wheresyoured.at

AI Bubble 2027

Edward Zitron

75–96 minutes

Soundtrack: The Dillinger Escape Plan - One Of Us Is The Killer

---

An MIT study found that 95% of organizations are getting "zero return" from generative AI, seemingly every major outlet is now writing an "are we in a bubble?" story, and now Meta has frozen AI hiring. Things are looking bleak for the AI bubble, and people are getting excited that this loathsome, wasteful era might be coming to an end.

As a result, I'm being asked multiple times a day when the bubble might burst. I admit I'm hesitant to make any direct timelines — we are in a time (to quote former Federal Reserve Board chairman Alan Greenspan) of irrational exuberance, where the markets have oriented themselves around something very silly and very, very expensive.

Regardless, Anthropic is apparently raising as much as $10 billion in its next funding round, and OpenAI allegedly hit $1 billion in revenue in July, which brings it in line with my estimate that it's made about $5.26 billion in revenue in 2025 so far,

The bubble "bursting" is not Sam Altman declaring that we're in a bubble — it's a series of events that lead to big tech pulling away from this movement, investment money drying up, and the current slate of AI companies withering and dying because they can't raise more money, can't go public, and can't sell to anybody.

Anyway, earlier in the year, a bunch of credulous oafs wrote an extremely long piece of fan fiction called "AI 2027," beguiling people like Kevin Roose with its "gloominess." Written with a deep seriousness and a lot of charts, AI 2027 makes massive leaps of logic, with its fans rationalizing taking it seriously by saying that the five authors "have the right credentials." In reality, AI 2027 is written to fool people that want to be fooled and scare people that are already scared, its tone consistently authoritative as it suggests that a self-learning agent is on the verge of waking up, a thing that is so *remarkably stupid that anyone who took this seriously should be pantsed again and again.*

These men are also cowards. They choose to use fake company names like "OpenBrain" to mask their predictions instead of standing behind them with confidence. I get that extrapolating years into the future is scary — but these grifting losers can't even commit to a prediction!

Nevertheless, I wanted to take a run at something similar myself, though not in the same narrative format. In this piece, I'm going to write out what conditions I believe will burst the bubble. **Some of this will be extrapolations based on my own knowledge, sources, and writing hundreds of thousands of words on this subject.** I am not going to write a strict timeline, but I am going to write how some things could go, and how (and why) they'll lead to the bubble bursting.

A Little Preamble Before The Premium Cutoff

For the bubble to burst, I see a few necessary conditions, though reality is often far more boring and annoying. Regardless, it's important to know that **this is a bubble driven by vibes not returns, and thus the bubble bursting will be an emotional reaction.**

This post is, in many respects, a follow-on to my previous “pale horse” article (called Burst Damage). Many of my original pale horses have already come true — Anthropic and OpenAI have pushed price increases and rate limits, there is already discord in AI investment, Meta is already considering downsizing its AI team, and OpenAI has done multiple different "big, stupid magic tricks," the chief of them being the embarrassing launch of GPT-5, a "smart, efficient router" that I reported last week was quite the opposite.

This time, I'm going to write out *the linchpin events* that will shock the system, and how they might bring about the bubble bursting. I should also be clear, and I will get to after the premium break, **that this will be a series of events rather than one big one, though there are big ones to look out for.**

I also think it might "take a minute," because "the bubble bursting" will be a succession of events that could take upwards of a year to fully happen. That’s been true for every bubble. Although people associate the implosion of the housing bubble with the “one big event” of Lehman Brothers collapsing in 2008, the reality is that it was preceded and followed by a bunch of other events, equally significant though not quite as dramatic.

One VC (who you'll read about shortly) predicted it will take 6 quarters to run out of funding entirely based on the current rate of investment, **putting us around February 2027 for things to have**

truly collapsed.

Here's a list of some of the things that I believe will have to happen for this era to be truly done, and the ones that I believe are truly *essential.*

- **NVIDIA's Growth Slows:** As discussed in the Hater's Guide To The AI Bubble, NVIDIA is the weak point in the Magnificent Seven (which collectively account for 35% of the value of the US stock market) specifically because of its success (19% of the Magnificent Seven's value) and value, which is driven entirely by its ability to sell more and more GPUs every single quarter. **It is inevitable that its growth slows, and once it does, the AI story crumbles with it.** Perhaps it has a down quarter then an up quarter — but that's three long months to make the markets happy.
- I will add that there's a chance the bubble reinflates a touch if NVIDIA crushes earnings. Or maybe the market doesn't care? We'll find out at the end of next week.

- **AI Funding Will Start Drying Up, And AI Companies Are The Most Funding-Dependent Companies of All Time:** As The Information wrote last week, the "dry powder" (available capital to invest in totality) has dropped, and venture capitalist Jon Sakoda of Decibel Partners believes that "...if VCs keep investing at today's clip, the industry would run out of money in six quarters." Removing OpenAI and Anthropic, the party will continue until 2028...but we all know those are the companies that are getting the money.

- **One of the major AI companies will collapse:** OpenAI and Anthropic both burn billions of dollars a year, and have shown no interest in stopping doing so, with Altman doing his best Lord Farquaad impression and saying that OpenAI is "willing to run at a loss" as long as it takes to get...somewhere. **For the bubble to burst, one of these companies has to die, and I will explain how this might happen.**

- In both of these cases, going public will be nigh-on impossible, and even if successful will expose what I believe to be their rotten economics.

- **Big Tech Will Turn On AI:** Meta's AI hiring freeze isn't enough: Meta, Google, Amazon or Microsoft needs to bring an end to their capex burn, and they need to be definitive that it's A) happening and B) that it's happening, at least in part, because they have "built enough" or "exceeded the opportunity in AI." These companies are barely making $35 billion in revenue from AI in 2025, and I doubt that revenue is increasing.
- **Another part of this: The Markets Make Big Tech Put Up Or Shut Up:** At no point have the markets really interrogated the revenue from AI. I can imagine some weird fucking games happening here. Microsoft, the only member of the Magnificent Seven outside of NVIDIA willing to talk about AI revenue, stopped reporting AI revenue in January, when it was at $13 billion "annualized" (so a little over $1 billion in revenue a month). One has to wonder if this means revenue is *flat* or *falling*, as annualized is monthx12.

- I can also see a scenario where these companies start putting out obtuse "AI-enabled" revenue stats. I do not think this works, and if it does, it only puts off the inevitable, **because they cannot avoid the capex crunch that's coming.**

- **This particular scenario is one that will happen in pieces.**

- **AI Startups Will Start Dying:** As I've discussed previously, AI companies are currently raising at suicidal valuations that make it impossible for them to ever sell or go public.

