4Q 2025 Co-Investor Letter

4 March 2026

Dear Co-Investor

As we close our fifth year managing the fund, we returned **+30.78%** in 2025 and **+15.39%** annualised since inception (benchmark return: 2025: +8.68% and +10.06% annualised since inception)[1]. The fourth quarter was volatile, resulting in a negative quarterly return[1] - our first since 2022 - serving as a reminder that compounding is seldom a smooth progression.

This letter also marks **our firm's tenth anniversary**, which serves as a waypoint on our 100-year journey to reflect on what we've learned, update you on our thinking, and discuss a few of our portfolio companies.

Firstly, what did not go well in the quarter:

- **Stronger ZAR:** the strengthening rand has been a headwind for our largely non-ZAR-denominated portfolio.
- **E-commerce Volatility:** turbulence in our e-commerce portfolio companies, Sea Ltd (Southeast Asia) and MercadoLibre (Latin America), amidst aggressive price wars.
- **Blue Label / Cell C**: continued uncertainty around Blue Label Unlimited (BLU) and the contentious Cell C IPO (more on this later).

Launch of Our USD Global Opportunity Fund

This week, we are launching our USD-denominated **SaltLight Global Opportunity Fund**. This new fund mirrors our existing opportunity set and will suit investors who've already externalised capital while opening a broader global investor base.

The fund launches imminently —please reach out to us if you would like further information.

Lessons from SaltLight's First Decade

When we speak to potential investors, the first hurdle is often conceptual: SaltLight is deliberately hard to categorise. We're not cleanly "growth", "value", or even "tech". This ambiguity is not an identity crisis but a design choice, born of a painful lesson from the firm's early years.

A decade ago, we ran what looked like a disciplined setup: one asset class (equities), one geographic market, and… focus. It felt logical. However, it was also structurally fragile. For five long years, we were confined to a market where inflation-beating returns were unavailable, regardless of how hard we worked or how "right" our research was.

1 *SaltLight BCI Worldwide Flexible Fund Class A1 Net of Fees. Benchmark is South African CPI + 5%. Annualised return is the weighted average compounded growth rate over the period measured. Fund inception date: 12 July 2021. Highest and Lowest calendar year performance since inception: High: 63.21% (benchmark: 12.85%), low: -34.46% (benchmark: 8.04%)*

In hindsight, the error was twofold:

- Overconfidence in predicting specific outcomes, rather than underwriting ranges of outcomes (more on this later)
- Constraining the opportunity set to a market and asset class where uncontrollable macro events dominated results

Five years ago, when we launched our current fund, we, by design, inverted those constraints. We adopted what we call an **"opportunity maximalist"** approach: **flexibility** as the defining characteristic, enabling us to invest across asset classes, geographies and opportunity sets. The intent is simple: direct capital **to where opportunities are**, rather than waiting (and hoping) for opportunities to appear inside a pre-defined box.

Now, that undefined flexibility does not always resonate with institutions that need a neat product label, which partly explains our boutique status. But it has also attracted a base of co-investors who, through natural selection, identify with an opportunity-driven approach to capital allocation.

To all our co-investors who have backed us over the last decade: thank you. We appreciate your trust in us to manage your hard-earned capital.

Reflecting on the Story Arc of the Last Five Years

Opportunities must be pounced upon. To put the last five years in context: in 2020, our research led us toward NVIDIA shortly after it completed its acquisition of Mellanox (a networking company). At the time, NVIDIA was still widely known as a gaming hardware business, but it had spent years quietly building a second engine: accelerated computing for *machine learning* (before "Artificial Intelligence" entered the zeitgeist).

The technical distinction was significant because neural networks—central to modern LLMs—rely heavily on matrix multiplication, and GPUs had become the primary hardware for that workload.

However, we were roughly on the right track, but it was too early. In 2022, our portfolio suffered a significant drawdown amid rising global interest rates. If you read our letters from that time, we admitted we didn't fully understand how AI would spread because of its complexity.

In late 2022, ChatGPT was released to the public as a research preview, and the world suddenly re-priced the probability that generative AI would matter.

Boom...the AI Epoch Begins

What surprised us wasn't that AI was effective; it was the slow initial response from the broader market, which viewed it as merely another hype cycle. Seeing near-term bottlenecks and AI's underappreciated impact led us to increase our position exposure to the AI epoch meaningfully.

