3Q 2026 Co-Investor Letter - "Intelligence"

18 November 2026

Our fund's returns are much the same as they were three months ago, which, in a quarter as turbulent as this one, feels rather like getting home from a long drive in the bush with the chassis intact. Whilst the portfolio return has been stable, a look under the bonnet reveals tremendous volatility amongst the companies we own. This volatility has allowed us to reallocate capital into a handful of businesses that we hope will reward us over the next few years (a topic for future letters).

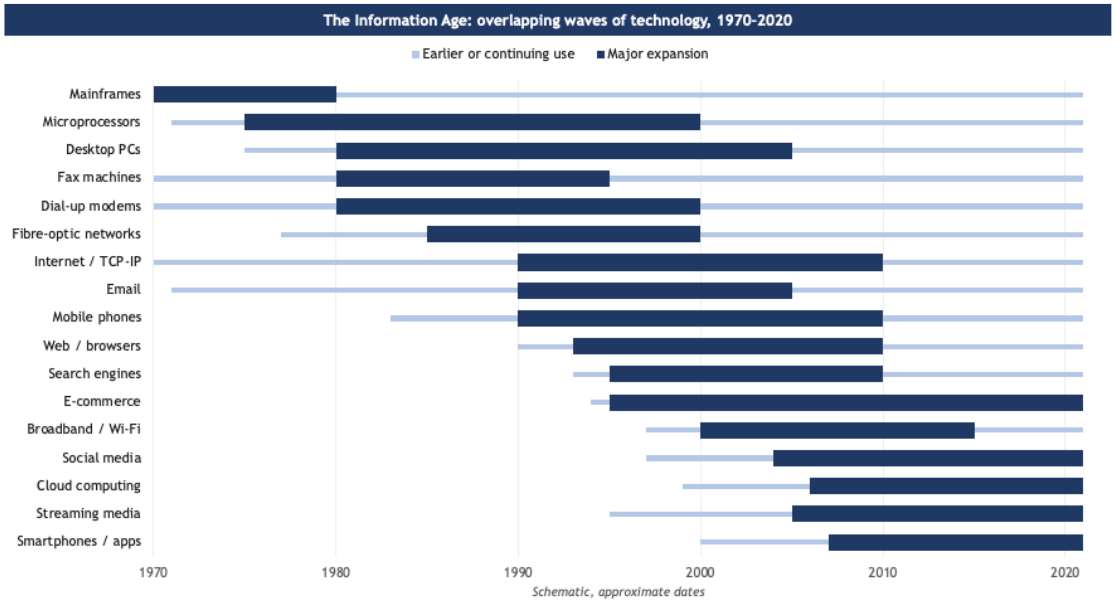
Given our unconventional portfolio and longer-term horizon, we have always maintained that our fund returns are likely to zig when the market zags (and 2026 has been a year of a good deal of zagging). We build a portfolio around durable, differentiated, and resilient businesses with multiple pathways to win. Our time horizon is focused on what a business might look like in 2031, while letting short-term market movements take care of themselves.

As always, our letters aim to share what is going on in our heads, with the standing caveat that we have no ability whatever to guess what the market will do next.

The Two Clocks of the AI Epoch

We maintain our view that we are very early in a five-decade-long AI Epoch. Our rough perspective is that we are likely in the equivalent of the Information Age's "mainframe era" in the late 1970s (for reference, the Information Age epoch ran from the 1970s to the 2020s).

How does one sensibly invest at the start of something that may take fifty years to play out? Long-term co-investors will recall our discussions of Carlota Perez's framework of Technology Epochs, which reminds us that epochs are not singular technologies, but clusters of technologies feeding off each other. History shows that, more often than not, a new micro-technology usurps a prior one (just as desktop computers usurped room-sized mainframes) through a process of creative destruction.



The investment lesson from the Information Age is that these micro-technologies move through their own rapid cycles long before the broader epoch has run its course.

Therefore, at SaltLight, we keep two clocks running. One measures the medium term and guides our tactical decisions; the other spans decades, keeping us focused on the long-term value we hope to capture throughout the AI epoch. From time to time, we act on the shorter clock (sometimes playing offence, sometimes defence) as individual cycles unfold, while keeping the broader horizon in view.

The Capital Cycle: A Looming Supply Glut?

