



2Q 2025 Co-Investor Letter

Calendar Year Results:

25 August 2025

	<u>Fund Return¹</u>	<u>SA CPI + 5%</u>
2020	+3.42%	
2021	+18.09%	
2022	-34.25%	
2023	+30.82%	
2024	+63.24%	
2025 - 1 st Quarter	+2.70%	
2025 - 2 nd Quarter	+24.52%	
Cumulative since inception	119.28%	57.90%

Dear Co-Investor

We had a pleasing second quarter, returning +24.52%, driven by positive contributions from **Roblox**, **NVIDIA** (re-entered during the "tariff tantrum"), and **Blue Label Telecoms**, which we think the market is finally starting to recognise the opportunity. At the end of June, we have 38% allocated to the US and continue to find opportunities in South Africa, China and Southeast Asia.

Forecasting where the general market is headed is notoriously tricky. There are compelling arguments for either direction.

Consensus among sophisticated investors at present is that the US market is due for a healthy correction. Frankly, we share this caution, and we are cognisant that most of the S&P's returns over the last couple of quarters have originated from the "AI Trade" by owning AI infrastructure and the Big Tech platforms.

On the other side of the argument, the right tail of the distribution seems ripe for exuberance depending on the trajectory of US interest rates, economic growth, and ongoing liquidity conditions.

Then there is a micro-debate that is more pertinent to our portfolio on what the AI opportunity set will look like and the market reaction thereof. For the last two years now, it seems the discourse in our letters is drenched in one overarching question: How Long Can the 'AI Trade' Go On?

At SaltLight, we decided long ago that we would not try to predict the market by positioning our portfolios for specific outcomes, whether up or down. Instead, what has worked well for us is to position a portfolio to manage a range of outcomes over five years—and let events unfold as they may be.

This thinking is threaded into the letter as we posit the scenarios on the longevity of the AI trade. We'll discuss in this letter: 1) update on our broad hypothesis of the AI opportunity, 2)

¹ SaltLight Worldwide Flexible FR Fund Class A1 return

risks, and 3) a recent investment in Google that will hopefully give perspective on how we're investing your capital within the range of outcomes we're facing.

Our Broad Hypothesis

Our overarching hypothesis is that AI is a multi-decade technology epoch spanning 40-60 years, akin to the railway age or manufacturing age. We use Carlota Perez's insightful work to unpack historical epochs into a discernible framework. We encourage interested readers to read her book, *Technological Revolutions and Financial Capital*².

Technology epochs follow similar patterns of evolution that typically unfold through two 15-20-year distinct phases. Within each phase, two microphases, each lasting a decade or so:

- **Installation Phase:** Marked initially by an "**irruption**" phase - such as ChatGPT's breakthrough moment, followed by a "**frenzy**" phase where the market recognises tangible value, enthusiasm swells, leverage builds, and inevitably, a speculative bubble emerges.
- **Deployment Phase:** Characterised by broader adoption, reconfiguration of industries, and genuine monetisation.

For historical context, Perez regards the Internet as one of several technologies of the Information Age, which began in the 1970s. The Information Age began with a three-decade *installation* phase, bisected by the typical *irruption* and *frenzy* stages, and then culminated in the Dotcom bust. Thereafter came the *deployment* phase that rose from the ashes in the early 2000s, and the mobile and internet era emerged as we have today.

The critical point for investors is the culmination of the *frenzy* stage. This moment brings a sharp bust that often damages both equity and debt tied to the epoch. It is precisely this phase where vigilance matters most to avoid a catastrophic fall.

Today's risks are higher than they were a year ago. Valuations are stretched, and volatility is on the rise. These are all signs of a late-phase market. At times like these, discipline is more crucial than ever. We stick to our 2030 hypothesis. Our portfolio construction approach guides us. We carefully size positions, diversify exposures, and only act where the expected return justifies the risk. This approach maximises our chances of capturing long-term returns while keeping downside risk in check.