Firstly, our early research and mapping of the AI stack had already highlighted likely bottlenecks to the technology's early irruption: memory and leading-edge semiconductor manufacturing. The market eventually converged on the same bottlenecks, but the timing advantage was meaningful.

Secondly, the opportunity was long-term and multifaceted. Carlota Perez's framework in her incredible book, *Technological Revolutions and Financial Capital,* was particularly clarifying for us in two ways:

- AI wasn't a single-application innovation like a new consumer appliance; it looked more like a general-purpose technological revolution with **multi-decade** second-order effects—across industries, labour markets, and national competitiveness.
- Technology epochs consist of a cluster of technologies that act as foundations or enablers for significant innovations, rather than being a single technology.

The *matter of size and impact* is important. As investors, we don't need to predict the winning application; we need to own the enabling constraints and the compounding downstream optionality. We can pick our spots by applying our investment process and understanding where companies generate durable, differentiated returns.

Lastly, AI is deeply intricate. Complexity can cause market participants to misprice assets as they focus only on what they understand. We continue learning daily, but this ongoing learning process has created opportunities for alpha.

*****

Range of Outcomes and Expected Value Investing

Our philosophical mistake as 'young' investors was to bet on exact outcomes. Our maturing investment philosophy recognises that we are always making bets on an unknown future in a complex adaptive system.

We therefore prioritise maximising expected value (EV) over engaging narratives or "being right". We build and construct a portfolio that can absorb a range of outcomes.

EV, in simple terms, is the probability-weighted sum of outcomes you would expect if you could repeat the **same decision** over **many events**. Subsequent events that *actually* unfold are the part we don't control. We don't know what will happen, and 'variance' is the inherent randomness of the financial system we operate in. What we want to ensure is that, in both good and bad events, our portfolio importantly survives and, hopefully, thrives.

We must be clear. We never know the exact odds, but we can infer ranges of outcomes. Securities that are priced for near-perfect outcomes can still go up, but the distribution is asymmetrically against you: limited upside and a meaningful risk of disappointment. Conversely, securities priced for overly pessimistic outcomes can offer the opposite asymmetry.

Poker serves as an excellent training ground for this mindset. The best hand can still lose; what truly matters are position sizing, odds, and bankroll management. Investing parallels this: knowing when to fold an overheated sector, when to hold during uncertainty, and when to aggressively pursue a mispriced opportunity—without mistaking confidence for certainty.

Three Paradoxical Observations.

Entering 2026, the AI epoch remains transformational, and, as a theme, it has dominated US equity markets in both size and mindshare. The market, however, is struggling to price it coherently. We see three paradoxes today — and, importantly, these interdependent observations cannot all be true at once.

- **Upstream supply-chain participants are increasingly signalling that 2028 (and beyond) capex could be materially *higher*, not lower**: This signal directly contradicts the widely held "digestion" narrative, which assumes that AI-driven capex will taper meaningfully over the medium to long term.
- **Market Scepticism Toward NVIDIA Analyst Revenue Forecasts:** Sell-side analysts have increased their cumulative revenue projections for NVIDIA by $1.1 trillion over the next five fiscal years. However, the current share price suggests that the market harbours significant doubt about the achievability of these elevated expectations.
- **AI ROI is unclear - yet software is being repriced as if disruption is certain:** SaaS multiples have derated sharply on fears that AI will disintermediate seat-based software, even as many investors remain sceptical about near-term AI payback.

If those statements feel like they're pulling in different directions, that's the point. The market is pricing in conflicting outcomes, which often creates asymmetric EV setups for investors willing to be explicit about what they think is wrong.

Later in this letter, we provide an example of our probabilistic thinking regarding NVIDIA and the software sector opportunity. First, let's dive into what we're seeing.

Upstream supply-chain participants are increasingly signalling that 2028 (and beyond) capex could be materially *higher* rather than lower.

Something is intriguingly emerging beneath the surface of AI capital expenditures. Until recently, consensus had settled on a comfortable timeline: hyperscalers made their big commitments in 2024–2025; those data centres would come online around 2026–2027; and then, inevitably, there would be a digestion phase.

The problem is that the supply chain is signalling a different path.

