



4Q 2024 Co-Investor Letter

Calendar Year Results:

27 February 2025

	<u>Fund Return</u> ¹
2020	3.42%
2021	18.09%
2022	-34.25%
2023	30.82%
2024	63.24%

Dear Co-Investor

The SaltLight FR Worldwide Flexible returned +20.35% in 4Q24 and +63.24% in 2024.

Year-end letters afford long-term investors to ruminate about where we are and where things could be heading. Too often, our industry's thinking succumbs to the tyranny of the urgent, mistaking fleeting headlines for enduring trends.

It's important to think in terms of decades, not quarters. It's also essential to get the big picture right. Communicating every quarter, we try to illustrate how our thinking evolves. Here is an illustration of how the AI Epoch has unfolded for us:

- [4Q20](#): We recognised that AI was entering a promising opportunity but highlighted the difficulty in picking early-stage winners and losers in AI models. Our view back then was that *"a more profitable approach in a nascent technological phase is to follow the wise adage of 'selling shovels to hopeful prospectors in a gold rush'".*
- [1Q2021](#): We wrote an in-depth letter arguing that NVIDIA is the optimal winner in this early phase of the technology, asserting that *"GPUs are structurally superior for AI problems"*.
- [1Q2023](#): As markets began their recovery, we framed AI as a defining technological epoch, pondering its long-term trajectory and the expansive scope of its impact.
- [1Q2024](#): We started to become niggly about valuations in the AI infrastructure companies that we had invested in prior years: *"Our investments in AI Infrastructure from three to four years ago have performed admirably for us; however, we became concerned with their high valuations and have pruned them significantly."*

Today, this once nascent AI opportunity is well-appreciated across the zeitgeist:

- **Market Dominance:** Big tech AI beneficiaries now have a market cap of \$19 trillion, 30% of the S&P 500's aggregate market cap². Including TSMC, that figure rises by another \$800 billion.

¹ SaltLight Worldwide Flexible FR Fund Class A1 return

² Apple NVIDIA, Microsoft, Amazon, Google, Meta, Tesla and Broadcom on 19 February 2025

- **Gigantic Capital Expenditure:** Last year alone, cloud giants spent \$216bn in capex, and they are projected to spend over a trillion dollars in the next three years, primarily on data centres and GPUs—key components in AI infrastructure³. A few weeks ago, OpenAI announced a five-year, \$500bn data centre project called “Stargate.”
- **Strategic National Advantage:** Global leaders have not ignored AI's strategic importance. Just weeks ago, an AI summit in Paris brought together 100 key stakeholders, including world leaders, to discuss AI's future. Concurrently, the US has been tightening AI export controls annually since 2022 to maintain its technological edge.
- **Valuation Concerns:** Despite these exciting developments, the valuation of the US market remains a concern, with the S&P 500 Shiller CAPE ratio mirroring peaks seen during the 2000 dot-com bubble and even the 1929 stock market crash⁴.

Given this backdrop, how are we at SaltLight thinking about the current market and the technology opportunity set ahead?

First off, we should always defer to four crucial quirks of financial markets: (1) outcomes are not always as relevant as are the shifting probabilities of the outcomes (2) all these expected value assessments need to be interpreted through the lens of what the market has priced in (3) every participant has different time horizons (ours is what the world will look like in 2030) and (4) the essential point, economies are complex adaptive systems, and we'll never be able to predict outcomes perfectly.

Our discourse considers some hypotheses about how things could shake out. We reiterate that we're most interested in the shifting odds of outcomes, and this is how we'll approach our letter. Furthermore, with new evidence, our minds could change on these views tomorrow, or we could be just plain...wrong.

- **The Bear View:** To be a bear where the AI bubble pops in the short-term, one needs to believe a) that the AI opportunity set is not real, b) that the financial returns will not materialise for a long time or c) compensate the capital provider for their cost of capital.
- **The Bull View:** To be a bull, with the knowledge that \$1 trillion (plus \$500 bn from Stargate and non-US data centres) will be spent over the next few years. This does not include the incremental energy capex being spent. The bull must believe in the prize: **AI companies will capture the value spent on human intelligence or see considerable GDP-enhancing productivity gains.**
- **The Uber-Bull View:** To be an Uber-Bull, one needs to believe that large language models (LLMs) are simply stepping stones into the real prize: Robots where intelligence operating in the digital world traverses into the physical world.