Astute co-investors will recall from our letters that over the past few months we have adopted a more defensive stance on AI infrastructure. The sector has performed exceptionally well for us over the past few years, but we're concerned about the sheer volume of compute supply expected to hit the market in the short term.

From a Clayton Christensen 'jobs to be done' perspective, AI models have reached the point where, at current prices, they are 'good enough' for 95% of current tasks, and we expect competitive arbitrage to eat into current healthy gross margins (see our 2Q26 letter for more). Societal adoption simply needs time to catch up with what the technology can already do, even as we wait for the next breakthrough.

Another clue to where we stand in a technology's cycle is the recent crop of early-stage IPOs. When venture capitalists and other early shareholders begin passing risk to public investors, it is often a reasonable moment to take money off the table.

Our primary challenge today is a disconnect between narrative and reality. The third-quarter earnings message from AI companies was uniform: demand still exceeds supply, spot prices for AI compute imply paybacks of just one to two years, and adoption remains strong.

Around a decade ago we came across a wonderful book, *Capital Returns*, Edward Chancellor's collection of letters from Marathon Asset Management, a firm we very much admire (recommended reading for any student of business). What we have found most useful is Marathon's knack for spotting a capital cycle as it turns. The difficulty, as Marathon describes it, is that investors extrapolate the present too far into the future and so fail to notice that the conditions they are extrapolating have begun to change.

Two lines from Chancellor's book are on our minds today:

1. "*Most investors devote more time to thinking about demand than supply. Yet demand is more difficult to forecast than supply*."
2. "*The negative phase of the capital cycle is characterised by industry fragmentation and increasing supply."*

Over the next two years, supply is forecasted to increase at a staggering rate. The table below, sourced from SemiAnalysis, shows the data centre capacity they expect to be added each year:

Year	Operational + Under Construction (GW)[1]	Planned (GW)	Net additions (GW)	Cumulative since 2023 (GW)	Rough value of cumulative invested capital[1]
2023	5.8	–	5.8	5.8	$290bn
2024	9.8	–	9.8	15.6	$780bn
2025	15.6	–	15.6	31.2	$1,560bn
2026E	24.6	0.4	25.0	56.2	$2,810bn
2027E	39.4	11.9	51.2	107.4	$5,370bn

Before 2023, the idea of "gigawatts of compute" would have drawn laughs. To its credit, the industry rapidly added compute to meet overwhelming demand during the epoch's 'irruption' phase.

Let's run our finger down the numbers in the table above and focus on the last two columns. Last year, roughly 30 GW ($1.5 Trn of cumulative infrastructure capex) had been installed. Now, from this year onwards, **IF** (the operative assumption here) SemiAnalysis's forecasts are accurate (though we suspect these forecasts may come down, given recent events at NVIDIA, which we discuss at the end of this letter), total

[1] *Source: "Datacenters 2027: The Year of Acceleration, +50GW in a Single Year" (19 October 2026). Global ex-China; annual GW of IT capacity added. SaltLight rough estimate of $50bn / gigawatt cost.*

installed compute is expected to more than treble by the end of next year ($1.5 Trn to $5.4 Trn), and 2027 alone will add more than 2023, 2024 and 2025 put together.

This is good for some but bad for others in the sector, and this unprecedented infrastructure build-out *_should_* affect AI pricing. We find lessons from other industries helpful.

Should compute be thought of as a commodity?

There is old commodity lore: *"There is no cure for high prices like high prices."*

Close your eyes for a moment and imagine the table shows not data centres but dusty copper mines (we're old enough to remember the China-led resources super-cycle). The AI world is already adopting a *lingua commodity*: '*spot prices*', '*GPU pricing curves'*, '*contango'*, and '*backwardation'.* If it walks like a duck, as they say.

If the commodity business saw its mining capacity treble in two years, analysts would spend little time admiring rosy demand forecasts; they would anticipate a collapse in copper prices, even with demand at record highs. *Why?*

Like miners, compute owners are incentivised to maximise the return on their invested capital and run their assets flat out! In any commodity, the marginal cost of supply sets the price. At today's spot prices, compute owners such as SpaceX can rent it out on very short agreements at rates that pay back the hardware within a year or two.

We suspect the industry (and market) is projecting today's scarcity pricing well into the future, ignoring that capacity is roughly doubling each year.