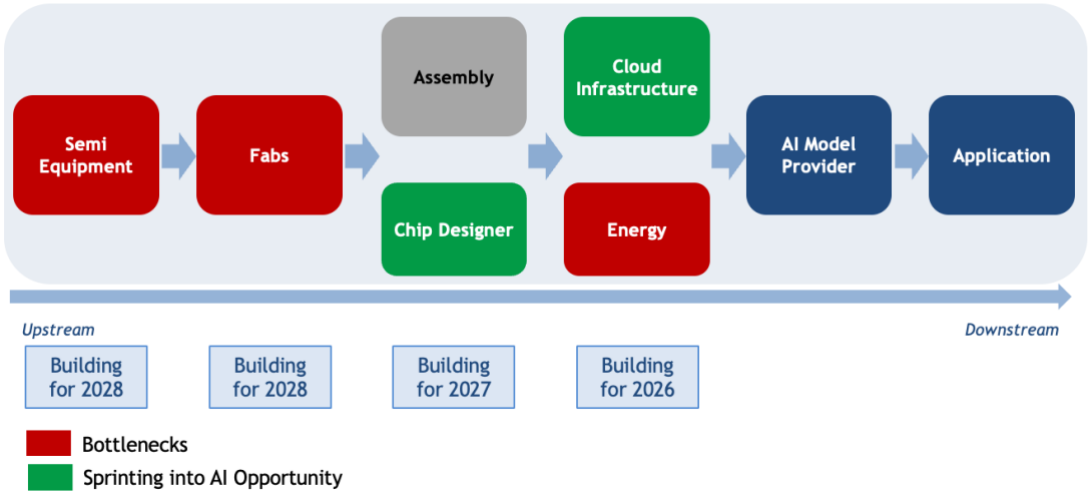

Stepping back, the AI supply chain is fairly neat and orderly:

- Equipment vendors (think ASML) enable foundries (TSMC, Intel) to add capacity.
- Foundries build chips for designers (NVIDIA, AMD).
- Hyperscalers (Amazon, Google, Microsoft) deploy those chips in data centres.
- Inference revenue then funds the next round of spend.

The key point for investors is that upstream orders act as leading indicators. If ASML's bookings and TSMC's capital expenditure suggest acceleration, it challenges the market's expectation of a 'digestion phase.'

We think something shifted in the 4Q25 earnings season: TSMC lifted capex expectations, ASML's bookings re-accelerated, and even Intel (a frequent punchline for missing cycles) found itself unexpectedly capacity-constrained in data-centre CPUs[2]. Historically, TSMC and ASML have been the cautious gatekeepers, hardened by semiconductor cyclical booms and busts. Therefore, we place great weight on any signals from them about acceleration.

What's particularly striking is that ASML, during its investor day back in 2022 (pre-ChatGPT mania), set out long-term 2030 guidance and hasn't yet revised it despite an AI-driven surge since. But with TSMC now breaking ranks, it seems clear that downstream expectations may need recalibration.

Which brings us to a falsifiable proposition: **Upstream bookings and capex imply multi-year visibility into demand**. If those signals persist, then one of three things must be true.

- Hyperscaler spend extends,
- New AI purchasers emerge (sovereigns, neoclouds, enterprise), or
- The utilisation/throughput assumptions used by the upstream participants were wrong.

Why does this matter? If it turns out that the Hyperscalers are ratcheting up capex, they are already deep into the build-out, and a sustained extension increases the odds of negative free cash flow — a meaningful regime shift for investors who instinctively underwrite hyperscalers as perpetual cash machines. The question would be whether the industry may be moving towards trillion-dollar annual data-centre investment sooner than the market is prepared to underwrite.

Portfolio Impact: This quarter, we took profits in our hyperscaler portfolio companies (Amazon and Google) and increased our position in NVIDIA (more about this later).

[2] *Interestingly, AI agents also need CPUs on top of AI compute*

Market Scepticism Toward NVIDIA Analyst Revenue Forecasts

NVIDIA is a useful case study in our expected-value thinking because there are currently meaningful dislocations in perspectives at the same time. The first is explicitly observable: sell-side revenue estimates. The second is inferred: the probability-weighted view embedded in the share price.

