

July 9, 2026


Dear Partners,

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

We are still in our very early innings, having only started just under two years ago. We request that our long-term performance be evaluated over a timeframe of three to five years or more.

Performance Overview.

Vision Capital Fund's unaudited gross and net returns for Q2 2026 were +6.2% and +5.7%, respectively, versus +15.2% for the S&P 500. The cumulative gross and net returns from inception (1 Oct 2024) were +12.2% and +6.1%, respectively¹, versus +33.0% for the S&P 500. The annualized gross and net returns are +6.8% p.a. and +3.5% p.a., respectively, versus +17.8% 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%
2026 YTD	-11.7%	-12.6%	10.2%	-22.8%
Inception	12.2%	6.1%	33.0%	-26.9%
Annualized (p.a.)	6.8%	3.5%	17.8%	-14.3%

Source: Vision Capital Fund, Net (Class A Shares after fees) as of 30 Jun 2026, since inception 01 Oct 2024

¹ The performance results for Q2 2026 (period ending 30 Jun 2026) represent Class A unaudited returns and may be subject to further adjustment. Net returns throughout are presented net of management fees, accrued performance allocation, and all other fund-level expenses. Gross returns throughout do not reflect the deduction of management fees or accrued performance allocations but are presented net of all other fund-level expenses. Each individual investor's actual performance may differ materially depending on factors such as the timing of the investment and the fee structure. Performance figures are not based on the performance of a specific investor's account. Instead, they are based on the performance of a hypothetical investor's account that invested at the inception of each Class, has made no subsequent contributions or redemptions, and remains invested. An investor could lose all or a substantial portion of its investment. Past performance does not guarantee future returns.

Fund Commentary.

The **top 5 QTD detractors** were: Palantir, Netflix, Lululemon, Meituan, and Nu. The **top 5 QTD contributors** were: CrowdStrike, TSMC, Pro Medicus, Wise, and NVIDIA.

It is disappointing that we have underperformed, particularly over the last 9 months, which dragged our overall cumulative performance. Short-term stock price movements of our holdings are beyond our control, and we prefer to focus on fundamentals rather than on short-term themes that involve constant sector and/or company rotation. While we have partially recovered from the Mar 2026 bottom, we have underperformed the recovery thus far.

This year, if one did not own memory or AI semiconductor supply chain stocks (e.g., top 10 2026 [S&P 500 YTD winners](#): SanDisk, Micron, Intel, Western Digital, Seagate, Marvell, Dell, Corning, Applied Materials, AMD, etc.), one would have underperformed the S&P 500. Unsurprisingly, we too have underperformed because we own none of these. While our underperformance could persist in the near term, we think it is more temporary and unlikely to be structural.

We had strong inflows in early Q2 from our largest investors (us included) seeking to take advantage of buying opportunities. We deployed some of that new capital to 10 of our existing positions (in alphabetical order): Cloudflare (NET), MercadoLibre (MELI), Meta (META), Nu (NU), Palantir (PLTR), Pro Medicus (PME), Sea Limited (SE), ServiceNow (NOW), Shopify (SHOP), and Zscaler (ZS). These companies continued to grow strongly and became more attractively priced with the larger price drawdowns. We exited 2 of our entire positions in Lululemon (LULU) and Paycom Software (PAYC) in Q2. We now have 25 holdings.

Not choosing Maxxing.

Maxxing is internet slang for “max,” short for “maximizing” a certain aspect of one’s life to the fullest. It has numerous variants over the years, evolving from “looks-maxxing” to “sleep-maxxing”, “money-maxxing”, and, more recently, to “token-maxxing” in AI.

Token-maxxing is the practice of maximizing AI token consumption, on the theory that higher AI token usage equals greater productivity. With scoreboards, employees are socially pressured to consume as many AI tokens as possible by running multiple AI coding agents simultaneously, using expansive context windows, and automating a wide range of workflows.