- **CoreWeave Will Die:** AI data center developer CoreWeave is a time bomb, burdened with, to quote analyst Gil Luria of DA Davidson, "deteriorating operating income." It's also a critical partner to OpenAI, providing compute as part of an $11.9 billion five-year-long contract, despite the fact that it's unclear if it has even come *close* to finishing its data center development in Denton, Texas. This has to be done by October, in part because that's when OpenAI is due to pay Coreweave, and in part because that's when it has to start paying back its massive, multi-billion dollar DDTL 2.0 loan.
- **The Abilene, Texas "Stargate" project fails to get off the ground...or OpenAI can't afford to pay for it if it does**. This alleged 4.5 Gigawatt data center will, when (if?) fully operational, allegedly lead OpenAI to pay Oracle $30 billion a year by 2028, which is more than OpenAI's combined venture capital and revenue to date. If this expansion doesn't happen it's bad, but even if it does, Oracle (unlike Microsoft) isn't going to accept "cloud credits," and for this to make sense, *OpenAI would need to be making upwards of $4 billion in revenue a month* ***and still have to raise a bunch of money.***

The common thread through all of these points is that they are **predominantly impossible to ignore.** So far this bubble has inflated because the problems with AI — such as "it doesn't make any money" and "burns billions of dollars" — have been dismissed until very recently as the necessary costs of the beautiful AI revolution. Now that things have begun to unravel, **the intensity of criticism will increase gradually, rather than in one big movement that makes everyone say "we hate AI."**

And it isn't just because of the money. CEOs like Tobias Lutke of

Shopify have oriented their companies' entire culture around AI, demanding in his case that "employees must demonstrate why AI cannot be used before requesting additional resources." Generative AI is, on some level, a kind of dunce detector — its flimsy and vague use cases having enough juice to impress the clueless Business Idiots who don't really engage with the production that makes their companies money. The specious, empty hype of Large Language Models — driven by a tech and business media that has given up on trying to understand them — symbolizes a kind of magic to these empty-headed goobers, and unwinding their "AI-first" cultures will be difficult...right up until the first guy does it, at which point everybody will follow.

AI has taken such a hold on our markets because it's symbolic of a few things:

- Executives and do-nothing middle managers' ability to control and suppress labor by suggesting a tool exists that can replace it.
- The future of automation, even though Large Language Models are absolutely terrible at it.
- The validity of the "ideas men" that run large parts of our society, who believe that their superior brains make them "above" labor somehow.

In any case, I am going to try and write the things that I think will happen, in detail. I'll go into more conditions in this piece, and as discussed, **I'm going to make some informed guesses, extrapolations, and give my thoughts about *how* things collapse.**

I predict that the impact of Large Language Models over the next decade will be enormous, not in its actual innovation or returns, but in its ability to expose how little our leaders truly know about the world or labor, how willing many people are to accept whatever the last thing a smart-adjacent person said, and how our markets and economy are driven by people with the most tenuous grasp on reality.

This will be an attempt to write down what I believe could happen in the next 18 months, the conditions that might accelerate the collapse, and how the answers to some of my open questions — such as how these companies book revenue and burn compute — could influence outcomes.

This...is AI Bubble 2027.

Table Of Contents

How I Am Justifying "Guessing"

Read any article about AI — for example, anything by Kevin Roose — and you will see constant guesstimates and extrapolations about "what will happen if AI gets powerful." As a result, I am justified in arguing the opposite: what happens if generative AI is a big, stupid waste of money, where everybody has joined in a ruinous, wasteful hype cycle behind a technology with terrible costs and limited outcomes. This is my newsletter, buddy!

NVIDIA, And What Will Happen Soonest

Due to the vibes-based nature of this movement, it's rather hard to tell exactly what might break the spell, but **the markets will dictate how long the AI bubble stays inflated.**

Even as I write this, tech stocks are climbing again, despite no significant sentiment shifts. As I'll say repeatedly, I'm not a stock analyst, so I am not going to tell you "how stocks work" or try to "time the market." Do not come to me for stock advice. I do not give financial advice. Help me, I hate writing about public stocks.

OpenAI and Anthropic, while "big" companies that get a lot of headlines, do not themselves move the markets. While their infrastructure is paid for by Microsoft (in OpenAI's case) and Amazon/Google (in Anthropic's case, though I believe Amazon handles most of its infrastructure, especially as it's building it a data center in Indiana), none of these companies intentionally break out the revenue they receive from them, in part because *doing so would likely reveal that they're paying themselves in cloud credits.* Microsoft also books OpenAI's compute — an estimated $10 billion this year — as direct Azure revenue, despite serving it at (or around) cost.

Microsoft and Amazon will next report earnings in the middle of October, but NVIDIA reports next week on August 27.

This is likely to give a temporary reprieve for those worrying about the bubble. NVIDIA will, likely, again crush estimates (around $45 billion to $48 billion in revenue). A year ago, its revenue was around $30 billion, which would make it likely it hit around 50-55% year-over-year growth.

This is likely to sate investors...to an extent, but a year ago its growth was *122%* year-over-year, which beat the previous year's 101% year-over-year growth.

Worse still, data center revenue, where NVIDIA books its GPU sales, (narrowly) missed estimates last quarter, making $39.1 billion from the segment, as opposed to $39.4 billion predicted. This accounts for 88% of Nvidia's revenue and is growing 73% year over year...which means growth is slowing year-over-year, going from 171% year-over-year in August 2023 (confusingly, NVIDIA marks its financial years a year ahead, so that's considered Q2 2024) to 154%

year-over-year in August 2024.

It's unclear how tolerant the market will be of a miss. NVIDIA comes with a certain level of exuberance. It makes a shit ton of money, it sells a physical good rather than access to models that burn billions of dollars, and Jensen Huang is a great salesperson.

Perhaps there's a scenario where things go back to normal for a bit. NVIDIA performs, the market says "oh, everything's okay" because the hyperscalers are still buying ass-tons of GPUs, and the desperate AI dance continues until another vibe-shift happens.

There's also a chance — not a huge one — that NVIDIA *misses* estimates, or its year-over-year growth looks stinkier than expected. The markets have forgiven misses before, and generally appear to see NVIDIA as untouchable, because the rest of the market is still buying the hype about AI, even though nobody talks about AI. NVIDIA has also halted production of its China-focused H20 GPU after the Chinese government told local companies to stop buying them, in part due to security concerns (such as the US putting tracking devices on some shipments), and in part because it wants to push them toward using GPUs built by local companies like Huawei.

To quote The Information:

> The production halt signals that despite the Trump administration's decision to allow Nvidia to resume selling the chips after a monthslong ban, the chip giant's hopes of maintaining its foothold in the Chinese market remain in limbo—now because of the Chinese government's policies. Chinese authorities fear Nvidia's chips could contain backdoors that funnel sensitive data from China to the U.S., The Information previously reported. As a result,

> Chinese authorities are encouraging local companies to use Chinese-made AI chips, such as those sold by Huawei.

This is, in a word, ***bad***. NVIDIA was poised to recoup as much as $15 billion in revenue from China in the second half of 2025 according to Laura Bratton of Yahoo Finance, specifically from the sale of its H20 chips after the Trump administration relaxed restrictions on sales to China. Despite the fact that NVIDIA is working on a new, more powerful chip for the Chinese market, it's clear that the Chinese government has made the decision to move away from NVIDIA entirely, possibly because of the 15% cut going to the US government, possibly for greater geopolitical reasons. Nevertheless, The Information also reported that a bunch of Chinese companies rushed in orders for 700,000 H20 chips following the reversal of the H20 ban.

Does this mean that NVIDIA is done selling chips to China? No. But it has the potential to hurt revenue for *the one company in the world that needs to continually grow.* I also don't know how much this will affect *this* quarter's revenue, or whether this is another can-kicking situation.