For readers outside the industry, the "sell-side" refers to analysts employed by investment banks who publish company forecasts. Those forecasts are aggregated into consensus estimates. They are often wrong, and they can be biased — not necessarily maliciously, but structurally, given the incentives created by investment banking relationships. They're useful for us to establish what one large collective market participant is thinking.

What matters here is not whether analysts are right, but whether the market's weight of a certain outcome is probabilistically accurate.

The Sell-Side is Very Bullish

Over the last six months, the sell-side has sharply raised its medium-term revenue expectations. Using consensus data from August 2025 and two February 2026 updates, analysts have cumulatively added roughly $1.1 trillion of incremental revenue across the next five fiscal years. As we've been invested in NVIDIA for many years, this recent upgrade is not a small "tweak" by the sell-side. Historically, they've consistently underestimated NVIDIA's forecasts since 2022. It could be that they've over-corrected by putting down some aggressive estimates, or they've seen something.

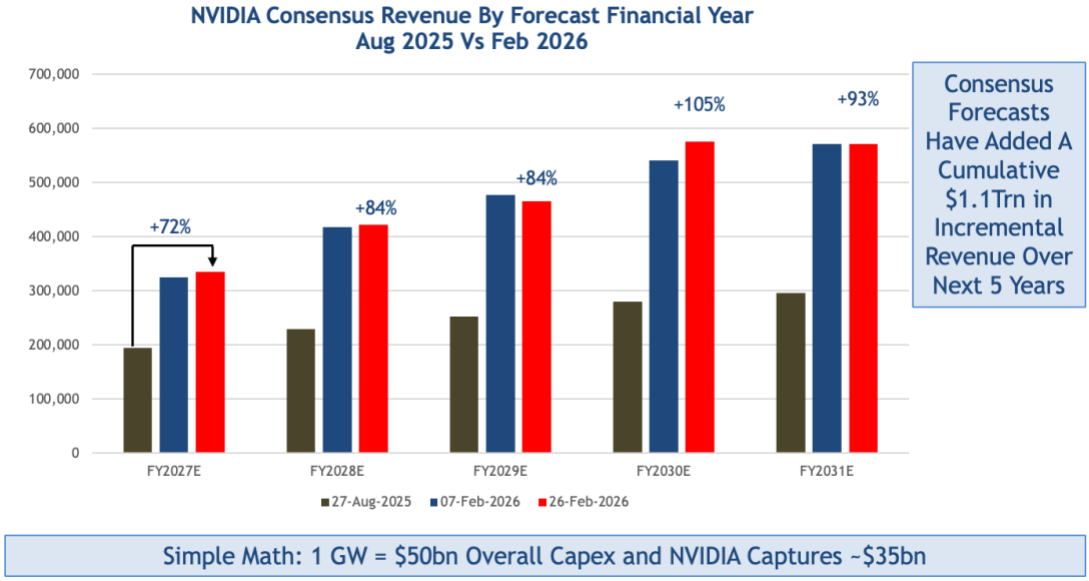

The Market is Not Buying These Numbers

Despite this significant upgrade, NVIDIA's valuation does not appear to reflect confidence in those figures. If you accept the consensus forecast as accurate, the stock is currently trading at a discount compared to the S&P 500, even though estimates suggest much higher growth.

A simple way to see this is through the multiple. NVIDIA trades at roughly 20x FY28 earnings. If the market believed the upgraded revenue trajectory was highly probable, then the multiple would typically be higher, not lower. *The clearest explanation is that the market is discounting the forecast path itself.*

This is precisely where expected value analysis proves valuable: we have two conflicting perspectives, and we can infer a probability-weighted distribution of outcomes from the price.

Why Scepticism is Rational (and Not About Demand)

The bearish case has merit, but it is less about end-demand and more about the physical and financial constraints of meeting expectations: **power supply, ROI, capital costs, and funding mechanisms**.

Start with what the forecasts *imply* in the real world. If you translate consensus revenue into data-centre build-out, you quickly arrive at requirements that look difficult to reconcile with current delivery capacity.

Management has suggested that a **1 GW** data centre costs roughly **$50–$60 billion**, and that NVIDIA captures around **$35 billion** of that via its share of the stack. The infamous Stargate campus in Abilene, Texas — a roughly **1.2 GW** complex — consists of eight warehouse-scale buildings and is scheduled for completion in **June 2026**.

