



Portfolio Managers' Quarterly Letter Q4 / 2025

2025: A YEAR OF THREE MARKET DECLINES

INTRODUCTION

2025 was a year marked by extreme price volatility. We weathered three significant declines, and at each one the market commentariat declared the AI trade was over. In moments like these, the pressure to 'do something' — to trade around the noise, cut winners, or move to perceived safety — is enormous. It would have been easy to capitulate.

We did not.

Our inaction was not stubbornness, but the philosophy we presented to you at the start of the year: we do not watch the waves on the surface of the Atlantic; we watch the Gulf Stream beneath them. In our Q1 2025 letter, written during the height of the China-tariff scare, we argued that the fundamentals of our thesis were unchanged. We stressed in particular Jevons' paradox — the economic reality that when compute costs fall, demand does not slow; it explodes. We held to that all year.

In 2025, the largest risks came from rational thinking. Software companies rationally protected their margins. Analysts rationally depreciated down hardware acquisition costs. Both decisions were wrong, because the economic rules of the AI era have turned upside down — at least partly.

This review constitutes marketing material prepared by Asilo Asset Management Oy. This review does not constitute an offer or solicitation to subscribe for, redeem, or exchange fund units. Investors should base their investment decisions on their own assessment and take into account their individual objectives and financial situation. Investment decisions should not be based on this review. While care has been taken to ensure the reliability of the information contained herein, Asilo Asset Management Oy cannot guarantee the completeness or accuracy of the information and accepts no liability for any errors or omissions. All fund investments involve financial risk. The risks are described in more detail in the fund's key information document and fund prospectus. The value of an investment in the fund may rise or fall, and investors may lose part or all of their invested capital. The fund's target return may not be achieved. Past performance is not a reliable indicator of future results. Prior to making any investment decision, investors should carefully review the fund's prospectus, key information document and rules, which are available from GRIT Rahastoyhtiö Oy and/or Asilo Asset Management Oy. The fund is managed by GRIT Rahastoyhtiö Oy.

THE REAL RISK: THE INNOVATOR'S DILEMMA AND THE MARGIN TRAP

To understand why rational decisions can prove fatal, we need to look at the framework developed by Professor Clayton Christensen: the Innovator's Dilemma.

Christensen studied industries from hydraulic excavators to disk drives and found the same pattern recurring throughout history — but his analysis of the steel industry offers the most chilling parallel to today's SaaS market.

He observed a battle between the massive, established 'integrated steel mills' and the small challengers, the 'mini-mills'.

The integrated mills were giant, high-quality operations with massive overheads. The mini-mills used cheaper technology (melting scrap metal), which initially produced lower-quality steel.

Here was Christensen's critical insight: **the integrated mills did not lose out of stupidity. They lost because they acted rationally.**

The mini-mills started at the bottom of the market by making **rebar** (reinforcing bars used in concrete). Rebar is a commodity with miserable margins. The integrated mills looked at their numbers and made the logical decision: 'Why fight over this low-margin garbage? Let the mini-mills have it — we'll focus on higher-margin sheet steel.' (Structural steel is a critical material for ensuring the durability and safety of buildings.)

They voluntarily retreated upmarket to protect their gross margins. On paper, this looked like a brilliant financial decision. Their profitability metrics in fact improved as the low-margin business was culled.

But the mini-mills did not stop at rebar. They used their cash flow to improve their technology. Soon they could manufacture **angle iron and bars** — and they could do so more cheaply than the incumbents. Once again, the integrated mills looked at their margin profiles and stepped aside: this market's margins were compressing, so they abandoned it and focused on the highest-margin product — sheet steel.

Time after time, the mills placed margin percentages ahead of market share.

Eventually the mini-mills found a way to make **sheet steel** as well. At that point the integrated mills had nowhere left to retreat. They had optimized their margins to the end — right up to the moment they became irrelevant. They were eaten from below, because they chose margins over market share.