Nevertheless, NVIDIA Will Accelerate The Collapse If Its Stock Falters

As I've said, NVIDIA sells physical things — GPUs — and those GPU sales are the only real barometer of the "health" of the AI bubble, as nobody else is expressing revenue (and certainly not profits). When the DeepSeek situation happened earlier in the year, NVIDIA dragged the rest of the Magnificent Seven (and stock market) down with it, because NVIDIA's revenue is basically the only tangible value in the AI trade.

If **anything** happens to NVIDIA's value — a bad earnings, a DeepSeek-esque panic in the market (such as the launch of its R2 model, which is expected late August), or even just a dramatic vibe shift against the AI trade — it will drag *every* stock down with it.

NVIDIA's current price-to-earnings ratio — a standard way of measuring how cheap or expensive a stock is by measuring their price-per-share versus their revenue-per-share — is around 50 as I write this, which suggests that it's likely overvalued (I'm not a stock analyst, to be clear), and the market could decide, on a dime, that it needs a haircut.

Alternatively, investors could believe that NVIDIA has room to grow from here, which would mean in about a year it'd be selling $65 billion in GPUs a quarter, and upwards of $80 billion to $100 billion a year after that. This would require the Magnificent Seven — who make up 42% of NVIDIA's revenue — to buy more chips, every quarter, without fail, possibly until eternity.

I don't know if that's going to happen.

And remember: **without NVIDIA, the only other support for the AI bubble *is vibes.***

So, in simple terms, NVIDIA's earnings are the next big turning point for the AI bubble. If the market accepts its earnings report, we're in this for another three months minimum. If the market starts getting pissy, things will begin to accelerate, ***because all of the companies in the markets who have been saying they love AI will have to explain why.***

Big Tech's Bubble Burst Moment Will Be When Growth Slows

It is, at some point, going to become impossible for big tech to continue hiding their AI revenues, and I expect them to play some extremely silly games when they do so. We may see terms like "annual contract value," as ServiceNow has, to append AI to other revenue by forcing customers to buy it, or engaging with fantastical terms like "experimental recurring revenue" to refer to companies piloting AI software, a term that's actually bad but will be dressed up as good.

Yet they have a reprieve until growth slows, because business journalists do not have object permanence. Google's last quarterly earnings were hailed as "its AI bet paying off" despite at no point disclosing its direct revenues from AI, as did betting on AI "pay off" for Microsoft, which also did not disclose its direct revenues from AI. Big tech should, however, be concerned. Some are saying that Amazon's AI bets — despite them beating earnings — are *not* paying off...*because the stock had only gained around 3% year-to-date,* ***not based on anything to do with their AI bets at all.***

This irrational, ridiculous thinking is extremely vulnerable to vibe shifts. Amazon is not doing anything different to Microsoft, Google or Meta, other than the fact that Microsoft, for the first time, revealed its Azure revenue, the kind of thing you do when you anticipate people might be pissed off at you otherwise.

In reality, it's because Amazon Web Services' growth was flat year-over-year, and despite making *way more* than Microsoft Azure, the markets are pissy because the markets care about very little other than *number always going up.*

The entire AI trade has been based on make-believe that big tech's growth *came from AI*, rather than things like price increases (or by gutting volume discounts, as Microsoft has) and massive armies of salespeople built to continually nag and upsell every potential customer until they give in, which *should also make it very clear how dire the situation is with selling AI.*

This means that the *second* growth falters — the *moment* that one of these companies fails to meet the Rot Economy's demands for growth at all costs — there will be little mercy from the markets, **because there really isn't much of a story to tell about AI.**

Capital Expenditures Are The Next Thing To Go

Back in March, I wrote about how Microsoft had seriously pulled back on data centers, a story that has only grown more concerning since then, with Amazon joining it in slowing deployments. We are, I admit, yet to see this translate into anything. We have not seen a drop in capital expenditures, nor have we seen a drop in revenue to NVIDIA.

This makes me question what, exactly, is going on with its capex figures. Why have they not dropped yet? And why has this drop in capex not directly hit NVIDIA? 47% of Microsoft's capital expenditures go into NVIDIA chips after all, and 60% of its capex is spent on equipment for AI inference.

There also might be a simpler answer: Microsoft is buying for the capacity it currently has, retrofitting current data centers with Blackwell GPUs (which requires a totally different cooling solution), and at some point will finish filling that capacity and slow down its purchasing. This may mean that it takes multiple quarters before the data center pullback has a knock-on effect on Microsoft's capex, or NVIDIA's revenues.

At some point, however, capital expenditures *will* slow — there are only so many places on earth to put GPUs, and if said GPUs are not leading to actual revenue, it will be hard to justify buying more of them. Microsoft itself [said back in May](#) that capital expenditures would grow "at a slower pace in its upcoming fiscal year starting in July," which means it'll likely take a few more quarters to show.

And at some point, big tech will *have* to show AI revenue, or *have* to reduce that capex. When the latter happens, it will show a lack of faith in AI, which will drag down this vibes-based trade.

I expect there will be a period where big tech attempts to continue justifying the AI trade — pointing to obtuse metrics, reducing (but not eliminating) AI capex and promising "an era of efficiency" [as Mark Zuckerberg did back in 2023](#). Doing so will hurt NVIDIA, which will drag down the Magnificent Seven, and at *some point* the markets will ask what comes after AI — and I don't believe there's anything waiting in the wings, [as big tech has been out of ideas for a while](#).

At this point, I want to introduce you to the concept of **chaos bets.** These are things that I will say that are *guesses of unlikely but not impossible scenarios that could happen that would dramatically change things.* These are **extrapolations, and should not be taken as guarantees or facts. Also nothing I say should be taken as financial advice.**

CoreWeave — A Timebomb For The Markets and AI Writ Large

I'm going to be lazy and somewhat retread what I said above. [AI data center developer CoreWeave is a time bomb](#), burdened with, [to](#)

quote analyst Gil Luria of DA Davidson, "deteriorating operating income," raising a total of $25 billion in debt and equity since 2024, and currently sitting on liabilities of $18.8 billion, with expected capital expenditures in this year of over $20 billion on expected revenues of just over $5 billion for the year.

As a reminder, CoreWeave provides AI compute for inference and training of AI models. I hate talking about this company.

Coreweave is also a critical partner to OpenAI, providing compute as part of an $11.9 billion five-year-long contract, despite the fact that it's unclear if it has even come *close* to finishing its data center development in Denton, Texas. As I noted above: This has to be done by October, because that's when OpenAI is due to start paying Coreweave, and because that's when it has to start paying back their massive, multi-billion dollar DDTL 2.0 loan.

> **Sidenote**: This post is over 14,000 words. It's long. I'm going to repeat some points I made earlier, not because I'm a shit writer or because my editor's phoning in this edit on the August Bank Holiday, but because I'm throwing a lot of details at you, and they're all important. I don't want you to miss anything.

CoreWeave is a *fucking mess* of a company. Its infrastructure relies on a separate partner — Core Scientific, a Bitcoin miner which has never built AI data centers before, and has only done so for one customer...CoreWeave, which is currently trying to *acquire* Core Scientific in an all-stock deal, but shareholders are potentially voting against it due to the declining value of CoreWeave stock.

Recent filings have also revealed that Core Scientific *worried about running out of money* if it wasn't acquired by CoreWeave, its only customer, adding that "there was no guarantee that additional financing could be obtained" and that Core Scientific "would not likely have contractual commitment from new customers."