Now map that to the implied build-out embedded in the current trajectory. The analyst path implies something like **~10 GW in FY27E**, **~12 GW in FY28E**, **~13 GW in FY29E**, and **~16 GW in FY30E–FY31E**. In other words, the market is being asked to underwrite multiple "Stargate-scale" campuses every year, for several years, while assuming NVIDIA retains a very large share of wallet.

That is a credible reason to discount the forecast path: not because the chips aren't wanted, but because delivering the supporting infrastructure is hard. It is easy to imagine reasonable failure outcomes:

- Energy capacity is delayed or not delivered on schedule.
- The cost of capital rises, tightening funding and slowing build-outs.
- Even if power is available, competitive dynamics shift and NVIDIA loses share (and we are not even assuming margin concessions to defend share).

Could Scepticism Already Be Priced In?

At the current multiple, we can postulate that the market is already up-weighting some of the above failure modes. The interesting question for us is whether the market has overreacted, placing too much emphasis on negative scenarios and too little on moderate or somewhat positive outcomes for a generational company.

Currently, we believe a more probable scenario lies between the market's bearish outlook and the analyst's bullish forecast. As time progresses, we expect to see more evidence that fears were exaggerated: ongoing execution, greater visibility into development, or a rise in capex could shift the market's implied probabilities.

As mentioned, we have increased our position in NVIDIA in the recent quarter.

AI ROI is unclear - yet software is being repriced as if disruption is certain.

If you've managed to keep up with our discourse so far, we must warn you about another probability discussion (*would the hecklers in the back of the room please refrain from expletives!)*.

The recent software sell-off on "AI disruption" has really given us reason to tap dance to work. It's a lesson in probabilities again, and we think it is quite a gift for software incumbents. Multiples are compressing as if AI disruption is inevitable. At the same time, hyperscalers and enterprises are still struggling to demonstrate clear AI ROI to their investors. *How can both statements be true?*

We have been thinking about AI's impact on legacy software economics and platforms for a while. In our 4Q24 letter, we argued that AI would pressure parts of SaaS economics and force management teams to adapt.

The core issues are 1) software has generally been an augmenter of the unit of "work", not a "doer" of work, and 2) the SaaS economic model is fragile to the upcoming platform shift.

A typical SaaS playbook was built around three pillars:

- **A stable, high-margin cost structure:** High gross margins, low variable costs, and a belief that usage growth did not meaningfully change unit economics.
- **Lose money now, expand later:** "Land and expand" through seats, modules, and renewals, with sales efficiency improving over time.
- **Deferred profitability.** Stock-based compensation and non-GAAP framing that kept the focus on revenue growth while GAAP profitability stayed "a tomorrow problem".

But AI is a challenge to all three.

- **AI eats into margins.** Inference is usage-linked. Even if costs fall, the economic shape changes: margins become a strategic choice, not a structural fact.
- **It attacks the pricing unit.** The more work that is automated, the less obvious it is that "seats" should be the unit of work. That doesn't mean spend disappears; it means it migrates to task-based outcomes, consumption, or new bundles.
- **It creates a cannibalisation problem.** The best defence is often to automate what you used to charge humans to do. That is strategically correct and financially uncomfortable.

So why do we see an opportunity now? Because the market has already done the painful part for investors — it has compressed multiples as if the worst case is the base case.

That matters for incentives. At high multiples, management teams are reluctant to take short-term pain and therefore avoid rocking the boat. At lower multiples, expectations of disruption are already priced in.

It is now up to management to lean into the platform shift. Self-disruption can now become a rational path for every software company to build a 'second act': move from augmentation to value participation, repackage and rebuild distribution, and accept near-term margin volatility to protect long-run relevance.

Investing in software today is therefore not a binary bet on whether AI "works". It is an expected-value bet: the market is pricing a broad swathe of SaaS as if disruption is certain, margins must collapse, and incumbents cannot adapt.

We think the survivors will have a few characteristics:

- Systems of record that can expose agent-safe actions,
- Vendors with deep governance, audit and compliance surfaces,
- Vertical software with proprietary data and where domain rules/workflows are complex,
- Platforms with real ecosystems, high switching costs and integration density.