AI is a Capital-Intensive Epoch


Market participants' concerns revolve around returns on invested capital. Big Tech's aggressive capital spending is the primary focus of the current AI narrative. This year's projected spending is expected to reach \$300–400 billion, around 1.3% of US GDP, and could rise by a further 20–30% next year. Many market participants misinterpret the purpose of these substantial investments. They only see conversational AI or software agents. However, these outlays are building strategic positions for a deeper future: robotics and 'physical AI', where artificial intelligence interacts directly with the physical world.

Many participants still hold strong views on the role of capital intensity in the Information Age. However, apart from the expansion of fibre and telecom networks, the previous era was

² *Technological Revolutions and Financial Capital* by Carlota Perez [link](#)

relatively capital-light. This is a stark contrast to the Railway and Manufacturing Ages, where investors readily backed capital-intensive expansion.

To put current developments into perspective, global data centre capacity today totals 59 GW³. Of this, only 7 GW is explicitly dedicated to AI workloads⁶. Notably, the industry now quotes AI data centres in power capacity (GW) rather than compute performance (FLOPS). Energy, not raw compute, is emerging as the critical currency of physical AI infrastructure. Brookfield projects a tenfold increase in installed AI data centre capacity, reaching 82 GW by 2034. This growth represents roughly \$7 trillion in total infrastructure spending over the coming decade. Of this, \$2 trillion is allocated to land and buildings, and \$4 trillion directly to compute, hardware and associated equipment.



Source: Brookfield internal research.

[Source: Brookfield](#)

These seem to be quite bold assumptions; however, what many are missing is that capex investments by sovereigns or companies have not yet taken place. Companies are concerned about data privacy, while sovereign nations might not share the same economic reasons and may instead see strategic value in maintaining domestic AI capabilities.

Narrative Misconceptions

There are two key misconceptions that we see as the underlying strategic rationale for spending so much capital:

Power: The primary 'gating' factor in AI now, at least in the US, is power. It takes an average of 24 months to build a data centre facility, but there is an *astounding* six-year interconnect queue to connect that facility to the power grid⁴. It could be argued that there is a strategic play, given these lead times, to invest in a "power" grab to ensure that they're not left behind.

Not just about GPUs: Another misconception is that headline numbers don't account for the asset's duration in the underlying capex. The physical data centre (the building, power, cooling, etc) will likely have 20-year lifespans (never mind the strategic value of power interconnect).

³ Goldman Sachs [Link](#)

⁴ Brookfield Whitepaper [link](#)

Next, around 40% of the capital expenditure is on the riskier GPU and network components, which are expected to generate a five-year revenue stream before becoming obsolete. Hyperscalers are keen to point out that the servers in their data centres are interchangeable across different workloads and can be easily upgraded.

Overall, these sources of demand will continue to invest in an AI vision, given the significant stakes for legacy and reputation. No tech founder wants to make the same mistake as Kodak by missing a generational shift.

From this perspective, for a tech founder, the greater risk is obsolescence rather than low returns on capital in the early stages of a multi-decade era.

In two years, there's a good chance that the initial years of capital expenditure on GPUs will be misallocated. A tech founder views this as “investing” or burning 4-5 years of free cash flow (equivalent to around 5-10% of their market capitalisation) as a call option to prevent disruption during a technology platform shift. They're willing to pay this price. The debate shareholders need to have is whether they're prepared to support that strategy, which could ultimately end the ‘AI Trade’.

Risks and Frenzy Phase Checkpoints

It all raises a question: isn't this a sign of frenzy?

In [our 4Q24 letter](#), we discussed typical “frenzy” phase markers that appeared in previous tech epochs:

- **Debt-fuelled capital investment:** Past epochs (e.g., railways, fibre networks) relied heavily on debt.
- **Speculative disconnect from fundamentals:** financial capital stampedes into the new technology, inflating assets beyond the earnings or productivity that the “new” economy can deliver
- **Mega M&A deals:** The classic financial engineering of merging old and new economies remains relatively subdued.

We'll briefly discuss each.

Frenzy Risk 1: Debt-Fuelled Capital Investment

Earlier this year, we pointed out that we hadn't yet seen the debt-fuelled investment typical of previous epochs' installation phases. Most recent AI capital expenditure has been funded by free cash flow from legacy businesses or venture capital. As a result, speculative risks have been relatively limited so far.

Now, we're starting to notice a growing awareness that debt funding is available for data centre builders.