³ Consisting of Microsoft, Amazon, Google, Meta and Oracle using consensus estimates on 19 February 2025

⁴ Source: <https://www.multpl.com/shiller-pe>



At SaltLight, we strongly believe that the AI Tech Epoch represents one of the most significant multi-decade investment opportunities of our lifetime. Investors should not underestimate this opportunity. The unpleasant truth is that those not exposed to the cloud, mobile, and social network eras missed out on considerable aggregate returns generated in the US and SA markets.

Over the last year, regarding AI infrastructure, we think the odds have increased of a 'frenzy' phase where AI Bears have a point. However, if we put on our 2030 binoculars, we increasingly raise our odds of that Bull View scenario.

A few comments on our portfolio construction: Over the last year, the odds for AI software opportunities have been more favourable. Our strong returns, particularly in the digital advertising market, directly resulted from this. Going forward, our focus is on who benefits from overspending on AI infrastructure.

We draw lessons from previous tech epochs, which have guided our thinking about where to place our bets in 2030.

For those less interested in getting into the technical weeds, you may conclude your reading here; otherwise, join us down the rabbit hole.

In our letter, we thread our discussion around two inter-dependent topics:

(1) Carlota Perez's framework revisited: In our [1Q23 letter](#), we introduced Carlota Perez's framework to size the opportunity and then frame the adoption progression for better portfolio construction.

We believe it's crucial to assess our position within this long-term technology cycle periodically. Although we are still in the early stages, constant vigilance is needed as the odds of a potential 'frenzy' phase increase, during which exuberance can overshadow prudent investment decisions.

(2) The Human vs. AI Intelligence Competition and AI Agents: We argue that the rapid advancements in AI technology are not merely speculative; they are becoming increasingly tangible and impactful - a necessary requirement for the AI tech epoch thesis to hold up. We delve into how AI is progressing up the intelligence curve and then tie this idea into AI Agents, an investable opportunity we're very excited about. We aim to argue that AI is moving beyond theoretical applications to real-world utility.

Carlota Perez's Framework Revisited: Chapter II of The AI Tech Epoch

Two years ago, in our [1Q23 letter](#), we wrote extensively about why we thought AI would be the next technology epoch⁵ and that the ChatGPT “reveal” would be as transformative as the Wright Brothers’ first flight and Henry Ford’s rollout of the assembly line.

Our map on the journey then and now is Carlota Perez’s seminal book *Technological Revolutions and Financial Capital*. In essence, Perez’s book presents a theory that all tech epochs (The British Industrial Revolution, the Age of Steam and Railways, the Age of Steel, Electricity and Heavy Engineering) follow a similar trajectory: they are revolutionary, multi-decade, and multi-phase opportunities.

Scale: Revolutions, Not Lifestyle Enhancements

Perez makes the critical distinction that tech epochs are not individual innovations like tumble dryers or toasters that make our lives a little easier. They’re technology revolutions with clusters of innovations that bring an “*upheaval in the whole fabric of the economy... propelling a long-term upsurge of development*”. The innovation is so impactful that it precedes a fundamental reconfiguration of economic, social and regulatory systems.

To orient oneself to history, Perez believes that the last tech epoch was not merely the Internet but the “Age of Information and Telecommunications” starting in 1971. Just think about the scope of change from somebody looking at the world in the early 1970s. It would have been technologically impossible to send a 30-second, high-resolution TikTok video to billions of people over the airwaves within milliseconds. It was Sci-Fi fiction.

It is crucial to note that when thinking about this new epoch, there were several interlocking technological innovations (cheap semiconductors, software, networking, and fibre) within the more significant epoch that preceded massive socioeconomic change from these combined innovations.