CoreWeave's reasoning for buying Core Scientific is simple: Core Scientific is its landlord for their new data center operations, saving CoreWeave "billions of dollars." The reason I am putting things in quotation marks is because the entire Core Scientific situation is extremely dodgy, and so is CoreWeave.

As Gil Luria of DA Davidson said in a recent note, "...**with operating income of $200 million and interest expense of ($267m), it appears that CoreWeave does not currently generate enough profit to pay all its debt holders, certainly not equity holders**." Luria also believes that CoreWeave will likely need to add $10 billion more debt this year to support data center expansion plans, and Core Scientific is critical to "scale capacity faster."

How Does CoreWeave Collapse?

We're doing AI Bubble 2027, so I'm going to hash out some scenarios that leads to CoreWeave's collapse:

The Runup

- If CoreWeave's Core Scientific deal doesn't go through, it'll likely savage the stocks of both companies, which will in turn make it difficult for CoreWeave to take on more debt, or to do an equity fundraise.
- Even if the deal goes through, CoreWeave's second big loan, which likely sits at around $8 billion, will come due in October 2025, just as OpenAI begins to pay its fees for compute...**which assumes that the compute will be available.** Having spoken to some infrastructure providers, there are doubts that this will be the case

— and plans are reportedly being made to provide redundancy if CoreWeave fails to deliver on time.

- Even if the compute is available, it is unclear whether inference or training loads are actually profitable for CoreWeave, which lost $290.5 million on $1.21 billion of revenue in its last quarter.
- CoreWeave's current revenue is mostly from Microsoft (and sources tell me that Microsoft's compute at CoreWeave is used by OpenAI) and a few other customers, making it *heavily* dependent on big companies' interest in generative AI.
- Remember, OpenAI picked up the option for compute that Microsoft didn't want.
- Everything comes down to the next two months: can CoreWeave buy Core Scientific? Can Core Scientific build the compute for OpenAI? And will OpenAI be able to actually pay CoreWeave?

The October Surprise

- Even if CoreWeave raises another $10 billion in debt, it likely has to sink every penny of it directly into building out the Denton, Texas data center for OpenAI. If this isn't complete in time, OpenAI will hit a compute crunch **and could potentially walk away from CoreWeave, or at the very least won't be paying it the money.** CoreWeave named OpenAI as its only counterparty credit risk on its S1/A, showing exactly how brittle this dogshit company is.
- The next few months are going to be chaotic, so it's possible that CoreWeave simply doesn't finish the Denton site in time, leading to OpenAI either delaying or entirely stopping its payments until it's done. CoreWeave *needs that revenue*, and any delay (or decay) in its revenue will only make things worse as we approach its mid-November earnings, though it won't have to report its October

revenues until early 2026.

- It's also not clear how much capacity CoreWeave actually has. According to Data Center Dynamics, as of August 2025 CoreWeave has "470MW of active power" across 33 data centers. Back in June 2024, CoreWeave signed a deal with Core Scientific to receive 200MW of capacity for AI, yet a July 2025 post about the acquisition said that Core Scientific had "840 existing gross MW supporting CoreWeave's HPC contracts," which is a lot more than CoreWeave apparently has active. CoreWeave loves to talk about "contracted capacity of 1.3GW" — but that just means "data centers we plan to build."
- *What the fuck is going on, exactly?*
- So, let's assume that the Core Scientific purchase goes through, and CoreWeave successfully borrows another $10 billion. Is that actually enough to successfully complete the Denton, Texas site? How long until that site is complete? How is it going, and indeed *what is going on there?* Have all the GPUs been bought?
- There is a scenario here where CoreWeave simply *runs out of money.* The Denton site is, from what I am told, behind schedule and a complete nightmare, as NVIDIA's new Blackwell GPUs require distinct specialized cooling, and this is quite literally Core Scientific's first rodeo. Its customers — Microsoft, Meta, and soon OpenAI — run services that *also* lose money, and are invested in AI based on vibes and market forces rather than any outcomes or revenue driven by the product.
- CoreWeave could, at any point, start running low on capital. In the event *anything* slows the development of the Denton site, CoreWeave will immediately lose the ability to grow revenue along

with its expenses, as OpenAI can (and will!) simply refuse to pay it, even if it has auxiliary support to cover the contract from another provider in place.

- CoreWeave's stock is also *heavily* vibes-based, as its stock value is driven by market hype around the growth potential of providing AI compute. If that narrative changes for the worse, CoreWeave's stock becomes less valuable as collateral for a loan, which in turn makes the company less likely to raise the capital it needs to continue providing services, which in turn lowers the value of the stock.
- If CoreWeave collapses, it takes with it a chunk of the AI narrative, **but also cuts into NVIDIA's revenue.** According to Gil Luria of DA Davidson, CoreWeave and other Neoclouds (like Crusoe, who is building OpenAI's Abilene Texas data center) make up around 10% of NVIDIA's GPU sales. NVIDIA is also an investor in CoreWeave, and also, for some reason, [a compute customer too].
- **Chaos Bet:** If CoreWeave runs low on funds, it could reach a point where it has to either choose to throttle operations for customers or simply stop building new capacity. If either of these happen, the company is toast: [OpenAI has contracted $11.9 billion of compute] with CoreWeave, and considering that OpenAI pays Microsoft over $10 billion a year for compute on Azure, and CoreWeave definitely doesn't have the capacity that Azure has, **it's very likely that CoreWeave has nowhere near enough compute to service OpenAI's contract.**
- If this is the case, OpenAI can (and will) walk away from the deal, and such an event will kill CoreWeave's customer story, almost certainly crash the stock, make it impossible for it to raise new debt, and bring about a sloppy, desperate end to the company.

- If this happens, this will unravel the public "AI is the next growth story" narrative for *every single company.*

In simple terms, CoreWeave is a timebomb, just like I said back in March.

The End Of AI Startup Funding

As I mentioned previously, The Information wrote last week that the "dry powder" (available capital to invest in totality) has dropped, and venture capitalist Jon Sakoda of Decibel Partners believes that "...if VCs keep investing at today's clip, the industry would run out of money in six quarters."

Sakoda adds that if you remove OpenAI and Anthropic, the money would continue through 2028... but let's be honest, those are the companies that are going to be raising the money, along with other bloated, overvalued AI firms.

The Information estimates there's around $221 billion left of dry powder in the hands of venture capitalists, and Sakoda believes that this will drop to $164 billion by the end of 2025.

The way VCs have gotten around their lack of capital is to use Special Purpose Vehicles — entities that can raise a chunk of money from a pool of investors — to pull together cash that would otherwise come from their Limited Partners and funds, allowing them to boost big bets without affecting their investment strategy (or pissing off LPs). Yet Anthropic, which is currently raising as much as $10 billion, has made it clear it will no longer accept SPVs, citing that their creators often have no connection to the company, taking fees for investors to participate — something that's *true* in some cases but a *peculiar* way to stymie the ability to raise capital.