Suppose the theory is that AI model training progression will continue along the predicted “scaling laws” trajectory. In that case, it will likely lead to exponentially higher capital expenditure on compute, memory, and power. Eventually, internal cash flows will no longer suffice, prompting firms to seek alternative sources of financing.

Private credit markets, backed by institutional investors rather than traditional banks, are set to take over. One of our portfolio companies, Brookfield, is actively engaged in this area as an alternative investment manager specialising in infrastructure projects. In their latest investor letter, Brookfield recently⁵ revealed a remarkable commitment of \$200 billion in capital over the next five years to just seven major projects: \$75 billion in the US, \$40 billion in Europe, \$20 billion in the UK, and \$50 billion in Canada. Although these figures are substantial, they will only cover 6 gigawatts (GW) of power capacity.

While staggering, the debt funding cycle still appears to be in its early phases.

The primary risk for our portfolio is that we are now on both sides of the customer and capital provider, and across the capital structure. We find reassurance in Brookfield's expertise in infrastructure finance. In their renewable financing, they typically mitigate risks by “hedging” out specific risks via off-take agreements or derivatives over the asset's lifetime. This results in a healthy yet low-risk cash flow stream over ten or more years. We anticipate similar strategies being applied to this opportunity set.

Brookfield has a distinct advantage, in addition to its funding, to deliver an energy solution for the power bottleneck. As the world's largest owner of renewable assets, it can also provide nuclear capabilities through its investment in Westinghouse.

Brookfield plans to offer GPU as a service, where it secures 4-5 year take-or-pay contracts with creditworthy counterparties and then owns or co-owns the hardware around the GPUs to capitalise on the strategic value of power, interconnect, and the lengthy build-out times.

We remain invested in Brookfield Corporation, which is diversified across asset classes and infrastructure assets. Although we don't directly take private credit risk, we benefit from the fees Brookfield earns in managing these assets. However, if Brookfield takes significant on-balance sheet risk on data centre credit, we may reconsider our investment.

Frenzy Risk 2: Speculative Disconnect from fundamentals:

This is the toughest risk to weigh *ex-ante*. Much of the speculative activity nowadays is happening in private markets, rather than public ones.

But then, how do we assess the disconnect from the fundamentals?


We'd argue that AI is already delivering significant productivity improvements in narrow areas such as software development, but the distribution of value remains very uneven. If SaltLight's anecdotal internal development is anything to go by, we pay roughly \$2,400 annually for a software coding agent, which replaces the employment of a \$75k mid-level software engineer. We gain substantial benefit, and it could be argued that AI software providers are not capturing enough of the value they generate.

However, we have seen this “tech” playbook before. Platform companies aggregate demand by offering a solution to a highly valuable problem at a cost that is vastly cheaper, quicker, and superior to existing options. They cultivate a drug-like dependence, whereby their users become completely “hooked” on the product, creating early signs of excessive consumer surplus and minimal shareholder returns.

⁵ Brookfield [link](#)

At this stage, well-meaning investors argue that this company can't create shareholder value. However, as the company scales and builds competitive barriers, it can raise prices, increase margins, and capture more value.

An anecdotal example is Uber, where most investors considered it a “terrible” business because drivers and consumers captured most of the value. Yet today, it's not a bad business. It has locked in its core customers and started to capture some of the value it's creating.



While Uber offers desirable services like convenience, AI delivers significant productivity improvements, and this habit will be very difficult to “kick” to revert to previous methods. In our software engineer example, the AI coding software provider could probably increase the price fivefold, and we would still see a considerable ROI.

At this stage, it's too soon to say if there's a speculative mismatch between fundamentals, so we offset risks through portfolio construction.

Frenzy Risk 3: Mega M&A Deals Between the 'Old Economy' and 'New Economy'

Nothing to note here.

Perils: A Possible AI Air Pocket?

Markets can be short-sighted, often causing significant volatility when valuations are stretched. Despite soaring AI spending, there is currently a clear gap between investment and tangible returns. The recent launch of ChatGPT 5 highlights this gap, showcasing significant advancements mainly focused on post-training improvements rather than groundbreaking innovations.

