

Prophets Speak, Compounders Deliver - The Oracle's Vision and Our Rational Choices

GTF Letter 2

or·a·cle

['ôrək(ə)l]: noun

1. a priest or priestess acting as a medium through whom advice or prophecy was sought from the gods in classical antiquity
2. archaic. A response or message given by an oracle, typically one that is ambiguous or obscure

If one reads the Oxford dictionary definition of the word "Oracle", one would think that Oracle's prophetic characteristics or his/her obscure messages would have been strict to the ancient world.

In the beginning of September, the world was caught by another prophetic message. A certain cloud computing company published an eye-popping RPO - a proxy for its backlog. The number caught the attention of multiple heavyweight professional investors from all over the world. It became the topic of tech companies' managements and boards. A very well-regarded sell-side analyst even said that he never saw anything like this in his ~30 years covering software. Such expressions remind us of a miracle. And the most ironic part of all of this is the company's name: Oracle.

While many investors got euphoric - as per the 36% pop in the stock the following day - we were intrigued to say the least. Is this an opportunity? Or was this an opportunity and we missed it? Or is this a sign of a top to the AI story, which some even refer to as a bubble?

We believe that this can bring material consequences to investors, and we decided to publish our thoughts on this theme and the entire AI megatrend. This work was produced by Sao Pedro Capital's GTF team and therefore fits our investment process and mandate. More than having precise answers, we leverage our multi-decade experience investing and working in technology to provide a compass to our investors.

We encourage you to read this full letter to understand our analysis, but these are the main points and conclusions:

- For the past 2-3 years, Gen AI has been a great tailwind to technology, in particular semiconductor stocks like NVIDIA. The company delivered massive EPS and FCF generation, leading to 1,634% and 1,498% returns from 2022 to 2024. Stock price appreciation, +819% in the period, was well correlated to delivered earnings growth on healthy balance sheets. However, Oracle (ORCL) arguably launched a new phase of the Gen AI megatrend, likely using substantially more leverage to build new, massive data centers. Most of this additional capacity should be allocated to Open AI (OAI). While we've seen other players with high levels of debt (like Coreweave), ORCL's absolute numbers are just much larger;
- This seems to push the risk level as the incremental capacity is not only aimed at debottlenecking the current demand but is also based on what appears to be an aggressive OAI business plan that assumes hyper revenue (and demand) growth for the next 4-5 years. Oracle's leverage adds fuel to a correction if Open AI misses its forecasts;
- Despite the nice run, we believe good prospective returns still lay ahead for several Gen AI stocks and we continue to invest in this theme. Nonetheless, as our strategy dictates and we feel is often necessary in thematic investing, we are increasingly selective. We like companies that have a clear rational approach to capital allocation and sensitivity to overcapacity risks, like Microsoft. We also like businesses that are "selling shovels in the gold rush", and a few monopolistic derivatives, like NVIDIA and TSMC. We invest in

companies that have current, and material profitable businesses augmented by Gen AI, like GOOGL and META. We are avoiding stocks like ORCL and CRWV, which are levered names building infrastructure to serve "second-hand" or "relegated" demand;

- We have been increasingly allocating resources/time in themes uncorrelated to Gen AI infrastructure or misunderstood because of Gen AI. Healthcare, software/SaaS and payments fall into this category.

We hope this helps,

São Pedro Capital's GTF Team

Written By: Thiago Kapulskis, Jose Medeiros and Gabriela Moraes

Part 1: Oracle's Prophecy: An RPO Like No Other Even Seen


Early September, Oracle surprised the market with a number that few could have imagined. In its September 9th earnings release, the company reported that its **Remaining Performance Obligations (RPO)** had jumped to an extraordinary **\$455 billion**, more than **three and a half times higher** than the year before.

For those less familiar with the term, RPO is essentially the value of **future business already under contract but not yet delivered**. It includes invoices that haven't been recognized as revenue, as well as longer-dated commitments that should be billed and fulfilled overtime. In simpler terms: it is the size of the pipeline that Oracle has effectively "locked in" with customers.