Nevertheless, venture capital is going to eventually run dry at this rate, because (as I discussed recently) the valuations of AI companies are such that they're too expensive to acquire, and their low-revenue high-burnrate businesses are tough-to-impossible to take public. Venture capitalists are simply not seeing returns on AI investments, which in turn will make Limited Partners far less likely to want to continue backing these deals, *especially if the market sentiment shifts against AI.*

Right now, there are no profitable generative AI companies. Yet Cursor/Anysphere is valued at $9.9 billion, Cognition is valued at $9.8 billion, Cohere is valued at $6.8 billion, Perplexity is valued at $18 billion, Glean is valued at $7.2 billion, Runway is valued at $3 billion, and these have somewhere in the region of $82.5 million in monthly revenue **combined,** or roughly 55 times their annualized revenue of $990 million. These companies also burn ungodly amounts of money. According to Tom Dotan at Newcomer, Cursor is sending "100% of its revenue" to Anthropic, and Perplexity lost $68 million last year on $34 million of revenue.

It is inevitable that at some point venture capitalists will stop handing these companies money. Once one of them has trouble raising, the contagion will spread to the rest of the valley, which only really invested in these cash-burning monstrosities because many VCs are incapable of holding more than one idea in their head at any given time.

How The Collapse Of Funding Begins — A Scenario (Chaos Bet)

The following is a scenario that represents a guess as to how things

could collapse. I am making no promises. This is just one of the many ways this could go.

Perplexity, burning buckets of cash, is getting desperate, raising at a $20 billion valuation (yes, this is already happening), just over a month after raising at an $18 billion valuation, still in search of a meaningful business model that doesn't involve giving away tens of millions of dollars in refunds and discounts.

It successfully closes the round, and things look good...until two months later, when Aravind Srivinas starts calling VCs looking to raise at a $22 billion or $24 billion round. At this point, the "AI bubble" narrative has fully set into the markets, with public investors and LPs feeling a great degree of anxiety that the "rapid growth" of AI may never set in. Srivinas insists that things will work out, and that, in fact, Perplexity will grow faster than any startup not called OpenAI or Anthropic.

The round doesn't close. Perplexity is able to squeak by with a $20 billion valuation on undisclosed terms, and begins to run low on capital, at which point it's forced to raise a down round, which begins panicking the market.

> Sidenote: A down round is essentially when a company raises funds at a lower valuation than previous funding rounds. For obvious reasons, this is bad.

A story leaks out about chaos at the company as funds run low. Perplexity was never made to live frugally, and thus its messy, chaotic operations begin to sap the company of cash, as its "growth" has come at a massive, continually-increasing cost as the cost of inference continues to grow thanks to the proliferation of reasoning models. No amount of flailing around talking about buying other browsers (or TikTok) seems to work, and a half-hearted attempt at forging a partnership with Microsoft to "take over Bing" never makes it off the ground, along with another end run to get Apple to acquire Perplexity.

Perplexity begins rate-limiting customers and laying off staff as a means of lowering its burnrate, raising another downround as a means of keeping the lights on. VCs make it clear that they cannot fund at the previous valuation, as LPs are now screaming in their ear that they need to see a *profitable* AI company, and Perplexity only burns cash.

Meanwhile, Cognition is running into problems of its own, despite having raised an undisclosed sum of "hundreds of millions of dollars" in March 2025 and another $500 million in August 2025. It's burning hundreds of millions of dollars without making much revenue, with Windsurf's piddly $82 million annualized and $100 million in remaining cash burning away quickly under the combined losses of training its coding LLM "Devin" and keeping Windsurf's customers happy. An attempt to raise at $11 billion falls through, and Cognition raises a flat — or down — round.

This cycle continues. Cursor raises at somewhere between $15 billion and $20 billion, and apparently sits on $900 million or more in cash, but finds itself burning more and more each month despite aggressive price changes as Anysphere attempts to train its own models to escape the burden of Anthropic and OpenAI, both of whom increased prices on the startup.

We then begin hearing rumours of subtle layoffs throughout AI startups, kept hush-hush due to the sensitive nature of the bubble, framing them as "strategic decisions focused on efficiency."

At this point, something happens at Replit, which had negative gross margins in April 2025 and was forced to change its entire pricing model in June 2025. With many "vibe coding" customers, Replit finds itself constantly burning capital, the reasoning-heavy nature of their models smashing against the idiocy of people that don't know how to code typing "make app better please sir." Even with "effort-based" pricing, Replit burns millions of dollars a month, and its July 2025 funding of $250 million begins to run low.

Its $3 billion valuation also feels a little much for VCs that are now beginning to worry that *none* of their bets will pay off, especially with Lovable, a company that does basically the same thing, which raised $200m at a value of $1.8 billion in July 2025. Behind the scenes, Lovable faces the same problems as Replit.

Across the board, Subreddits complain of product degradation. Models don't feel like they used to, products are less efficient and burn more tokens to do the same jobs to a worse standard. Users begin to trail off.

As we approach the end of the year, layoffs have begun at some AI startups, and a few downrounds have begun to get VCs nervous. Dry powder begins to deteriorate — not necessarily because it's running out, but because VCs begin to get desperate. A "cash crunch" narrative begins to poison the valley, making even "positive" investments come with concerns around the viability of generative AI-backed companies, leading to more layoffs, less early investments, and at least one company simply being allowed to *die.*

At some point, Replit and/or Lovable are absorbed into OpenAI or Anthropic for a small or undisclosed sum, their customers effectively told to go fuck themselves and find another platform to burn money. These will be *branded* as positive acquisitions, but news will come out that it's a raw deal for investors and employees alike, *and people will start getting very, very nervous.*

Once the nerves fully set in, these companies will get desperate, seeking acquisitions wherever they can be found. In a buyer's market, big tech chooses, at best, to pick up one or two failing AI companies, but likely stays put, the markets frowning at any suggestion that they "need more AI."

Cursor, at this point, begins to run low on cash, unable to continue raising the hundreds of millions of dollars necessary to burn every year. Toward the beginning of 2026, Cursor is forced to start layoffs, unable to raise at anything approaching their current valuation.

It's unclear *who fails first*, but a startup will fail, because the venture capital funds will simply stop coming. We will see a cascade of desperation that takes the valley by storm, flushing away any chance of an escape through venture or acquisitions. As a result, we will enter a dark era of low funding, down rounds and death.

Costs Are A Brewing Scandal In Generative AI — Another Chaos Bet

Never in my career have I ever seen an industry given more grace than generative AI, with reporters gladly using figures like "$50 million annualized revenue" to make a company like You.com seem worth $1.4 billion, despite that working out to $4.16 million, or the revenue of four Dunkin' Donuts franchises.

To be clear, I am not talking about any specific startup when I say this, but something doesn't make sense about the burn rates on these companies, or at the very least, the only way it makes sense is if they're apocalyptic. Why does Cursor need $900 million, and then needs to raise more funding? Where is the billion-plus dollars of funding that went to Perplexity, and why does it need more of it already?

No, really, let's look at Perplexity. In July, [it raised $100 million in financing after], [according to Bloomberg], trying to raise $1 billion in March and, I assume, failing to do so. This leads me to two questions:

1. Why couldn't Perplexity raise $1 billion in March, and why did it raise a tenth of that in July?
2. What the fuck does Perplexity need this money for?

I believe that beneath the surface of these companies are incredibly bad costs, so much worse than anybody has accounted for. [While Perplexity's costs look "bad" based on The Information's reporting], with a net loss of $68 million, it's more than likely based on what it's raising that costs are both much higher *and worse than the reporting suggests.*

Remember, these companies (I am not saying this directly about any firm) have been allowed to run fast and loose with their costs, which more than likely means that *somebody* is undercounting or representing them in a more positive light than they should.