Notably, the training used NVIDIA's previous-generation Hopper GPUs, while the rollout of the newer Blackwell architecture is only now gaining momentum, with initial deployments in enterprise cloud infrastructure and specialised servers set to begin in mid-2025.

The more pressing challenge lies in enterprise adoption. AI agents, despite their potential, remain complex and challenging for widespread deployment. The supporting software ecosystem also remains underdeveloped, creating additional barriers to adoption. As a result, we anticipate an "air pocket" of uncertainty and volatility, testing our long-term investment hypothesis through 2030. While this may introduce short-term fluctuations, it also presents opportunities to acquire high-quality assets at attractive valuations.

Navigating the Long-Term with Near-Term Uncertainties: Google

To illustrate our approach to navigating these uncertainties, we turn to our recent investment in Google (Alphabet), which exemplifies balancing innovation risks with established strengths.

Innovator's Dilemma

Google is grappling with an Innovator's Dilemma as it protects its \$200 billion search business from a significant technological shift. To put it plainly, Google Search's primary purpose is to act as a 'match-maker', guiding users to the best source for their query on the open web. However, artificial intelligence is changing this role, with AI handling much of the searching, synthesis, and answering for the user, reducing the need to visit destination websites to gather information. A natural tension is emerging.

Humans naturally gravitate towards the path of least resistance, increasingly depending on AI to undertake cognitive tasks for them. This development poses challenges for content providers and for Google itself, which derives advertising revenue from these interactions.

Our bet is that Google is taking appropriate steps to address this Innovator's Dilemma:

- 1) Create a world-class foundational Large Language Model (LLM)
- 2) Differentiate on other vectors, such as cost and distribution
- 3) Value Accrues at the Application Layer

Considering the 'air pocket' risks we've identified, we view Google as a key enabler for advancing and participating in the spread of AI. If Google can protect or at least maintain its core search business (a big 'if'), we see significant upside for relatively limited risk, given the valuation multiple we paid.

Creating a World-Class Foundational LLM

It is worth acknowledging that Google's AI journey, as widely documented, has encountered numerous setbacks. We have added our voice to the critiques of their development efforts, even though they pioneered the transformative technology that now threatens disruption.

Yet, despite the wobbly start, Google has made an intensive effort to improve its models to become highly competitive. Our journey of researching the Google opportunity began when we started to use the Gemini models for our internal research agents.

Google has exceeded our expectations by releasing [Veo 3](#) for video generation (with audio!) and [Genie 3](#), which generates dynamic worlds on the fly.


Why do these matter? Both developments strongly suggest that their models are grasping real-world physics. We are inclined to view this as evidence of their data advantage, particularly from YouTube, in comprehending vision and audio, thereby offering substantial potential in the long-term Physical AI domain (remember Robots?).

Regrettably, we do not believe the model alone will resolve the Innovator's Dilemma, as we anticipate the foundational layer will ultimately become commoditised. Google's opportunity lies in de-commoditising other segments of the value chain, enabling it to compete effectively in this emerging technological epoch.

Differentiate Along Other Vectors

Google is distinguishing itself through vectors such as infrastructure costs and distribution.

1. Infrastructure Costs



Source: [link](#)

While much of the AI infrastructure narrative is on NVIDIA GPUs, Google's custom AI chips, called Tensor Processing Units (TPUs), are highly competitive and don't carry NVIDIA's 80% margins. Since 2016, Google has already produced six generations of TPUs and will shortly add its 7th generation (Ironwood) TPU.

There are significant advantages for Google that they can extract performance and cost gains by having their custom hardware and then an extensive internal and external captive customer base to scale their development.

Combined, they offer a cost-effective, high-quality, and low-latency product that sets them apart from other hyperscalers in terms of intelligence-per-dollar cost. Estimates suggest TPUs can reduce operational costs by 20 to 30 per cent compared to equivalent GPU setups.

We would not discount the possibility of Google marketing its TPUs to external clients, given NVIDIA's commanding market share and elevated gross margins.