AMAZON'S LESSON AND SAAS'S AI MARGIN TRAP

We saw the same pattern repeat in the early 2000s. Bricks-and-mortar retailers looked at Amazon and laughed. They saw a company burning cash on a margin structure that made no sense compared to their efficient physical stores.

Retailers fell in love with their margin percentages. Amazon, by contrast, focused on **absolute** cash flow. Amazon accepted a structurally lower margin to lock in customers and to reach enormous volume. The retailers who prioritized their 'superior' efficiency metrics ended up sitting on flawless margins atop shrinking revenue — exposed to the competitor that had already captured their customers.

Today the SaaS sector stands on the brink of its own 'mini-mill moment.'

For the past decade, software has been the best business model in history. You write the code once and sell it a million times at near-zero marginal cost. This created companies with gross margins of 80–90%. Investors learned to demand them — and to punish any compression.

Generative AI breaks this math.

AI is not 'write once, run an infinite number of times.' It is 'compute again every time.' Every inference, every reasoning step, every agent action consumes expensive GPU compute. A successful AI-native software company's gross margin is on the order of 40 percent, not 80 percent.

Think of the industry's giants: Salesforce, ServiceNow, HubSpot, GitLab, Atlassian. They have every advantage: proprietary data, deep customer relationships, and massive distribution. They are perfectly positioned to create the AI agents that will define the next decade.

But to do so, they have to accept a painful truth: the new products will be lower-margin than the old.

Instead of driving the transition forward, we see them hesitating. By clinging to their historical margin profile, they are ceding ground to AI-native competitors. These new entrants build agents that reach into the incumbents' workflows from the outside, pull their data and do the work — ultimately rendering the underlying 'system of record' invisible or, worse, obsolete.

The situation calls to mind Stephen Elop's famous 'Burning Platform' memo from 2011, when he was Nokia's CEO. He described a man standing on a burning oil platform faced with a choice: stay on the platform and burn, or jump into the freezing water.

The SaaS incumbents are standing on a burning platform. The 'freezing water' is AI's new economic reality — a world where compute is currency and margins compress. They need to jump in order to survive — and yet they refuse, because rationally, the freezing water looks worse than the warm platform they know.

This refusal creates a fascinating paradox. While the software giants feared jumping into the water, the markets looked at the ocean itself — the massive GPU infrastructure under construction — and declared it a mirage.

THE DEPRECIATION PANIC AND WHY 'A GPU IS NOT A TAXI'

At the same time that investors were busy punishing companies for GPU acquisitions, they should have been punishing those who refused to make them.

In early October a narrative spread that the GPUs in data centers were losing value too quickly — that the billions flowing into investments would soon be written down to zero as newer chips arrived.


We did not believe this story for a moment.

The market's mistake was to assume that AI infrastructure follows the mechanics of traditional machinery, in which value is purely a function of age. This ignores the unique nature of today's technology stack.

Consider a traditional asset whose value drops reliably with age: a taxi. When a taxi operator buys a new Škoda Octavia for €30,000, they already know almost exactly what the car will be worth in five years — perhaps €5,000. The yardstick is predictable: the odometer climbs, the body rusts, the engine wears, newer models arrive — and value falls accordingly.

This is the entire basis of depreciation accounting: allocating bookkeeping cost to real economic wear and tear. The rules allow depreciating the car evenly over five years — €5,000 per year — because that mirrors physical reality. If the owner suddenly chose to depreciate over ten years instead, he would 'save' €2,500 in annual costs and make the income statement look better on paper. But he would only be postponing the moment of truth: when that €5,000 scrap heap finally goes to the baler, someone pays the difference.

A Taxi Company's Annual Income Statement — 5- vs. 10-Year Depreciation



Revenue is always €60,000. Operating costs are €50,000 in both cases.

A €30,000 car's depreciation schedule changes from 5 years to 10 years → €4,000 in profit becomes €7,000 in "profit."