Incentives drive outcomes. While these were initially used to rapidly drive behavioral and adoption change, such scoreboards eventually backfired when they drove employees to quickly game the system, generating unnecessary AI workloads to boost their visible statistics, resulting in AI budget blowouts rather than meaningful business outcomes.

Because raw token consumption is a weak proxy for productivity, many companies are shifting away from “token-maxxing” to “value-maxxing” or “outcome-maxxing”, focusing on cost efficiency, right-sizing AI token consumption for positive and valuable business outcomes.

The pattern is always the same. Maxxing begins with a lack, overshoots abundance, and turns the excess into a new lack. Maxxing is effectively maximizing without constraints and with unlimited freedom (and resources), rather than optimizing within given limitations.



Overfilled and Filled.

Similarly, the equivalent of “token-maxxing” would be “return-maxxing” in investing and “profit-maxxing” in business. Both of which we abhor and avoid doing or investing in companies that do. Let’s elaborate on this further.

Avoiding Return-Maxxing and Bottleneck “Investing”.

“Return-maxxing” is the practice of seeking to maximize investment returns, often by seeking higher returns by chasing rapidly rising price momentum, companies in hypergrowth, constantly moving on to the next newest shiny thing, trading (buy low, sell high), and/or employing leverage through margin/borrowing and derivatives such as buying options (to boost upside) and selling options (to boost returns).

“Bottleneck-investing” has become the latest investing game, which we have intentionally avoided. Many “investors” have done well by employing “bottleneck-investing” across the AI stack, from AI chip designers (NVIDIA, AMD, Broadcom, Marvell, MediaTek, etc) to semiconductor foundry/advanced packaging (TSMC) to memory (SK Hynix, Samsung, Micron, Sandisk, Western Digital, Seagate, Kioxia, etc) to fiber optics (Corning, Lumentum, Coherent, etc) to power and grid (GE Vernova, Vistra, Constellation Energy, Vertiv, etc).

The primary issue with bottlenecks/constraints is threefold. The three deepest lessons from The Theory of Constraints (TOC) are that: (1) a system has exactly one constraint at a time, (2) the system can only move as fast as the bottleneck, and (3) bottlenecks never disappear, they keep migrating. Fix one choke point, and the constraint moves to the next slowest possible step. Every system has a bottleneck, only its location changes.



The Narrowing.


“Bottleneck investors” are short-term renters, not long-term owners. The bottleneck investor merely rents (for months, maybe 1-2+ years), and the lease is often much shorter than he thinks. The moment one bottleneck starts to clear, and another appears, the bottleneck investor keeps rotating and chasing the next shiny little thing.

We far prefer to own Apple than the players in Apple’s supply chain (or such equivalents).

Apple does not manufacture any of its hardware products (e.g., iPhones, iPads, Macbooks, etc). Instead, it designs them and gets contract manufacturers to produce them.

Apple operates with greater bargaining power, as its suppliers provide negative working capital, and it has higher operating leverage and profitability. In the modern global economy, supply chain players tend to be more oligopolistic than monopolistic. Pure monopolies/duopolies in supply chains are exceptionally rare. In semiconductors, there are TSMC in advanced foundry, ASML in advanced EUV lithography, and Synopsys and Cadence in EDA design.

While it is certainly possible for Apple’s supply chain to ride Apple’s success and do well, it usually does less well. The two 20-year charts below (vs Apple’s first iPhone in 2007) illustrate this best. On a stock-price returns-alone basis, Apple returned 154x and 28.7% CAGR, significantly exceeding all 27 supply chain players. Notably, some of the top performers evolved to do more than just supply Apple. They outgrew their largest customer, stopped being suppliers, and became platforms in their own right.




Hopefully, you should know by now that we do our best to avoid “return-maxxing” and that you are invested with us, knowing the type of game we prefer to play. We own high-quality, durable, secular growers. Profitable businesses that sell products and services that people value. Guided by win-win-win thinking. Top dogs in winner-take-all markets.