2. Distribution and Services

Google's \$200 billion core advertising business is still vulnerable to significant disruption. This is perhaps the greatest threat to our hypothesis. That said, given Google's existing strengths – a cost-effective, standardised LLM model combined with seven platforms (each with two billion monthly active users) – it's well-placed to capture value by integrating intelligence into its established distribution channels. Imagine email services offering personal AI agents, YouTube serving up personalised ads, or Chrome, Google Maps, and Android all enhanced with AI capabilities.

Value Accrues at the Application Layer

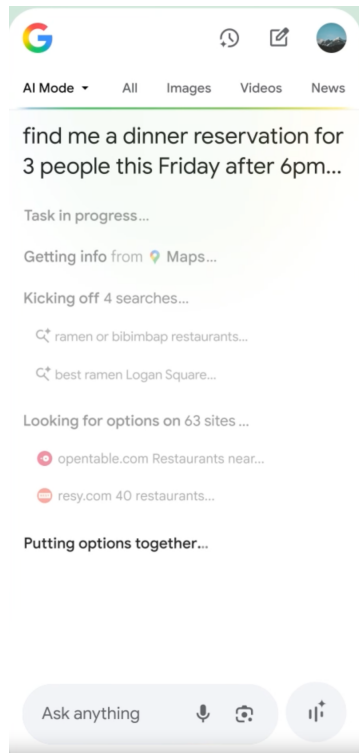
Regarding the core business, it's becoming increasingly clear that the search product will shift from helping users search and navigate websites (where revenue comes from click-through rates) to keeping users engaged within the Google interface (focusing on engagement time and Lifetime Value (LTV)).

Google will interact with the outside world on users' behalf through agents – for example, creating a holiday itinerary based on insights from Gmail data and booking flights and accommodation. The revenue stream may involve taking a cut from each transacting provider.

Most intriguingly, Google is in a favourable position due to its longstanding search operations on the open web. LLM-based search has become crucial to AI platforms to ensure responses stay up-to-date beyond the model's training data.

However, content providers are resisting this shift, as traffic no longer reaches their sites for monetisation; instead, information is generated by LLM models⁶. Yet Google already indexes this information for its search engine, placing it in a privileged position to deliver superior web-search responses.

⁶ For instance, Cloudflare is setting up a marketplace that requires AI bots to compensate content providers



The Banchan Table

4.6 ★ (512) · \$20–30 · Korean
2980 W Fullerton Ave

This welcoming Korean spot offers a taste of traditional home cooking, and is well known for its bibimbap, which diners say features a perfectly crispy rice bottom and a wide selection of fresh vegetables. [🔗](#)

Available on Friday, August 22



Source: Google⁷

Google Thesis Risks

That noted, elements of our thesis carry risks. Google awaits a ruling on remedies in the antitrust action initiated by the US Department of Justice concerning its search monopoly, with a decision anticipated imminently at the end of August. This could necessitate divestitures from specific business segments.

Additionally, in April 2025, a court ruled against Google in a separate AdTech antitrust case, finding it had monopolised open-web digital advertising markets.

And then there is the valuation-compression risk that occurred in China's Baidu, which experienced a platform shift as open-web activity consolidated onto Tencent's Weixin platform (Tencent being a portfolio holding).

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

⁷ Google: Video source: [link](#) and [blog post](#)



"Collective investment schemes are generally medium to long-term investments. The value of participatory interest (units) or the investment may fluctuate. Past performance is not necessarily a guide to future performance. Collective investment schemes are traded at ruling prices and can engage in borrowing and scrip lending. A schedule of fees, charges, minimum fees, and maximum commissions, as well as a detailed description of how performance fees are calculated and applied, is available on request from FundRock Management Company (RF) (Pty) Ltd ("Manager"). The Manager does not guarantee the capital or the portfolio's return. The Manager may close the portfolio to new investors to manage it efficiently according to its mandate. The Manager ensures fair treatment of investors by not offering preferential fees or liquidity terms to any investor within the same strategy. The Manager is registered and approved by the Financial Sector Conduct Authority under CISCA. The Manager retains full legal responsibility for the portfolio. FirstRand Bank Limited is the appointed trustee. SaltLight Capital Management (Pty) Ltd, FSP No. 48286, is authorised under the Financial Advisory and Intermediary Services Act 37 of 2002 to render investment management services."