Why did this figure suddenly explode? Oracle has been signing a wave of **multi-year, multi-billion-dollar deals** to provide cloud infrastructure for artificial intelligence workloads. Some of these are linked to the largest names in AI - most prominently **OpenAI** - and involve commitments that stretch years into the future. Taken together, these contracts have turned Oracle's backlog into something closer to what you might expect from an infrastructure provider, rather than a traditional software company.

Moreover, the sheer scale is unprecedented. A single quarter pushed Oracle's RPO to levels no other major software vendor has come close to reporting. Industry analysts and financial press alike described the number as "stunning" and "without precedent," underscoring how it reshapes the narrative around the company.

Figure 1: Oracle's Historical RPO




Source: Oracle, São Pedro Capital.

But it's important to understand what this actually means. **RPO is not cash in the bank**. It represents commitments that will be realized only as Oracle delivers capacity and as customers consume it. Some of those obligations may stretch many years into the future, and there are risks attached—both in terms of execution, customer follow-through, financing and return on capital. [Ratings agencies were quick to point out that Oracle's newfound backlog is concentrated on a relatively small number of massive AI deals, making it impressive but also potentially fragile.](#)

In short, Oracle's latest results gave investors a glimpse into a future flooded with contracted demand. The backlog is historic in size, yet it carries the uncertainty inherent in any prophecy: whether the future will unfold exactly as written.

Figure 2: Oracle vs Peers RPO/backlog - Oracle is now larger than all its peers*



Source: Company Reports, São Pedro Capital

Note on comparability: these are not identical metrics. Oracle and Alphabet disclose RPO (revenue backlog); Microsoft's figure is Commercial RPO (company-wide commercial contracts across Azure, M365, etc.); AWS reports AWS backlog. They all point to contracted, not-yet-recognized revenue, but coverage and scope differ.

Part 2: What does this mean for MSFT, NVDA and the entire AI theme?

Such story triggered a lot of work and a huge debate within our team. "Why did MSFT let this go?", "is this an opportunity?", "is this the beginning of a bubble?". These were just a few questions we had among ourselves.

Let's dig in each of the most important points related to this story.

2.a) Why did MSFT let this go?

We believe that this question is key to this entire story. As a reminder for anyone not familiar with the OpenAI and Microsoft relationship, the latter invested in the company in 2019, with a series of extensions and new deals. Here's the full timeline:

- **July 22, 2019 — first investment & partnership.** Microsoft announced a **\$1B** partnership with OpenAI and became OpenAI's **exclusive cloud** for training and running its models on Azure.
- **September 22, 2020 — exclusive GPT-3 license.** Microsoft received an **exclusive license** to commercialize GPT-3's underlying technology in its own products (separate from OpenAI's public API).
- **January 23, 2023 — "third phase" investment.** Microsoft announced a **multiyear, multibillion-dollar expansion** of the partnership (widely reported as about \$10B).
- **January 21, 2025 — partnership update.** Microsoft and OpenAI said they were evolving the partnership (e.g., working together on "Stargate"). [The Official Microsoft Blog](#)

Interestingly and as part of the January 2025 renegotiation, Microsoft and OpenAI moved from blanket Azure exclusivity to a model where **Microsoft has a Right of First Refusal on new compute capacity** for OpenAI - **with one key exception: OpenAI's API workloads remain exclusive to Microsoft/Azure**. Here is how it works in practical terms:

- When OpenAI needs **new compute capacity**, Microsoft gets **the first option to supply it** on Azure.
- If Microsoft **declines or cannot meet** the need on acceptable terms, OpenAI can procure that capacity **elsewhere** (e.g., Oracle) — and Microsoft has granted **waivers** so OpenAI can also **build additional capacity** (notably the “Stargate” initiative).
- **API workloads** (the traffic behind OpenAI's public API) **stay exclusive** to Azure, regardless.
- Separately, Microsoft holds an exclusive license to OpenAI IP (with carve-outs, including AGI-related areas), plus mutual revenue-sharing provisions for commercialization of OpenAI IP.