The Software Opportunity Set is Likely to Double or Triple in the Next Few Years

The average knowledge-based enterprise spends ~6-7% of operating expenses on software and ~60% on headcount. Software has always been an augmenter of work rather than the "doer" of work. Over time, as AI replaces some work, we're likely to see that gap narrow, implying that the software total addressable market will grow 2-3x in the next couple of years.

Our job is to find the cases where that is too pessimistic — where distribution and workflow ownership still matter, where value capture shifts rather than disappears, and where management teams use this moment to earn a second act.

*****

Blue Label Unlimited and Cell C

A last word about a new position in the portfolio – Cell C: it is quite an old friend, but we've never owned it directly.

How We Got Here via Blue Label Unlimited

We first came to Cell C through **Blue Label Unlimited (BLU)**. BLU remains a compelling opportunity, even if it has been an exceptionally volatile business to own over the last three years: we have seen the market capitalisation move from roughly **R2.5bn to R15bn**, and then back down to around **R9bn** today.

We have written before that BLU's core business is a high free-cash-flow, capital-light distribution operation, currently valued at roughly 4x FCF. The challenge has always been that BLU's decade-long investment in Cell C has made the group difficult to underwrite. Until November last year, the accounting complexity and significant unconsolidated debt kept most investors away.

The Cell C listing

In November last year, BLU listed Cell C by selling down its 95% stake. The market reaction was surprisingly poor: early hopes of achieving an R18bn valuation were dashed. In the end, management sold 30% of the business at a valuation of R9bn.

At a high level, we think the bear thesis had three pillars:

- **Hysteresis**: Cell C has historically been the laggard in the sector: the smallest player (**c.** 7% market share) yet, historically, required to maintain a tower network

comparable to the incumbents. It has suffered from weaker network quality, low consumer mindshare, and low prepaid ARPU.

- **Uncertain terminal value**: Cell C has shut down its tower network and now relies on roaming agreements with the two largest operators. These agreements expire in 2029 (Vodacom) and 2032 (MTN)**,** creating the risk - at least in theory - that the business could be left without network capacity.
- **A mature mobile market**: South Africa's mobile market is mature amid ongoing regulatory pressure on pricing. Competition remains intense, particularly in prepaid, where Cell C has historically competed on value.

What the Market Missed

We have been fortunate to have had the opportunity to get to know BLU and Cell C management over the last three years. We have spent significant time discussing their strategic repositioning and the opportunity ahead. As BLU shareholders, we wanted management to maximise the value of any sell-down, and we attended several IPO roadshows. What stood out was that most participants could not get beyond the bear thesis to the new economics of the business.

That matters because the Cell C of the past is not the same as the Cell C being built today.

What Cell C Could Be

Cell C's new economic model radically transforms it into a capital-light FCF generator with mid-teens FCF growth over the next couple of years.

- **Balance sheet and execution:** It is now essentially debt-free, aside from asset-backed financing related to post-paid contracts and lease liabilities, and is supported by a top-tier management team.
- **Free cash flow and dividends:** We estimate the business could generate roughly R1.6–R1.8bn of free cash flow in 2027, and we expect 50% of this to be returned to shareholders via dividends.
- **Second act via MVNO platform:** Cell C is building a "second act" as an MVNO platform growing at roughly **20% p.a.**, as retailers and banks increasingly "white-label" mobile offerings and mobile/banking continue to converge

The terminal-value concerns around the roaming contracts are real. That said, we think the incentive structure is better than the market assumes it to be. The wholesale fees Cell C pays represent a meaningful profit stream for the larger operators, driven by higher utilisation of their tower assets. In addition, Cell C's spectrum can enhance network capacity for roaming partners. The result is closer to a mutual dependence than a simple supplier/customer relationship — and that is typically more durable.

Why We Participated in the IPO

BLU proceeded with the IPO at a valuation of R9bn. Based on our estimates, that implied a **c. 19% free-cash-flow yield** — meaning new IPO investors would capture most of the discount, not BLU shareholders. The value left on the table was sufficiently compelling that we decided to participate directly in the IPO.

*****

As always, we'd like to remind our co-investors that your manager's liquid wealth is invested in the same fund as yours. We share the inevitable ups and downs with you.

If you have any questions, please don't hesitate to get in touch.

David Eborall - Portfolio Manager