Therefore, our base case is that AI pricing (measured in $/1m tokens) has likely peaked, allowing new value to accrue elsewhere in the AI stack. We return to this later in the letter.

Co-investors will likely ask us a good question: *what about the hyperscalers we own*? We think they are better placed than rivals selling unbundled, short-term compute. Partly because they already own workflows from enterprise customers or consumers, and cheaper tokens mean lower serving costs. Furthermore, because they can absorb capacity internally and distribute AI through existing enterprise contracts rather than competing purely on spot rates, they have a buffer against falling token prices, and we suspect that they have structured contracts to earn an adequate return on capital rather than maximise short-term profits. Even so, we will keep a close eye on them.

The Intelligence Quandary: Fruit Flies vs Adaptability

If data centres are the factory, the factory production is "intelligence" in the form of tokens. With Anthropic coming to market (and OpenAI presumably in the queue), we have recently had some time to think through the 'intelligence production' business.

We face two, really good, opposing arguments, and we are still on the fence as we watch how this plays out.

Model training is a treadmill

An AI model stops learning the day its training ends. Conversely, humans continue learning (one hopes) after paying for their education. Economically, a model company must keep investing in training.

Mathematically, training costs follow an exponential curve, which differs sharply from the linear model of training humans. With an exponential cost curve, you'd expect some enduring value from this training. Sadly, not. Because the model's knowledge is fixed, rivals can reproduce model differentiation easily through reinforcement learning and distillation. Therefore, on the one side, you have exponentially increased training costs and on the other side, you have a rapidly declining value after the model is released and competitors come in. This dynamic, unfortunately, is the standing cost of staying in the game rather than a foundation laid once and then enjoyed.

The other uncomfortable pattern is that the temporary economic benefits flow disproportionately to the most capable models… *for a short time*. Last year Google's Gemini 3 was, for a time, the leading model, and within months their advantage was competed away. Customers simply switched to the next best model.

Would-be owners of these model businesses must live with the fact that the principal asset, the trained model, has an economic life expectancy that marginally exceeds that of the fruit fly. Then again, as anyone who has left a bowl of peaches out over a warm weekend will know, fruit flies do not lack staying power as a species, *and they are not fussy about where they eat*. We suspect the model companies will have to behave in much the same way to capture enduring value.

Generalisable, adaptable intelligence enables a cheaper land-and-expand strategy

An opposing 'bullish' case can be made for the intelligence business. Fruit flies, as we have mentioned, are incredibly adaptable and can spread everywhere (under idealised conditions, a single pair of fruit flies could give rise to millions of descendants in just three months). Likewise, intelligence, bottled inside a large language model, may be the most *adaptable* technology product any of us has come across.

AI models can wander into the business adjacencies with ease, which, as students of competitive advantage, we find a little unsettling. Ben Horowitz, speaking on the All-In podcast, put it better than we could:

> *"Whatever you're holding in the equity markets, whatever you're holding, OpenAI and Anthropic are threats to that. Almost like there's almost nothing that you can hold in a portfolio that those two companies don't threaten. Because they're selling intelligence broadly, and every company is based on intelligence."*

So while a model's intelligence goes stale quickly, its natural adaptability may yet prove the advantage that lasts, and while the value of each training run fades just as fast, its makers may find they hold a structural advantage over those who build in the layers above them.

Perhaps a concrete example will help co-investors here. Why is it that enterprise software sold on a subscription basis has been such a marvellous business model, with its recurring revenue and handsome gross margins, yet so many of the fast-growing public companies still lose money? A large part of the answer lies in sales and marketing, which in a fast-growing software company can consume more than half of every gross profit dollar. Here are typical enterprise SaaS income statements illustrating how revenue and costs are distributed:

Scaling public SaaS (growing 20–40%)

% of revenue	
Revenue	100.0%
Cost of service	(20.0%)
Gross profit	**80.0%**
Product R&D	(25.0%)
General and administrative	(12.5%)
Sales and marketing	(45.0%)
Operating profit	**(2.5%)**

Mature large-cap public SaaS

% of revenue	
Revenue	100.0%
Cost of service	(20.0%)
Gross profit	**80.0%**
Product R&D	(17.5%)
General and administrative	(8.0%)
Sales and marketing	(30.0%)
Operating profit	**24.5%**

Moving from a scaling software business to a mature business is a tried-and-true process of what the industry calls the "land and expand" strategy. Land a software subscription inside a company by delivering a single product with a narrow use case. Then salespeople build relationships with company employees to learn what else the customer needs. This feedback helps the SaaS company build a solution, sell it, and find more customers with similar needs. *Rinse and repeat.*

Over time, the fixed costs of the enterprise sales force are spread across a widening product offering. It is, after all, far cheaper to sell to a customer you already have than to go out and find a new one (as every shopkeeper with a loyalty card has long understood).