For example, Anthropic's gross profit margins don't include the costs of training its models, and Perplexity doesn't include the cost of its free customers in *its* gross profit margins. Playing silly buggers with accounting only suggests that things aren't good enough under the hood.

I am — non-specifically — suggesting that at least one generative AI company is likely misrepresenting the significance of their costs, and genuinely believe that at least one is overstating their revenues.

In any case, there're two very obvious companies that I believe are doing so: OpenAI and Anthropic.

The Curious World of Anthropic and OpenAI, And My Own Suspicions About Their Businesses

Questions About Anthropic's Revenue and Costs

A few weeks ago, I went over in detail how OpenAI and Anthropic leak annualized revenue as a means of obfuscating how much they're actually making.

> I have serious concerns about both companies' ability to do business as an ongoing concern, and writing out this story has made me truly worry about their business fundamentals. While OpenAI might *seem* big, it is only at around $5.5 billion in actual revenue for the year, with a dramatic climb needed to reach its projections, and an even more dramatic one to grow further.

In Anthropic's case, I do not see how it continues to grow at this rate, to the point I am actually concerned about the validity of its reported numbers. I won't go as far as to say they're wrong — just that if they're right, Anthropic is engaging in dark financial magic, rugpulling both consumers and its enterprise customers, showing a mercilessness rivalled only by their willingness to raise money from sovereign wealth funds and "dictatorial governments."

My concern about the validity of Anthropic's numbers is simple:

> Anthropic went from making either $72 million or $83.3 million in January 2025, to making around $333 million in June 2025, to

> making around $416 million in July. This is a *dramatic* leap, so dramatic that I genuinely would question its validity had it not been reported by multiple reputable outlets. Where did this money come from? New sales? New contracts? New products? That's more than 500% more than it was making at the start of the year!

Anthropic's revenue growth is peculiar. In July, it allegedly hit $5 billion annualized revenue — or around $416 million — with $1.4 billion annualized ($116 million) of that coming from their top two customers (Cursor and Github Copilot) and $400 million annualized ($33 million) of that coming from Claude Code according to The Information.

To which I say...*that's it?* People in the valley won't shut the fuck up about Cursor and Claude Code, they're two of the most popular products in AI coding, and Anthropic jacked up prices on Cursor in May. Where is this revenue growth coming from? Are you telling me Anthropic went from making $72 million in January to $116 million in February to $166 million in March to $250 million in May (no, we don't have April) to $333 million in June to $416 million in July?

The following is my gut instinct, and not based on any privileged knowledge: these revenues do not make sense to me in any way, shape or form. Anthropic has never, ever shown this speed of growth, and I am having a lot of trouble working out *how* it's even possible it hits $5 billion annualized so quickly, and indeed how it might potentially hit as much as $9 billion annualized. Were Claude Sonnet 3.5, 4 and Opus 4/4.1 *really that successful?* Were they so much more successful that Anthropic will, by the end of the year, increase its run rate by *1000%?*

The numbers don't even add up right. Say it's at $1.4 billion annualized from Cursor and Github Copilot, and $400 million annualized from Claude Code, that puts Anthropic at about $150 million in monthly revenue, or just over double where it was in January.

Where is the extra money coming from? Cursor and Copilot's revenue isn't new, either — they were both customers as of March, when Cursor was at $200 million annualized revenue and Anthropic was at $2 billion annualized (so $166 million in revenue). Did Anthropic...double its prices? That doesn't make sense either.

You're telling me that subscriptions to Claude and API calls — where Anthropic makes the bulk of its money — made Anthropic *$266 million in revenue* in July? And that number is expected to nearly *double* by December?

This would be remarkable — truly historic — if it's true. It would be faster than OpenAI. Faster than any startup, ever. And I am sorry, while I'll report as if these numbers are true, **I have genuine suspicions about the revenue numbers that Anthropic is sharing. They do not make sense.**

I believe that Anthropic is, on some level, misstating its revenues, because these numbers are truly obscene. The only previous justification I could find was from the price increases on Cursor and other companies and the launch of Claude Code — yet those numbers don't come close to explaining its growth.

How About Anthropic's Costs?

After it was reported it'd raise "up to" $5 billion, Anthropic is now raising as much as $10 billion, after raising $3.5 billion in March 2025, $1 billion in January from Google, and having previously raised

$750 million from Google, $750 million from other VCs in 2024, $4 billion from Amazon, and Google having invested $2 billion in 2023.

That's $13 billion dollars, the majority of it raised in the last year, with the potential for it to increase by *ten billion dollars,* bringing the total to $23 billion on revenues of, in totality, no more than $2 billion to date, and that's assuming you trust what Anthropic has leaked.

The Information reports that Anthropic burned $5.6 billion in 2024 and expects to burn $3 billion in 2025, at which point — **again, gut instinct** — I say "bullshit." How does that make sense? Even if Anthropic's 2024 revenues were at a full-year $1 billion runrate (they were not), that's...one billion dollars, and you lost $5.7 billion making it.

How the fuck are costs *lower* now despite Anthropic launching multiple reasoning-heavy (read: computationally complex and more expensive) and popular models? Anthropic allowed Claude Code users to burn so much compute that there's a leaderboard for it where somebody has burned $43,231 in compute on a $200-a-month subscription — and costs are going to be ***two billion dollars lower?***

And on top of that, Anthropic is raising $10 billion...for what reason, exactly? Because it wants a lot of cash in the bank for its lower burn rate? Or is it deliberately leaking lower numbers to muddy the waters?

Who knows, but I am not buying it.

An Important Note: Anthropic and OpenAI's Costs Are Dramatically Underrepresented Because Neither Company Pays For The Construction Of (Or

Owns) Their Infrastructure

This is a short, obvious one: **Anthropic and OpenAI's costs would be in the region of $30 billion or more a year if they were made to build their own servers, but Microsoft, Google and Amazon shoulder these burdens.**

As a result, their burn rates are *deeply ugly*, *and the value of their companies is overinflated.* Anthropic is "worth $170 billion" despite not owning a single part of its infrastructure. That's bad!

Questions About OpenAI's Revenue And Costs

As with Anthropic, OpenAI's revenue — that just hit $1 billion for the month of July, apparently! — is extremely strange. It is projected to hit $12.7 billion in revenue this year, yet by my estimates, it's only hit around $5.26 billion through the end of July.

It's done so by having 20 million paid subscribers, 5 million "business customers," and 700 million weekly active users, a number that has stayed consistent for over a month, even with the (disastrous) launch of GPT-5. OpenAI does not break out what differentiates a "paid" and "business" customer, but if we estimate that "paid" means "plus" ($20 a month) and "business" means Teams ($25 a month annually or $30/month monthly), Pro ($200-a-month), and "Enterprise" (starting at $60 a month, though OpenAI is offering discounts to close deals), the math becomes questionable.

"Business" also includes education.

- **Plus:** 20 million people paying $20 a month is **$400 million a month in revenue**.
- **Teams:** I am going to guess there are 3.75 million of them, **or**

$131.25 million in revenue.