It is why we own NVIDIA and TSMC thus far, and not memory stocks or the AI supply chain. We do not chase the highest returns and ignore the downside. We do not pretend to spot every

bottleneck and trade it well. That strategy simply costs too much. It requires excellent market timing of both to buy low and sell high, which we are poor at. We hope you will come to increasingly appreciate this over time.

Prefer not to own “profit-maxxing” businesses.

One “profit-maxxing” business that comes to mind is Broadcom (AVGO) with its AI compute and networking growth runway, supported by accelerating inference demand, expanding strategic customer engagements, and its strong technology leadership across ASIC design, SerDes, packaging, and Ethernet switching.



Broadcom is a technology giant that built its “indispensable technological leadership” on numerous serial acquisitions, from hardware and semiconductors to software and infrastructure SaaS. CEO Hock Tan’s repeated M&A playbook to strengthen technological competencies: cleaning up acquisitions, shutting down unprofitable operations, cutting costs, firing employees, raising prices, prioritizing margin over innovation, making each more profitable than before, and reinvesting the proceeds to acquire new, larger companies.

We were considering increasing our AI semiconductor exposure beyond GPUs to include Broadcom's custom ASICs. Broadcom is the #2 global AI semiconductor supplier (after NVIDIA), #1 custom chip ASIC supplier, and #1 networking semiconductor supplier. With its rapidly growing XPU and networking business, Broadcom effectively owns the “crown jewels” of the networking and AI infrastructure. While we were initially attracted by its “strong” financials and strong AI Semiconductor growth stemming from XPUs, we were unattracted by its slower organic growth in its non-AI semiconductor and infrastructure software business.

Philosophically, we found ourselves fundamentally at odds with how they operated. Many of its customers valued it for its unparalleled technical execution, supply chain reliability, and “gold standard” performance. The higher premium pricing customers pay is viewed as a

necessary investment for mission-critical infrastructure where failure is not an option. While we prefer businesses with “scaled economies shared”, Broadcom is the equivalent opposite with “monopolistic margin extraction”.

Based on numerous expert interviews, Broadcom has a long-standing reputation for aggressive business practices, including drastic price increases (e.g., with VMware), deteriorating user experiences, and forced transitions that prioritize its bottom line over its customers’ workflows or budgets. Customers rapidly become frustrated and seek to leave at the first opportunity.


Such a win-lose-lose (profit>customers>people) philosophy is not long-term sustainable compared to a win-win-win one, which we very much prefer. Unsurprisingly, when Google diversified beyond Broadcom for its TPUs to add MediaTek (and possibly Marvell), we were not surprised. Broadcom could easily continue to do very well as a key AI beneficiary and leader in custom ASIC designs (Google TPU, Meta MTIA, OpenAI, etc), versus Marvell (Amazon Trainium, Microsoft Maia), and MediaTek (Google TPU).

However, Broadcom’s philosophy of “profit-maxxing” does not align with ours. We prefer businesses to generate growing value for their customers and then gradually increase prices to monetize a portion of the higher incremental value they have created. **There are many ways to make money. It is about making money the right way, that is sustainable, win-win-win for a long period of time.** We just have to work harder to make our money elsewhere.

Scarcity versus Abundance and Constraints versus Freedom.

Maxxing seeks to maximize without constraints, rather than optimize with constraints. Being trained in neoclassical economics, we far prefer the much more realistic mental model of constrained optimization given limited budgets and constraints (which can change over time), rather than the unrealistic one of maximization with unlimited resources.

This brings us to two mental models of scarcity versus abundance and constraints versus freedom, which we recently wrote about [here](#). Scarcity versus abundance is a mental model that we regularly employ to discern trends from fads and avoid investing in hype or fad stocks.