Now, an adaptable *intelligence* product sold to the enterprise could turn this model on its head. What the customer buys is no longer a single product for a single need. It is closer to a Swiss Army knife, a tool that adapts to whatever the organisation needs from it, and organisations, we have noticed, are rarely short of needs.

SaaS product development economics also fundamentally change. When building a product, a software company wants a large market so it can spread development costs across it. With generalisable intelligence, the model provider can sell to a market of one. The customer spends their OWN time codifying their internal intelligence into their product. Anthropic's filing shows what this looks like in practice:

Anthropic Disclosed 3Q26 Income Statement Data

% of revenue	3Q26
Revenue	100.0%
Compute and infrastructure	(58.2%)
Gross profit	**41.8%**
Personnel (proxy for R&D)	(10.1%)
Other operating and administrative	(21.2%)
Sales and marketing (ex-partnerships)	(1.8%)
Cloud / channel partnerships	(12.1%)
Operating profit	**(3.3%)**

Source: Anthropic S-1 (3Q26 Data, SaltLight Adjusted)
Note: Cloud/channel partnership revenue is fees paid to distribution partners such as the hyperscalers

Therefore, once an intelligence product has landed, the customer drives most of the expansion, and the list of use cases to move into has no obvious end. Anthropic's IPO filing discloses 900% net revenue retention (NRR)[2]. What this means is that a customer acquired a year ago is currently spending $9 from $1 a year ago; an astounding example of 'land and expand'. Anthropic's income statement (3Q26) shows how this new business model has a structural advantage. Recall how traditional enterprise SaaS sales and marketing swallowed half of every gross profit dollar? The

[2] *Source: Anthropic S-1 link*

new intelligence business spends around 4% of every gross profit dollar on sales and marketing and expands rapidly.

So where does all this leave us? Still on the fence for now, if we are honest, though leaning in one direction. If compute behaves like copper (and we rather think it will), the price of a token has probably seen its best days, and today's high prices are already paying for the capacity that will cure them. That is poor news for those who rent out raw capacity by the hour, better news for the hyperscalers who can put it to work themselves (or at least have diversified long-term enterprise agreements), and good news for anyone who buys intelligence by the token. Who says the hyperscalers couldn't also capitalise on intelligence's go-to-market advantage? Microsoft has recently written about what they call the 'infinite SaaS Factory'[3]. We also believe there is value in having the context layer that intelligence uses to make proprietary decisions – not just general knowledge.

The model makers sit awkwardly in the middle. Their principal asset spoils faster than the fruit in the bowl, and they must keep paying to replace it, yet the economics of distribution are compelling. Whether that adaptability can outrun the treadmill is a question we cannot yet answer with any confidence. For a company like Anthropic, they should move up into the application layer to capture more value.

Today, we watch and wait, keeping our conclusions flexible.

******

Addendum: Jensen Huang and NVIDIA (Extracts from our May 2021 1Q21 Letter)

Last week, we heard the difficult news that Jensen Huang is leaving NVIDIA. Words cannot express how disappointed we feel. We first bought NVIDIA in 2020. In May 2021, we spent most of our 1Q21 letter on NVIDIA's AI opportunity (remember ChatGPT came out later in 2022). NVIDIA's market cap was roughly $300bn. At the time we said:

> *"As an aside, we think co-founder and CEO Jensen Huang is one of the most impressive CEOs that we've come across. It is customary amongst tech CEOs to have a signature clothing ensemble; Jensen is no different, and he frequently sports a 'glossy' black leather jacket. Since the COVID era, Jensen has presented all of NVIDIA's product launches from his kitchen*

[3] *Source: The Infinite SaaS Factory by Satya Nadella, Microsoft Link*



….