This means that the Open AI deal that ended up resulting in a great part of Oracle's RPO was analyzed by Microsoft first. So, why did MSFT let this exclusivity go? While there might be a few reasons, we believe that the centerpiece lies in MSFT CEO Satya Nadella's comment in the FY25 Q1 earnings call in October 2024:

Kash Rangan (GS): “Satya, when you talked about the investment cycle... models are getting bigger and more expensive, but monetization happens in inference. Where will that inference-phase revenue show up in Microsoft's P&L?” (full question paraphrased for brevity).

*Satya Nadella: “...we are not taking what I would— **there's a huge adverse selection problem today**... tech companies using VC money to buy a bunch of GPUs... we are **not** participating in most of that; we're going to the **real demand** in enterprises and our own products.” (excerpt)*

While Satya is probably not talking about Open AI (there are many small neoclouds and labs out there), there is an important hint here. Adverse selection is not an expression of software companies - it is used by **banks** and other financial entities like insurers. It means a situation when banks, due to imperfect information, end up lending more to bad borrowers than to good ones.


Interestingly, there are two ways to interpret Satya's comment. (1) He might be referring to supplying scarce resources (like GPUs) to customers whose demand may be speculative, unfunded, or low-quality. If you're not selective, you risk filling your capacity with the wrong kind of demand. However, (2) he may be simply avoiding supplying clients with an uncertain ability to pay for their commitments.

Let's analyze both angles. As one may see in figures 3 and 4, based on data published on The Information newswire, Open AI's business plan forecasts a very aggressive demand ramp-up over many years. It involves new products and more usage of ChatGPT - a killer app that already has ~1 billion monthly average users (MAUs), so not exactly an early-stage product. And such demand is not only necessary to fulfill the contracts but also generate revenues to pay for them.

From the perspective of cloud computing players like Microsoft (and Oracle), it's important to understand both risks (demand and credit) as they allocate capital on data centers needed to supply Open AI - similar to a bank allocating capital to provide loans. MSFT decided to leave such opportunity to others.

Moreover, the largest expense in data center capacity are GPUs which are inherently (somewhat) rapidly depreciating assets. This is not like fiber optic in the dotcom capital cycle, which received a significant overinvestment but was ultimately absorbed over the years as their obsolescence rate was much slower vs accelerated computing systems.

Figure 3: Open AI Revenue Estimates (Business Plan)




Source: The Information, UBS.

Interestingly, NVDA decided to back OAI up last week as it announced a USD100 billion investment in the company. So why did this happen? There are two main reasons in our view:

1. **NVIDIA could be supporting Open AI to improve their chances of fundraising elsewhere.** The USD100 billion comes as a key vote of confidence in the company and its aggressive plans. As we heard from Jensen Huang himself on [BG2's podcast this weekend](#) (on Open AI's plans to build at least 10GW of AI datacenters): "10 gigawatts is like \$400 billion... that \$400 billion dollars will have to be largely funded by their offtake... their capital, the money they've raised through equity and whatever debt they can raise... smart investors and smart lenders will consider all of these factors... There's the revenue side... and has nothing to do with the investment side... The investment side... is not tied to anything. It's an opportunity to invest in them."
2. **NVIDIA could be envisioning some "verticalization" to compete with other full-stack players.** NVIDIA has been partnering with Microsoft and Oracle to ultimately serve Open AI while the latter doesn't have its own data centers for some capacity. In a way, this group of companies compete with verticalized players Google (Gemini + Google Cloud + TPUs) and Amazon (Anthropic + AWS + Trainium)

Figure 4: Open AI FCF Estimates (Business Plan)




Source: The Information, UBS.

2.b) Could this be the beginning of a bubble?