- **Pro:** As of January 2025, "Pro" users were contributing $25 million a month in revenue, which works out to be 125,000 paying customers. **I'm going to be generous and double this number to $50 million, so 250,000 of them.**
- As an aside, this article suggests that "Pro" customers are not included in the "business" users. I don't think this changes the math in OpenAI's favour by very much.
- **Enterprise:** This leaves us with an overly-generous 500,000 $60-a-month enterprise customers, **or $30 million in revenue.**
- **Education:** From what I can tell, OpenAI has made about $17.8 million in total on contracts in 2025 from educational institutions. **Nevertheless, at least 500,000 of the 5 million business users are paying an effective $2.50 a month.**

That's only $629.05 million, and I believe I'm massively overcounting the revenue from Enterprise and Pro customers, leaving us with $370.5 million of monthly revenue from OpenAI's API customers... Which sounds unlikely, as it'd work out to roughly $4.3 billion in revenue for 2025, or more than twice the amount that OpenAI projected to make ($1.8 billion) this year from API calls.

These numbers also *only ever seem to trend up*, with OpenAI leaking in January 2025 that it had hit 15.5 million paying subscribers in December 2024. That's a 29% increase! That's whacky!

I'm going to be honest, something doesn't make sense here! How is OpenAI's subscription churn? According to The Information, 77% of OpenAI's paying customers are retained after 6 months (as of June 2025), meaning that OpenAI is effectively adding a ridiculous amount of customers every month — effectively meaning that

OpenAI has had to add over a million new paying ChatGPT subscribers every single month to have hit 20 million by April to keep up with churn, with as many as 1.5 million added in March (20 million was announced on April 1 2025).

Remember, the calculation is not just "OpenAI added 4.5 million subscribers." That churn likely means that 3.8% of users — hundreds of thousands a month — are canceling their subscriptions. These aren't *bad* churn rates, but I am *astonished and suspicious* at how fast OpenAI has been able to sell subscriptions, to the point that it doesn't really make sense.

OpenAI Is Using Cal State's "Edu" Contracts To Pump Its Paid Business User Numbers — And It's Unclear How Long It Keep A User In Its Numbers

What also doesn't make sense is that despite all of the headlines, OpenAI doesn't seem to want to update this number, instead choosing to say it had gone from three million paid business users a few months ago to five million paid business users as of August 2025 — a jump that still feels confusing to me but feels *plausible* with the addition of Education contracts like the Cal State system, which added another 500,000 seats to its deal with OpenAI in July 2025, costing $15 million for the year, **which works out to $2.50 per-seat-per-month.**

In fact, LAist laid out the revenues pretty nicely:

California State University's OpenAI deal

How CSU's deal with OpenAI compares to other schools' contracts

Campus	Total cost of deal	Number of users	Length of contract	Cost per user	Notes
Arizona State University	$300,000	1,000	12 months	$25	

California State University	$1.9 million	40,000	6 months	$7.92	This contract was in place from Jan. 1, 2025 to June 30, 2025.
California State University	$15 million	500,000	12 months	$2.50	This contract went into effect on July 1, 2025.
University of Massachusetts Lowell	$21,000	350	2 months	$30	This contract was in place from June 19, 2024 to Aug. 18, 2024
University of Massachusetts Lowell	$126,000	350	12 months	$30	This contract went into effect on Aug 19, 2024.
University of Nebraska Omaha	$360,000	1,000	12 months	$30	
University of Texas at Austin	$126,000	Undisclosed	Undisclosed	Unknown	

Source: Public records requests • Credit: Julia Barajas and Ross Brenneman | LAist

While most of these deals are priced at the regular Teams pricing of either $25 or $30, the largest of them — the Cal State deals — are at $2.50 or $7.92 a month, or a fraction of the $20-a-month a regular Plus user would pay.

To put this more bluntly: **10% or more of the 5 million "paid business users" that OpenAI states are paying $2.5 a month.**

What About Government Contracts?

From what I can find, OpenAI has signed a few different contracts — $108,000 of licenses to NASA, $70,000 to the National Gallery of Art, and $100,000 to the US Department of Agriculture, all in 2024. It also signed a $200 million contract with the Department of Defense "with an estimated completion date of June 2026."

I also think it's worth considering if OpenAI considers "Government" a "business user," as it's giving away ChatGPT for a dollar-a-year to the US government.

OpenAI Is Using "$1 for a Month" Subscription Deals On Teams To Juice Business User Numbers — And Offering Deals For ChatGPT Plus For $10-a-month To Stop User Churn (REMINDER: OpenAI Loses Money On Every User Anyway)

When GPT-5 launched, I went to try and access it via purchasing a "Teams" subscription (so that I could see whether there was still access to legacy models), only to find that OpenAI was running a special deal where I could get a month of Teams for a dollar.