It is becoming clearer that NVIDIA's strategy is to use its leading position in GPUs (or what really should be called 'AI' hardware) to build an AI ecosystem by tying in hardware users into their software (middleware and applications). We believe that this integrated stack could build a considerable moat over the next decade.

…

NVIDIA's CUDA SDK is the primary software that interfaces with their GPU. As they have the most advanced GPUs, by virtue of being the first mover, scientists repurposed gaming GPUs for their particular use case. Seeing a new market of users, NVIDIA provided the software to maximise the performance of the GPU for free. As the AI opportunity has become more established, CUDA has become the de facto parallel processing software SDK used by developers.

This is a critical place to be to build a moat and 2.3m developers already use NVIDIA SDKs. But we've seen this movie before…"

He was, and is, an incredible visionary who made remarkable strategic moves to capture the early part of the AI epoch. Over the years, we have written about how NVIDIA's moat, particularly CUDA, was at risk of being eroded. In a way, its extraordinary success and high margins are what made the moat a target. AI is fast becoming good enough to write a replacement for CUDA itself. Warning bells went off for us when Google announced it was selling TPUs to Anthropic, because it signalled that Anthropic was investing in its own low-level code to run on competing chips.

We don't know what will happen going forward, and we think the range of outcomes is now much wider than it was a few weeks ago. Over the last few quarters, we had already substantially reduced our NVIDIA holding. Overall, our portfolio has not been significantly affected by turbulence in the AI infrastructure sector.

*****

As always, please feel free to get in touch with us should you have any questions. We always remind co-investors that most of our liquid wealth sits in the same funds as yours. We, therefore, share in the 'ups' and inevitable 'downs' alongside you.

Fundrock Collective Investment (RF) (Pty) ("FR") Ltd is part of the Apex Group Ltd. FR is a registered Manager of the Fundrock Collective Investment Scheme, approved in terms of the Collective Investments Schemes Control Act, No 45 of 2002 and is a full member of ASISA. Collective Investment Schemes in securities are generally medium to long-term investments. The value of participatory interests may go up or down, and past performance is not necessarily an indication of future performance. FR does not guarantee the capital or the return of a portfolio. Collective Investments are traded at ruling prices and can engage in borrowing and scrip lending. A schedule of fees, charges and maximum commissions is available on request. FR reserves the right to close the portfolio to new investors and reopen certain portfolios from time to time to manage them more efficiently. Additional information, including application forms and annual or quarterly reports, can be obtained from FR, free of charge. Actual annual figures are available to the investor on request. Performance fees are calculated and accrued daily based upon the daily outperformance, in excess of the benchmark, multiplied by the share rate and paid over to the manager monthly. Performance figures quoted for the portfolio are from Morningstar, as at the date of this document for a lump sum investment, using NAV-NAV with income reinvested and do not take any upfront manager's charge into account. Income distributions are declared on the ex-dividend date. Actual investment performance will differ based on the initial fees charge applicable, the actual investment date, the date of reinvestment and dividend withholding tax. FR retains full legal responsibility for the third-party named portfolio. Although reasonable steps have been taken to ensure the validity and accuracy of the information in this document, FR does not accept any responsibility for any claim, damages, loss or expense, however it arises, out of or in connection with the information in this document, whether by a client, investor or intermediary. This document should not be seen as an offer to purchase any specific product and is not to be construed as advice or guidance in any form whatsoever. Investors are encouraged to obtain independent professional investment and taxation advice before investing with or in any of FR's products. Access the FR Privacy Policy and the FR Terms and Conditions on the FR website (www.bcis.co.za). Full details and basis of the awards are available from the manager. SaltLight Capital Management (Pty) Ltd (FSP no 48286) is authorised under the Financial Advisory and Intermediary Services Act 37 of 2002. Trustee: Standard Bank Custody and Trustee Services Division is the appointed trustee.

The SaltLight Global Opportunity Fund is a separate portfolio from the SaltLight FR Worldwide Flexible Fund: while it draws on the same investment philosophy and research process, its mandate, currency, fees, tax treatment and risk profile may differ, and the performance of one fund is not indicative of the other. For the Minimum Disclosure Document, fee schedule, prospectus, and details of the South African representative appointed under section 66 of the Act, please reach out.