal, tax or investment advice) as an offer to sell or a solicitation of an offer to buy or subscribe to any funds managed by Galilee Investment Management Pte. Ltd. ("Galilee"). An offering of interests in the Fund will only be made under a confidential offering memorandum or similar written material and the Fund's subscription documents (collectively referred to as the "Material"), which will be furnished to accredited or institutional investors (and their employees and agents) on a confidential basis at their request for their consideration in connection with such offering. None of the information or analyses presented is intended to form the basis for any investment decision, and no specific recommendations are intended. No reliance may be placed for any purpose on the Information provided or the accuracy or completeness thereof, and no responsibility can be accepted by Galilee, and/or any of their respective affiliated entities to anyone for any action taken on the basis of such information. Whilst Galilee shall use reasonable efforts to obtain information from sources which we believe to be reliable and up to date, Galilee gives no warranty as to the accuracy, completeness or reliability of any information, opinions or forecasts contained in the Information. No responsibility or liability can be accepted for any errors or omissions or for any loss resulting from the use of the Information. Past performance of the managers and the funds, and any forecasts on the economy, stock or bond market, or economic trends that are targeted by the funds, are not indicative of future performance. Investment in the Fund will involve significant risks, including loss of the entire investment. The Fund will be illiquid, as there is no secondary market for interests in the Fund and none is expected to develop. There will be restrictions on transferring interests in the Fund. Investments may be leveraged and the investment performance may be volatile. Whilst Galilee shall use reasonable efforts to obtain information from sources which we believe to be reliable and up to date, Galilee gives no warranty as to the accuracy, completeness or reliability of any information, opinions or forecasts contained in the Information. No responsibility or liability can be accepted for any errors or omissions or for any loss resulting from the use of the Information. Galilee may update, revise, delete or modify the content and information herein without notice. The material should only be considered current as of the time of initial publication or as otherwise stated in the Material without regard to the date on which you may access the Material. These exclusions of liability do not apply to the extent that such exclusions are invalid or ineffective under any law or regulation applicable to Galilee. Before deciding to invest in the Fund, prospective investors should read the Material and pay particular attention to the risk factors contained in the Material. Investors should have the financial ability and willingness to accept the risk characteristics of the Fund's investments, including any risk factor, forward looking statements as set out in the Material. Holdings are subject to change at any time.

No Commercial Exploitation: The copyright and other intellectual property rights in the Information are owned by Galilee. Any use of the Material for any purpose is accordingly prohibited except as stated below. You may not reproduce, transmit, modify, store, archive or in any other way use for any public or commercial purpose any of the Information without the prior written permission of Galilee.