Time: Multi-decade Opportunities With Similar Patterns

Perez suggests that major technological epochs can span four to six decades and comprise two primary phases: an ‘*installation*’ period and a subsequent ‘*deployment*’ period, during which the technology is mass-adopted.

⁵ *Technological Revolutions and Techno-Economic Paradigms*: “A technological revolution can be defined as a powerful and highly visible cluster of new and dynamic technologies, products and inspires, capable about bringing about an upheaval in the whole fabric of the economy and of propelling a long-term upsurge of development”




Figure 1 - Source: *Technological Revolutions and Financial Capital: The Dynamics of Bubbles and Golden Ages*

The installation phase is further subdivided into two micro-stages: an **irruption phase**, marked by a surge of initial enthusiasm and a **frenzy phase** where tangible value is observed. This phase attracts financial capital and likely leads to a divorce between paper and fundamental values.

Tangible Value Being Created

An essential ingredient is that the new technology exhibits signs of real value. One marker in the dot.com irruption phase is that from 1995 to 2003, labour productivity grew 3x the rate in the preceding decades⁶. This tangible evidence often results in financial resources being redirected from mature industries to new, dynamic 'new economy' opportunities. This transition causes a profound "techno-economic split" between the old and new sectors.

Typically, with a lower cost of capital, **this new paradigm fosters a wide range of new products, processes, and services while simultaneously revitalising existing industries.** This dynamic attracts further investment into the burgeoning 'new economy' (in the 1990s, IT investment was \$131bn, and by 2000, it was \$401bn⁷).

However, the installation phase faces a significant challenge: adoption rates remain low as the economic and social frameworks have yet to adjust to the new norms.

Perez identifies a 'decoupling' during this phase as market participants recognise that despite rapid growth in new sectors, these areas still comprise only a tiny fraction of the overall economy. Consequently, the returns from these burgeoning sectors do not yet correspond to the substantial capital invested.

⁶ *The Boom and Bust in Information Technology Investment*: "Labor productivity grew at an average rate of 1.4 percent between 1973 and 1995. Since 1995 through the end of 2003, labor productivity growth perked up to 3.1 percent" [link](#)

⁷ *The Boom and Bust in Information Technology Investment*: [link](#)

Frenzy Stage Risks

So where are we? Are we in the irruption or frenzy stage? The short answer is that we don't know.

We'd be confident in assessing that we are in the installation phase, but given the backdrop, we started this letter with, it becomes less clear whether we are in the irruption or frenzy stage.

There are several considerations for or against the 'frenzy' stage:

- Early innovations in one phase may be supplanted by an improved innovation in a subsequent phase (mainframe compute was overtaken by decentralised compute in the latter stages of the Age of Information).
- Great technologies for humankind don't always translate into great businesses (Warren Buffett quipped in his 2007 Letter: "If a farsighted capitalist had been present at Kitty Hawk, he would have done his successors a huge favour by shooting Orville down.")
- We have seen promising tangible benefits in some early use cases (such as advertising). Still, we have not yet seen the *raison d'être* outcome that the AI Bull needs - meaningful signs of productivity enhancements or labour substitution.
- In previous epochs, debt was extensively used to finance the installation phase (e.g., fibre lines and railroads). Today, most capital expenditures are funded from free cash flow or retained profits.
- A common sign in frenzy stages are the mega M&A deals and the mergers of 'old economy' and 'new economy' businesses. This is the financial engineering Wall Street engages in the old-new decoupling.

It's worth noting that profits (or meaningful revenue) were often absent from the installation phase in previous tech epochs. Instead, monetisation and profits usually appear in the deployment phase after mass adoption and wide reconfiguration.

So, while we are mindful of the frenzy risks in AI infrastructure, the result is that the cost of inference is declining by two orders of magnitude per annum. This will likely be setting the stage for new innovations to appear.

Remember, we are interested in new evidence that allows us to shift our odds of various scenarios. We see two new innovations that we think are vital to arriving on the pathway to real-economy viability: (1) AI's advance up the intelligence curve, and (2) the demonstrable capabilities of AI Agents.