This is why depreciation schedules matter: they should reflect the actual wear of the asset. If you see a company arbitrarily stretching depreciation from five to ten years, that is a warning sign — they are dressing up their reported profit by pretending the asset lasts longer than it really does. The market panicked in Q4 precisely because this is what the hyperscalers appeared

to be doing: stretching GPU depreciation schedules while 'everyone knew' AI infrastructure would soon be obsolete.

This taxi logic is precisely how most people analyze the value of GPUs — but in reality, that taxi-stand logic takes you the long way round and straight off the track.

THE GPU — A MACHINE WHOSE VALUE CAN RISE WITH AGE

Let us step into a different reality.

Take Nvidia's A100 GPU, which shipped in May 2020. In November 2025 — five and a half years later — it is still running at full utilization in the clusters of hyperscalers and cloud providers. And it does not merely run; it produces more money than ever before.

Why?

Because the software has improved exponentially. CUDA, TensorRT-LLM, FlashAttention, quantization, vLLM, continuous batching — each update has made the same piece of silicon 4–6× more efficient on typical LLM inference workloads. In 2020 a single A100 produced perhaps 30–40 tokens per second on a 70-billion-parameter model. In 2025 the same card produces 150–180 tokens per second on models of the same size class. It is as if an old taxi suddenly began burning half the fuel, driving twice as fast — and creating entirely new demand at the same time.



This is the perfect opposite of the traditional law of wear. In the old world, value falls with time. In the GPU world, value rises with time, because software extends the hardware's economic useful life faster than the silicon physically degrades.

That is why hyperscalers and cloud providers have justifiably extended GPU depreciation from four to 5–6 years — and why NVIDIA itself hinted on its Q3 call that the true economic useful life may be longer still. This is not accounting trickery. It is honest financial reporting that reflects a new physical and economic reality.

FEAR VS. EVIDENCE

While the markets panicked in Q4 about an 'AI bubble' and accused companies of depreciation games, we listened to the Q3/2025 earnings calls:

- **NVIDIA:** "A100s delivered six years ago are still in full use today — now driven by a significantly stronger software stack."
- **Alphabet:** "Inference-driven demand continues to exceed all expectations; older generations are running at full capacity."
- **Bloom Energy:** Fuel-cell installations for AI data centers are sold out for years ahead; gross margins are growing rapidly.
- **Astera Labs:** Q3 revenue rose 20% over the prior quarter; demand exceeded forecasts so severely that they cleared their inventories to fill orders. Multiple product generations are shipping at volume in parallel.
- **Nebius (formerly Yandex):** The European AI cloud is scaling faster than planned; GPU utilization is near 100% on older generations.
- **CoreWeave:** Customers are begging for more capacity; raised \$12 billion in funding at a valuation of over \$30 billion in the fourth quarter — and did all of that amid the deepest panic. Sophisticated infrastructure investors voted with real capital.

The earnings releases and the value chain offered an enormous amount of data, but in our view the market picked up the wrong things from it. The relevant core was, and remains, in conflict with the Q3/2025 bubble rumours.

The fundamental drivers we have identified did not merely hold steady — they strengthened quarter by quarter. We made one meaningful sale — Palantir — not because we lost faith in AI, but because the valuation had reached an extreme level detached from any sensible scenario. We reinvested the proceeds into positions where we saw better risk-adjusted returns within the same structural theme.

OUTLOOK FOR 2026

The AI era will produce many more misunderstandings like these. The markets will repeatedly panic about things that look perfectly rational through the lens of the old world, but that are meaningless or wrong in the new world. Each such panic is an opportunity for those who have done the homework — who can tell the difference between yesterday's framework and today's reality.

We expect 2026 to bring continued price volatility as the markets struggle to fit old frameworks to new realities. SaaS companies will face growing pressure as AI-native competitors demonstrate what can be achieved with lower gross margins. Infrastructure bottlenecks — energy, memory, interconnects — will remain binding constraints on AI adoption, which creates durable pricing power for our portfolio companies.

Our job is not to chase every small mispricing. Our job is to hunt for precisely these structural misunderstandings — because that is where the greatest alpha hides.

Thank you for your trust,
Henri Blomster & Ernst Grönblom