When it comes to a business's products and services, there must be a dynamic, beautiful balance between scarcity and abundance. Abundance (of supply) can rapidly devalue incremental satisfaction and reduce demand. Too much of something, and it quickly becomes less desirable. This is generally evident in most consumer goods.



Scarcity and abundance.

A recent example would be Pop Mart's meteoric rise with Labubu/The Monsters, where a monster doll went from a niche collectible to a global obsession overnight. Celebrities started carrying them, social media went viral, and peak consumer mania followed as hype demand overwhelmed the initially limited supply. Pop Mart subsequently capitalized on inflated resale prices, rapidly ramped up production volumes, and increased customer availability.

Revenues for The Monsters exploded 57X to RMB 9.3bn over 2 years, from 1H23 to 2H25. With massive restocks, panic-stricken scalpers sold. As resale markets were flooded, consumer interest cooled rapidly, and resale prices collapsed. At the peak of the Labubu mania, the Labubu hype could follow similar cycles to NFTs in 2021 or to sneaker reselling.



While Pop Mart's business model primarily relies on generating hype for collectible products, often driven by a few viral products whose hype eventually dies down, it eventually has to be replaced with new hype products. Despite Pop Mart's surging revenues and profitability, investors piled in, chasing the stock higher. "Foolishly", we refused and stood on the sidelines, choosing to quietly pass on chasing hypes and companies generating hypes.

Applying Bottlenecks and Constraints to Memory.

Although memory appears structurally tied to this secular AI trend, we previously discussed our rationale for not buying memory stocks (e.g., SK Hynix, Samsung, Micron, the DRAM/HBM oligopoly) in our 2025 [letter](#), and we continue to hold our ground and decline to chase it.

Membulls (memory bulls) argue that this time is different and that memory is no longer commoditized and cyclical. The fundamental question is: can memory growth persist indefinitely despite memory players' profit-maxxing, being less aggressive with fab capacity expansion, rapidly raising ASPs, and profit margins hitting all-time highs?

Let's go back to first principles: **Why do LLMs require more and more memory?** Before 2017, the best models were based on recurrent neural networks (RNNs) or convolutional neural networks (CNNs). RNNs read a sentence one word at a time, slow and forgetful, each word waiting on the last. CNNs were the total opposite. They ran in parallel, but struggled to link distant words.

The 2017 Google paper "[Attention is All You Need](#)" proposed a new Transformer architecture, based solely on attention mechanisms. Attention did three things. It ended the waiting and unlocked parallelism, so models could scale and learn from huge amounts of text faster and more cheaply. Distant words stay connected no matter how far apart. It further enabled expansion beyond text to images, audio, and video.

Attention is the single foundational thread that runs through most transformer-based LLMs, from OpenAI's ChatGPT (Generative Pretrained Transformer) to Anthropic's Claude, Google's Gemini, DeepSeek, Grok, etc. Advances in larger, better models (software) and better compute (hardware) have allowed LLMs to improve significantly in intelligence and cost.


While chip compute raced ahead, memory crawled, and the main bottleneck is no longer compute but memory bandwidth, especially with the growing shift from training to inference workloads. **There is a fundamental flaw in our current transformer architecture.** Try having an extremely long conversation with your LLM, and you rapidly hit limits. Re-ask just one question in the same context window, and before you know it, you hit your limits again.

LLMs are inherently memory-hungry by design. The larger the model, the more parameters it has and the more weights it must store. Model weights are largely a fixed cost. Weights never leave once loaded. KV (key-value) cache is the hidden memory killer. And it goes back to attention being the real culprit for the memory wall.

When an LLM generates text, it works autoregressively. It produces one token at a time, and each new token feeds back in as input for the next. To avoid redoing old math, it saves the keys and values of every past token in the KV cache.

Each new token must read the entire KV cache before generating the next one. So the KV cache grows with every word, and memory grows without end. Not compute. Memory is the bottleneck. The KV cache is the rent, and autoregression is why the meter keeps running.