The AI Intelligence Curve and AI Agents

To properly frame the AI Bull Hypothesis, we must identify what we think are the presuppositions to a successful end game: **AI should augment or supplant human intellect in the digital and physical worlds.**

The critical success factors are (1) intelligence must continue progressing up the intelligence curve to compete with humans, and (2) the cost of that intelligence must come down on raw capex and energy consumption.

If these success factors do not continue to improve, the odds increase that the AI Bears have been right all along – the return on the industry’s gigantic capital spending spree will not exceed the cost of capital.

In this section, we’ll argue that AI intelligence is rapidly improving and becoming a veritable competitor to humans. We’ll then discuss how AI agents use this incremental intelligence to the point where we think they are ready to be adopted by enterprises and see the necessary labour productivity improvements the AI Bulls need.

Reasoning Models: The Junior Employee and Test-Time Compute

It is worth setting the scene that AI output can be divided between training and inference to make sense of our discussion. *Training* an AI model is a costly upfront process to create the model. *Inference* is the run-time execution on the trained model when a user makes a request. Economically, training involves upfront costs for researchers and compute resources, while revenue is generated through inference as users pay on token usage⁸ or via a subscription.

Where are we today on the intelligence curve? To be crude, until a couple of months ago, we thought LLMs had reached the same level as a fresh-faced university graduate (book smart but needs guidance). But this has recently changed with reasoning models and test-time compute.

Until last year, the prevailing thinking was straightforward: intelligence scales log-linearly with the addition of more data and computing power, known as the “scaling laws.”

The charts below demonstrate that the data points closely align with a log-linear trend. Although they appear harmless, they represent logarithmic growth in computing power and data.




Figure 2 - Source: *Scaling Laws for Neural Language Models* (Kaplan, 2020) [link](#)

The implication is that an increase in these inputs is not just 10% more but 10 or 100 times more. For investors, these scaling laws represent astounding capital expenditure leaps per model version, from \$100 m for version 1.0 to \$10bn for version 3.0.


Aside from the staggering economic costs, there were concerns that intelligence gains were diminishing despite the substantial capital expenditure. Firstly, training was already saturated with the world’s high-quality text-based data⁹. Secondly,

⁸ A token could simplistically be equated to several characters in the input and output.

⁹ There is additional data that can be used from video but it is extremely costly to process. Furthermore, there have been additional gains from synthetic data generated from other AI models.

at inference, models had only one chance to provide an intelligent output—similar to answering a board game question with the first idea that comes to mind—limiting their capacity to reason deeply and effectively.

Then came the breakthrough: chain-of-thought (CoT) prompting at inference time¹⁰. By forcing models to articulate intermediate reasoning steps and giving them more time to “think”, researchers discovered that the intelligence output improved significantly¹¹ - called “test-time compute.”



Source: artificialintelligence.ai ([link](#))¹²

Essentially, the longer a model can think and reason, the more its neural network can yield sharper and more nuanced outputs.

AI practitioners will likely berate us for our over-simplification, but we wanted to convey a key message: **AI can now plan, reason and self-correct.** This is an essential ingredient for our discussion later on AI Agents.

For investors, the critical takeaway from reasoning models and test-time compute is that it drives up demand for AI workflows that capture more “intelligence share” from humans. Fortunately, it also changes how revenue across the value chain is distributed. Test-time compute queries can utilise 10-20 times the number of tokens compared to a traditional query (read: exponentially more revenue opportunities for cloud providers). This will go a long way in monetising significant investments for cloud providers.

All this put together; we think test-time compute has pushed up the odds for the AI Bull thesis that AI is dramatically moving up the intelligence curve (and starting a monetisation pathway for capex).

¹⁰ Models such as DeepSeek V1 are trained on reasoning rather than exclusively carrying out additional forward passes at inference time.

¹¹ The technical explanation is additional forward passes rather than single-shot passes.

¹² The Graduate-Level Google-Proof Q&A is a benchmark designed to evaluate the capabilities of large-language models in answering complex, graduate-level questions in different knowledge domains.