Oracle took the opportunity that Microsoft left despite having a much more fragile situation. First, the company already starts off with leverage - it currently has a net debt to EBITDA of ~4x even before starting the heavy lift of building large data centers. Moreover, Oracle now has the largest RPO ("commitments") among all cloud players. This means that Oracle will have to build data centers in an unprecedented scale. While we believe they have capacity to execute, it is certainly a risk.

Figure 5: Net Debt to EBITDA - ORCL vs Peers



Source: Company Reports, São Pedro Capital.

So why would ORCL do this? We see a few reasons:

- Despite having a very profitable software business, it is no longer a “high growth” one.
- Moreover, part of the business face risks in the not-so-distant future, depending on what the future of AI brings.
- Also, they start from a distant 4th position in the hyperscaler league table. This announcement has immediately given them a big bump in credibility.
- Oracle’s founder Larry Ellison is competitive, incredibly smart and not known for conservative incrementalism... ORCL is his “treasury”

What happens if Open AI business plan doesn't go through - a new competitor comes in more aggressively (like Gemini or Grok) or AI adoption just takes longer to happen? In our view, Oracle brings significant more risks than peers. Microsoft, Amazon and Google are using their own cash and cash flow generation to fund their capex but Oracle will probably raise even more debt. Interestingly, this also allows another bold and historically aggressive player, who has been vying for a seat at the table to step up with more prominence – Softbank (and Masayoshi Son).

At the end of the day, we should see a new stage in AI, in which aggressive estimates and debt become part of the story. This adds fuel to risks - if things turn wrong, the downside is substantially higher, leading the asymmetry to the downside. On the other hand, at the risk of sounding like Ben Bernanke’s ill-fated prognostication during the Great Financial Crisis, we could add that this risk still feels somewhat “contained” as the bulk of industry investment still comes from healthy cash flow and not from balance sheet leverage. The conclusion is that stock picking seems key to equity investors in these conditions.

2.c) What is the read to the AI story - NVDA, TSMC, MSFT and others?

Despite all the abovementioned issues, we believe that Gen AI continues to be a great space to invest as it is likely to change the world in significant ways that we admit we can’t fully appreciate yet. However, we believe that active investing and a thoughtful approach to risk management is key to being successful in this new environment.


ORCL's numbers bring even more confidence to an investment in MSFT. It's likely that its ROI will be superior to ORCL's and if there is so much "relegated" demand, this means that MSFT's heavy capex deployed should not lead to overcapacity. Furthermore, MSFT now trades at a discount to ORCL (29 vs 35 P/E CY26).

NVDA has had the benefit of being an asset-light business model compared to the infrastructure players. As we wrote in the beginning, NVDA's market cap expansion has been very healthy as it was almost entirely given to EPS growth. Given that NVDA should be the most important supplier of silicon to Oracle, trading at a discount to it, it seems to be a better way to be exposed to ORCL's growth than ORCL itself.

However, earnings could contract prominently whenever such supercycle has a digestion. Moreover, NVDA's recent USD100 billion investment in OAI adds additional risks as it has now a financial commitment to support OAI. That said, this partially "locks-in" Open AI demand to NVIDIA GPUs.

Along the same lines, TSMC's growth also depends on such story. Still, a combination of a lower multiple, a more diversified client base (Apple, uncorrelated to this AI theme, is the #1 client) and a great track record of capital allocation are all pillars in favor of this name.

Figure 6: NVDA is not a bubble; NVDA EPS expansion and stock price increase are correlated



Source: IBES, LSEG, JP Morgan.

Material de caráter exclusivamente informativo e opinativo, elaborado pela equipe de gestão da São Pedro Capital. Não constitui, em nenhuma hipótese, oferta pública ou recomendação de compra ou venda de valores mobiliários, nem convite ou proposta para investimento em quaisquer fundos ou produtos geridos pela São Pedro Capital. As informações aqui contidas podem ser alteradas sem aviso prévio e refletem exclusivamente a visão da equipe na data de sua divulgação. Investimentos em ações e no exterior envolvem riscos, inclusive de perda do capital investido e variações cambiais.