As a result, **KV cache grows linearly with context.** Double the context, double the cache. The cache stores one key and one value for each token at each layer. That is why long LLM conversations often crash first. The transformer won in training. In inference, it pays the bill with memory, and the bill keeps rising. Accordingly, estimates from [Epoch AI](#) indicate that memory's share of total AI chip component costs (logic + CoWoS + HBM + auxiliary) has risen rapidly from 51.9% in 1Q24 to 63.4% in 4Q25. This could head even higher into 2027/28 if little is done.



Compare LLMs to the human brain. The human brain does not need to get larger as it learns and writes more. For perfect recall, the LLM wins. It will quote you word for word. The human brain cannot remember every detail. Ask for the exact sentence, and it fumbles.

For everything else, the brain wins, and not by a little. It never crashes when the conversation runs long. It runs on a light bulb, not a power plant, and it remembers for decades. The human brain is not a transformer that stores every token and remembers everything. It keeps a small, fixed state, remembers rules and what matters most, and forgets the rest.

Since then, there have been numerous patchwork software and hardware fixes targeting the same enemy: the KV cache. They address the symptoms but not the root cause. Software fixes have ranged from storing fewer tokens to compressing, lowering precision, looking at fewer past tokens, looking at earlier tokens, and managing the KV cache more efficiently, but most end up reducing quality and detail to varying degrees as a tradeoff.


Hardware fixes have ranged from increasing off-chip memory capacity and speed with high-bandwidth memory (HBM), the core driver of the current memory “boom”, to integrating memory on-chip (e.g., SRAM in the Groq LPU and the Cerebras wafer-scale chip).

Hardware raises the ceiling, but it does not move the wall. The software, the transformer with its attention mechanism, still insists on storing every token. A bigger, faster, closer memory just lets you pay more and more rent before you crash. The AI industry is unlikely to remain hostage to memory players for long because of this fundamental architectural flaw.

Constraints driven by “profit-maxxing” by memory players will inevitably spur disruptive innovation. Something has to give. We don’t know when and how. Perhaps either a brand-new model architecture or a hybrid architecture, with one that refuses to store everything in the first place, or intelligently stores only what is most relevant.

Stepwise software or hardware fixes should gradually lower the memory cost per token and drive [Jevons Paradox](#), as cheaper tokens get used more, not less. But a disruptive architecture does not threaten memory capacity, it threatens the high-bandwidth memory (HBM) premium.

Kill the growing KV cache, and the very bottleneck that makes HBM precious gets softened rapidly. The risk to HBM was never that the world needs less memory, but that it stops needing memory to be this fast. That is the risk from above. There is another from below.



When the top oligopolistic players gave up commodity memory to pursue premium memory, it opened the door for Chinese players (e.g., CXMT in DRAM and YMTC in NAND) to fill the gap over time. When that happens, it usually does not end well. **So HBM sits squeezed. Architecture risk from above. China risk from below, and premium in the middle. Constraints driven by scarcity that spur disruptive innovation are necessary**, and we are getting to that stage. Hopefully, we get boxed in enough to think out of the box.



The Root Finds a Way.

Fund Actions

We exited 2 of our entire positions in **Lululemon (LULU)** and **Paycom Software (PAYC)** in Q2.

Lululemon (LULU) was the most disappointing for us and hurt the most. It was sized as our 2nd-largest position when we started the fund (5.1% as of Oct 2024). We added twice to it in June and August 2025 in smaller positions. Our complete exit resulted in a 55% loss from our average cost. In hindsight, our Lululemon presentation in August 2024 at WOA Best Ideas Day introduced an undesired commitment and consistency bias. Public conviction hardened our private judgment, and we sized the position larger than we should have. We should have weighed more heavily the rapidly decelerating growth in the Americas business from 2023 to mid 2024, and the pushback against increasing competition from other brands. We failed to realize that the same CEO, Calvin McDonald, who had revived Lululemon, struggled to reinvigorate the US business. The ship was sinking, and it was increasingly hard to save it.