AI Agents



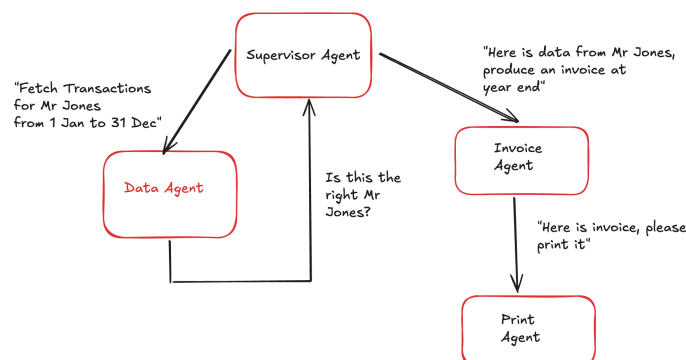
Figure 3 Scene from Wargames in 1983 (10)

Science fiction has been a wonderful tool for expanding the human imagination and innovation. Of course, writers can be too early: they have predicted moon colonies, flying cars, and household robots. But it often sparks an idea in a young Sci-Fi lover who turns it into reality.

AI, in the true sci-fi sense, is capable of thinking, planning, and reasoning independently of human supervision¹³.

We believe the ingredients for “agency” are here. As discussed, AI models can now **independently plan, reason, and self-correct** with little need for human supervision. Secondly, this intelligence has become cost-effective enough to be deployed in enterprises, where it will see tangible economic and productivity benefits **today**.

Deterministic Software to Agency



We often forget how deeply tied software is to human agency. Historically, every computer action was deterministic: if X happens, do Y—every possible scenario is spelt out in code. Humans remained the ultimate operators, exerting control through graphical user interfaces (GUIs). Whether it is an accountant printing an invoice or a

¹³ Source: <https://youtu.be/Ylh41wZEd5c?si=BvcLf8ewVEBzrrY1&t=147>

customer clicking “add to cart,” the code behind the scenes obeyed distinct, human-driven actions.

But here is how LLMs have changed everything: (1) new self-reflection and self-correction provide a ‘great leap’ in capability, (2) creating a robust, generalisable intelligence that slashes development costs and expands the total addressable market, (3) making domain data and systems easily accessible.

Self-reflection and Recursiveness

LLMs operate in natural language, humanity’s most refined tool for conveying high information value, enabling them to understand better and contextualise the world. AI feels like it genuinely gets what we’re asking because we interact with it on our terms rather than through the rigid constraints of a GUI designed for the deterministic world.

Beyond comprehension, LLMs can plan, act and then dynamically select the right tools for the job. Crucially, they possess the intelligence to self-evaluate and self-correct, iterating until they achieve the desired outcome. For your portfolio manager who has been coding for three decades – we cannot understate how dramatic this capability is.

Generalised Intelligence as a Primitive Layer

From a development cost perspective, the transformer (the fundamental architecture behind modern LLMs) can now inject every piece of code with a baseline level of intelligence. This fundamentally changes software development: intelligence is no longer painstakingly hand-coded for each scenario but instead generalised, slashing costs and accelerating deployment.

Tesla’s self-driving evolution is a case in point. Early approaches relied on deterministic programming: “See a red sign? Apply the brakes.” But reality is messy - no amount of if-else statements can account for every possible road condition. By shifting to transformers, Tesla replaced vast swathes of its C++ code with neural networks that learn from real-world driving, leading to a system with deeper contextual understanding and better decision-making.

A similar progression is emerging in robotics. Historically, robots thrived on factory floors, where tasks seldom change, but failed in unpredictable environments like a family kitchen. By embedding transformers that can understand images, language, and even unstructured contexts, robots can better adapt to varied environments. Reasoning and decision-making capability, coupled with large-scale data, make a profound difference in development economics. It will open a vast, addressable market for robotic applications.

Making Domain Data Accessible

Yet, most readers know that LLMs like ChatGPT have no idea what is happening in their data. An LLM is limited if it can’t access domain data. Unless the underlying connections exist, it cannot simply access an accounting database or understand how an enterprise operates.

Our primary research on agent orchestration platforms has revealed some interesting findings. LLMs are very good at understanding unstructured data and

adapting to the nuances of enterprise systems. However, the real magic happens when LLMs are given direct hooks - live access to data, applications, or real-world actuators (driverless cars or household robots). At that point, software is no longer just a passive tool waiting for human input; it gains agency.

The implications are profound. If an LLM can understand a request like “*Print Mr Jones’ invoice for December 2024,*” it can also infer which database to query, generate the invoice, and send it to a printer without a dedicated UI. This capability is **existentially disruptive** for the entire SaaS industry, which has long relied on domain-specific graphical interfaces (payroll systems, ticketing dashboards, point-of-sale terminals) as the bridge between humans and machines.

Historically, these interfaces existed because humans needed a structured way to translate real-world tasks into machine commands. **But what happens when an AI agent becomes the interface—eliminating friction, cutting costs, and streamlining workflows?** A vast swathe of SaaS businesses, predicated on providing a bridge between humans and machines, becomes redundant.

New Paradigm, New Impact

Imagine an AI-powered accounting system that doesn’t just wait for an accountant to request an invoice—it proactively checks the date, determines that an end-of-year invoice is due, and generates it automatically. Mr Jones can request it himself, or his tax AI agent can do so on his behalf. In this scenario, multiple AI agents from different applications coordinate behind the scenes—fewer seats to pay for and fewer clicks in a GUI.

Winners and Losers

It’s still early to call the biggest winners in AI automation, but we have an idea of who may lose out.

Cloud providers offering inference primitives (the compute layer that powers these agents) will likely be the first beneficiaries.

Then, we can speculate that “GUI” providers will have a tough time avoiding disintermediation. The business logic will likely shift down a layer to the AI agents, bypassing the GUI. The AI agents will directly interface with database storage, data lakes, or IoT devices. For those “GUI” businesses that don’t adapt, the ones that monetise on seats, their revenue stream is at risk.

In a defensive move, we wouldn’t be surprised if legacy enterprise software firms started locking down their APIs, turning once-free integrations into a gated, monetised service. Decades ago, software companies charged hefty fees for “connectors” that allowed different systems to communicate with each other.

The Broader Shift

It’s easy to say, “*We’ll keep the human in the loop.*” *Still*, a fundamental lesson of technology over the last century is that quicker, faster, cheaper - and *lazier* - solutions generally win. Of course, the reliability and trustworthiness of these agents will be a key consideration for enterprise customers. However, users consistently migrate to more straightforward ways of doing things, even if this erodes old business models.



The same logic that put text-based search on every smartphone now applies to back-office tasks like accounting or front-of-house tasks like customer service.

In the long run, this shift goes beyond chat interfaces. As AI agents become more capable of context awareness - seeing not just the immediate request but the bigger ecosystem in which it resides - they can proactively drive tasks from start to finish. That means reordering roles in the enterprise, from how we staff accounting departments to how we manage customer support. This is not about the end of humans but the end of drudgery. Repetitive tasks increasingly become the domain of AI agents, freeing humans to focus on higher-level responsibilities.

Frenzy Stage, AI Infrastructure

Back to our questions on the market, valuations and a focus on AI infrastructure. We'll reiterate - our focus is on 2030. What's vital for our thesis above is that intelligence goes up, inference costs go down, and the tooling for AI continuously reduces friction for the broadest possible realm of customers. We think the odds of that have improved significantly.

Another element that drives our conviction is our adoption of AI in our business. We have implemented rudimentary AI agents across our systems and have been wowed by the productivity and cost improvements. We think this epoch is the real thing, and we're pleased to be able to invest in it.

As we always remind investors, most of our liquid wealth is in the same fund as yours. We share in the inevitable ups and downs right with you. Please feel free to reach out if you have any questions.

David Eborall

Portfolio Manager



"Collective investment schemes are generally medium to long-term investments. The value of participatory interest (units) or the investment may go down and up. 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."