We defended Lululemon in our [3Q25 letter](#), and shortly thereafter, further struggles ensued. Calvin McDonald stepped down as CEO on 31 Jan 2026. Following a four-month search, Nike executive Heidi O'Neill was appointed on 22 Apr 2026 and will join as CEO on 8 Sep 2026.

Given her strong background across consumer sports apparel, women's, global leadership, brand marketing, marketplace, and direct-to-consumer, her appointment was deliberate. We hope she can revive Lululemon. We will be rooting for her from the sidelines.

In the meantime, Lululemon's troubles did not seem to end, which leads one to question whether they are transitional or increasingly structural. Sales continued to moderate, dragged by negative social media commentary (persistent) and weaker-than-expected product launches. Given long product development cycles (>12-15 months), Lululemon may need 12-24+ months to fix this, if they can.

Americas (66% of 1Q26 revenues) struggled throughout 2024, with low single-digit revenue growth, and swung into 3 consecutive quarters of low-single-digit negative growth (i.e., -2.1%, -4.1%, and -3.2%). Management is guiding for it to worsen further in 2Q26 to -10% to -13%.

Even Lululemon's burgeoning China business (19% of 1Q26 revenues) which had been growing revenues very strongly at >20% alongside strong profitability (66% GPM and 42% EBIT margins, versus America's 55% and 25% and ROW's 54% and 19%), saw momentum slow in Q1 and early Q2 driven by spikes of negative commentary over PFAS and faced strong backlash over using a Japanese taiko drum during its Great Wall Yoga wellness event.




Lululemon is an expensive lesson: investing in consumer brands is really hard, as consumer trends can evolve rapidly. In addition, we have intentionally sought to avoid single-stock pitches that let us pound our chests and the unnecessary introduction of negative biases.

Paycom Software (PAYC) distinguishes itself in the highly competitive human capital management (HCM) market through its proprietary “single-database” architecture and its pioneer status in employee-driven automation. While competitors often struggle with the complexity of integrating acquired technologies, Paycom's platform was built from the ground up as a unified system, ensuring that data entered once flows through the entire employment


lifecycle in real time. This architectural simplicity was a major differentiator, driving significant ROI for clients through automated payroll. But as Paycom moved upmarket from SMBs, its “closed” ecosystem, while secure, can be a reason for attrition among larger clients who require open APIs and complex third-party integrations.

When we first bought Paycom in Oct 2024, our thesis was that the slower quarterly revenue growth in 1Q24 (+10.7%) and 2Q24 (+9.1%) was more transitory than structural, and we expected it to reaccelerate. While it did reaccelerate slightly to +11.2% and +13.6% in 3Q24 and 4Q24, growth subsequently decelerated further into the 6-10%+ range.



While Beti's automated payroll system increased efficiency for customers, it cannibalized Paycom's revenue and reduced payroll-related billable events, such as corrections and unscheduled runs. In addition, to drive Beti adoption, client relationship representatives who typically focused on cross-selling new modules were redirected to spend time training existing clients on Beti, resulting in a booking headwind as cross-selling activity declined.

While Beti could arguably be more transitory than structural, the persistent rise in client attrition, particularly among smaller clients, which Paycom focuses on, has been a significant drag on its revenue growth since 2024. Paycom's retention rates, which historically sat between 93% and 94%, began declining to 90% and even 89%. Paycom's revenue growth has lagged peers such as Workday, Paychex, and Paylocity, and its market share has stagnated.




Paycom was a “turnaround” play. Thankfully (like PayPal), we had sized the position to be the 2nd-smallest at 2.3% in Oct 2024, and we were reluctant to double down throughout, as its business fundamentals remained poor. As we became increasingly disappointed quarter after quarter, we decided to call it quits and sell our 0.8% position in 2Q26 at a ~18% loss. Overall, while the total loss was small and did not hurt, Paycom (like PayPal) was another reminder to us that turnaround businesses generally do not turn, and it's far better for us to avoid them.

New Fund Write-ups

We wrote two new memos in Q2 2026 on [Tencent Holdings \(HK: 700\)](#) and [Mastercard \(MA\)](#). You may find all of our memos in [Partner Materials](#), accessible with your login credentials.


Tencent Holdings (HK: 700) is China's most dominant and embedded internet company, woven into the daily lives of over a billion people across messaging, gaming, payments, video, music, and commerce. It understands Chinese consumers deeply and has a highly competent management team focused on relentless innovation. Weixin/WeChat, its unrivaled super-app with 1.4bn monthly users, commands roughly 76% of social app time. Tencent's ecosystem takes about 55% of all mobile time in China, with Weixin alone at 35%. It is a super-app no company has replicated. WhatsApp, PayPal, Instagram, Uber, Amazon, and your bank in one app you never leave. Tencent is also the world's largest games company by revenue, owning evergreen titles like Honor of Kings alongside Riot, Epic, and Supercell. Its moat is not a single advantage but many that reinforce one another. Network effects, high switching costs, massive scale, a traffic dividend for every new venture, and cornered resources in gaming IP, proprietary data, and licenses. Rivals compete with pieces of Tencent. None competes with the whole.



Tencent was seen as slow in generative AI. That was timing, not weakness. It has now reached for its historical playbook, copy fast, innovate faster, disrupt the incumbent, and it is running the same move into the agentic era. Tencent chooses an open ecosystem and prioritizes user experience over exclusivity, a pragmatic aggregator rather than a walled developer. Its breadth


across mobile, PC, and cloud lets it deploy agents that work across devices and apps far better than any centralized rival. Weixin could well be China's largest consumer AI opportunity.

Tencent has compounded through every regulatory cycle, macro downturn, and competitive assault of the past two decades, and its business today is stronger, wider, and more profitable than ever. With 10/12/14% revenue growth over five years, 30/32.5/35% FCF margins, 15/20/25x EV/FCF, and a 20% haircut on its investment portfolio, we see a 15-32% CAGR, with a 24% midpoint. Tencent is not an AI loser. Higher AI capex will weigh on near-term profit, yet it already shows up in stronger unit economics, and Tencent could increasingly be AI-priced. A founder, Pony Ma, who owns 8.8%, still shows up after 28 years to keep Tencent's moat compounding. That is why we own Tencent.




Mastercard (MA) operates the second-largest card payment network (35% of GDV) outside China, in a near-duopoly with Visa (55% of GDV), which controls over 90% of global cards and spend. Visa was built by one. Mastercard was built against one. It does not lend. It does not issue cards. It does not take credit risk. It sits beneath every card payment transaction and takes a thin slice (31bps net). Mastercard rides two powerful tailwinds.


The first is the electronification of payments, growing faster than GDP and consumption as cash continues to die. The second is a fast-growing, high-margin value-added services business, now \$13.3 billion, or 41% of net revenue, larger than Visa's both in size and share. Many feared fintechs would displace the payment networks. Instead, the payment networks ended up powering them. Mastercard serves over 80% of the top fintechs, including Nubank, Wise, Mercado Pago, Monzo, and Revolut. What was built against a rival is now the network of networks.



Mastercard is more international than Visa, with roughly 71% of revenue earned abroad, and its international revenue continues to grow faster than in the US. Its economics are exceptional.



Mastercard's ROIC of 55% is nearly twice Visa's. If we could own only one of the two, despite Mastercard being smaller, its higher ROIC and faster growth profile edge it over Visa.



The real threat is not buy-now-pay-later or stablecoins but the rise of domestic real-time public rails like Brazil's Pix and India's UPI, alongside regulatory pressure on interchange. We size those risks and still like the business. At an undemanding 24x EV/FCF, if revenue grows at 12/14/16% into 50/52/54% FCF margins, into 25/30/35x EV/FCF valuation multiples, forward

returns could range between 15% and 30% a year, with a 23% midpoint. The more the world spends, the more it earns, and the more it can reinvest. That is why we own it.

Continued Confidence and Enthusiasm in our Holdings.

Overall, we remain more optimistic and enthusiastic about the long-term returns of stocks and of our holdings in the Vision Capital Fund than we were previously. 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, remain largely intact and are expected to continue to drive the majority of our long-term returns rather than changes in valuation multiples.

Our holdings are faster-growing and more profitable, with **29.4% trailing 12-month (TTM) revenue growth** and **26.5% TTM FCF margins** on a weighted-average basis (see below).

No.	Holding	Percentage Allocation	1Y TTM Rev Growth	1Y Fwd Rev Growth	TTM FCF Margin	Secular Theme
1	Palantir Technologies Inc	8.9%	67.7%	72.4%	51.5%	Software (Data Integration)
2	Taiwan Semiconductor Manu	8.8%	30.7%	34.0%	25.7%	Semiconductors, AI
3	CrowdStrike Holdings Inc	7.0%	23.2%	23.6%	29.6%	Software (Cybersecurity)
4	Shopify Inc	5.9%	31.9%	28.1%	17.1%	Ecommerce/Payments
5	Nu Holdings Ltd	5.8%	34.5%	41.7%	18.2%	Payments/Digital Banking
6	NVIDIA Corp	5.1%	70.7%	81.8%	47.0%	Accelerated Computing/AI
7	Pro Medicus Ltd	5.0%	30.3%	28.7%	52.9%	Software (Medical Imaging)
8	Meta Platforms Inc	4.8%	26.2%	25.8%	22.5%	Digital Advertising
9	MercadoLibre Inc	4.5%	42.1%	40.1%	37.2%	Ecommerce/Payments
10	Wise Group PLC	4.3%	19.2%	15.6%	33.3%	Cross-Border Payments
11	Amazon.Com Inc	3.7%	14.2%	14.9%	-0.3%	Ecommerce, Cloud, Advertising
12	Cloudflare Inc	3.6%	31.6%	29.8%	13.3%	Software (Connectivity & Security)
13	Tesla Inc	3.5%	2.3%	8.2%	7.2%	Electric Vehicles/Solar/Robotics
14	ServiceNow Inc	3.2%	21.7%	22.0%	33.2%	Software (Workflow Automation)
15	Sea Ltd	2.7%	40.5%	30.5%	18.6%	Ecommerce/Payments/Gaming
16	Zscaler Inc	2.3%	24.6%	24.6%	30.4%	Software (Cloud Security)
17	Adyen NV	2.0%	17.9%	16.0%	44.7%	Payments
18	JD.com Inc	1.8%	10.3%	8.8%	1.8%	Ecommerce
19	Spotify Technology SA	1.8%	8.0%	9.7%	18.1%	Streaming (Audio)
20	Netflix Inc	1.7%	16.7%	13.8%	25.4%	Streaming (Videos)
21	Meituan	1.6%	5.5%	13.6%	-12.0%	Ecommerce (Services)
22	MasterCard Inc	1.6%	16.8%	13.1%	52.4%	Payments
23	Tencent Holdings Ltd	1.4%	12.8%	13.4%	30.0%	Digital Advertising/Gaming/Payments
24	Microsoft Corp	1.3%	17.9%	17.0%	22.9%	Software, Cloud
25	Trade Desk Inc	1.0%	15.6%	9.7%	28.4%	Digital Advertising
	Cash	6.6%				
Weighted-Average			29.4%	30.3%	26.5%	

Source: Koyfin and Vision Capital Fund (as of 2 Jul 2026)

Final Words

Thank you for your continued trust and support in allowing us to invest for the long term. We 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 would like to 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. It 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.