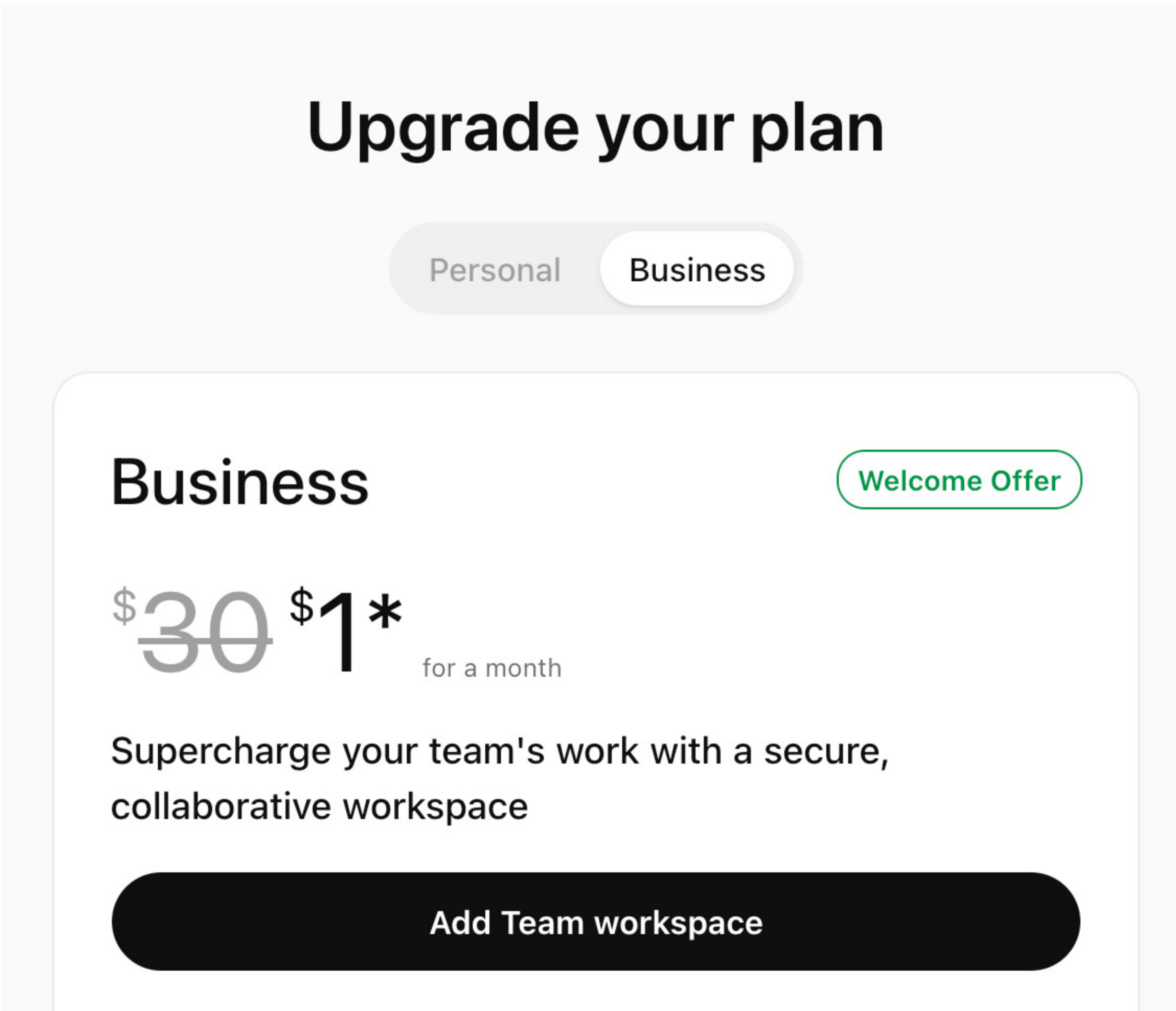

I do not know the scale of these offers, nor how long OpenAI has been offering them. A Redditor posted about the deal a few months ago, saying that OpenAI was offering up to 5 seats at once. In fact, I've found a few people talking about these deals, and even one adding that they were offered an annual $10-a-month ChatGPT Plus subscription, with one person saying a few weeks ago saying they'd seen people offered this deal for canceling their subscription.

So I went to try and cancel my ChatGPT Plus subscription, and *wouldn't you know*, OpenAI offered to drop my subscription from $20 a month to $10 a month for three months.

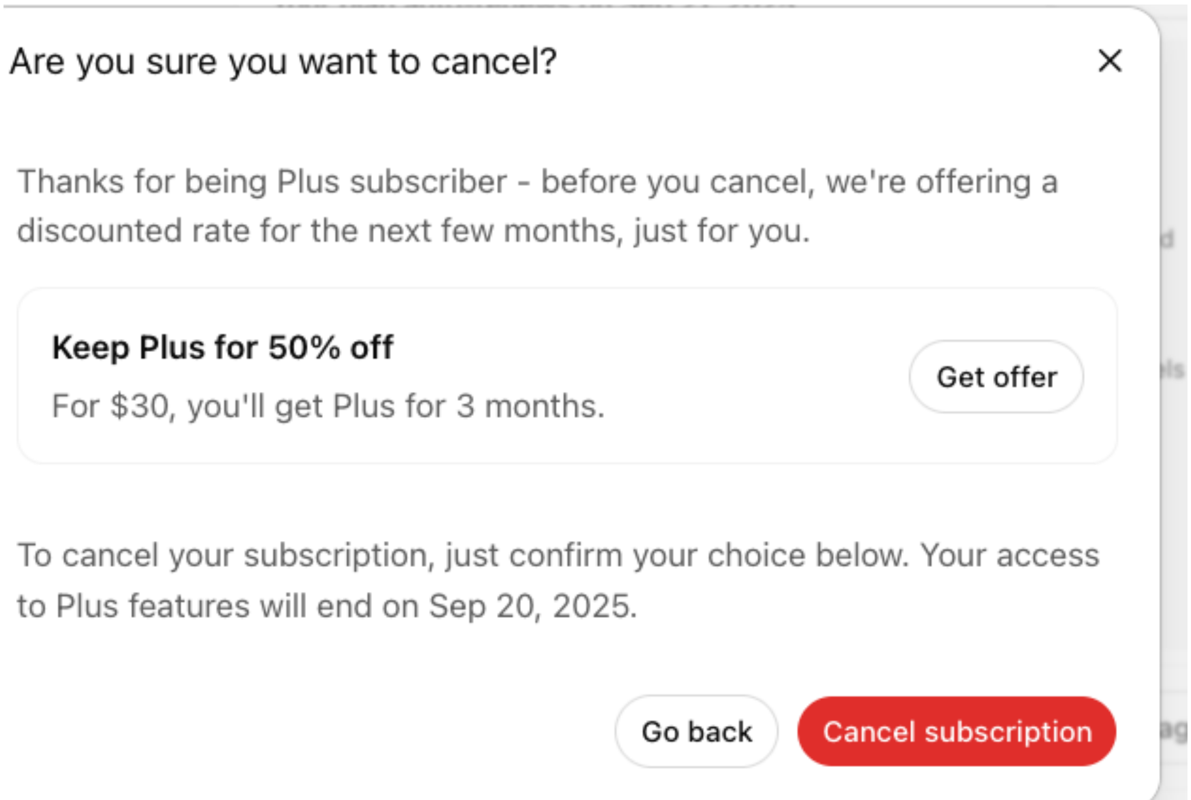

There are plenty of people saying they've been offered this deal, and these posts all are fairly recent — 10 days ago as of writing this sentence — with reports going back as far as two months ago. While these are standard churn tactics for a software company, **they make me even more suspicious of OpenAI's stated user numbers and leaked revenue numbers.**

I Believe OpenAI's Costs Are Worse Than They Seem — Where Is All The Money Going?

This is, to be clear, an opinion, but OpenAI's supposed (IE: leaked)

"$8 billion burn" in 2025 sounds like complete horseshit to me.

Last year, OpenAI burned around $5 billion on less than $4 billion of revenue, and is projected to — though I am not convinced it will — hit $12.7 billion in 2025.

In 2024, OpenAI raised $6.6 billion in equity funding and $4 billion in debt for a total of $10.6 billion, along with $250 million from ARK Investments. Since then, it's had to raise $11 billion in June, then another $8.3 billion in August, wants to raise at least another $20 billion from SoftBank by end of year (assuming it converts to a for-profit, which I'll get to in a bit), and has told investors it wants to raise another $17 billion in 2027.

Where is all the fucking money going? Where is it going, Sammy?

No, really! Even if OpenAI *makes* $12.7 billion — which I am not sure it will! — it clearly *needs* to raise this much money. Where is it going? It isn't going to the Abilene Texas data center. Oracle, Crusoe, and Blue Owl are funding **the entire operation**, and Microsoft (and CoreWeave!) are paying for the rest of OpenAI's infrastructure.

In the space of a year, OpenAI has raised *$26.15 billion,* and claims that it'll only lose $8 billion. Fuck you! Fuck off! That's nonsense! I haven't even included the $4 billion in debt! And now these chucklefucks want to raise *another $20 billion from SoftBank?* Where is the money going?

Why did OpenAI have to raise another $8.3 billion so quickly after the $11 billion closed? I get that there's a level of "let's get this done, why wait," but at the same time...what's the rush? Why do you need all this cash? Where is it going?

I am not suggesting anything illegal: **I am simply stating that I**

believe OpenAI's costs are likely very, very high, much higher than we are currently aware of. There is no rational reason why this company — which lost $5 billion in 2024 — needs to raise a total of $26.15 billion *and a four billion dollar credit line*. Stargate LLC has not been formed, no money has changed hands, and based on reporting, Clammy Sam Altman has cunningly convinced everybody else to pay for his infrastructure buildout.

OpenAI Will Burn At Least $3 Billion On Salaries In 2025, And May Spend As Much As $8 Billion — With Any Layoffs Guaranteeing An Industry-Wide Panic

Let's take a look at The Information's reporting on OpenAI's projected costs from 2024.

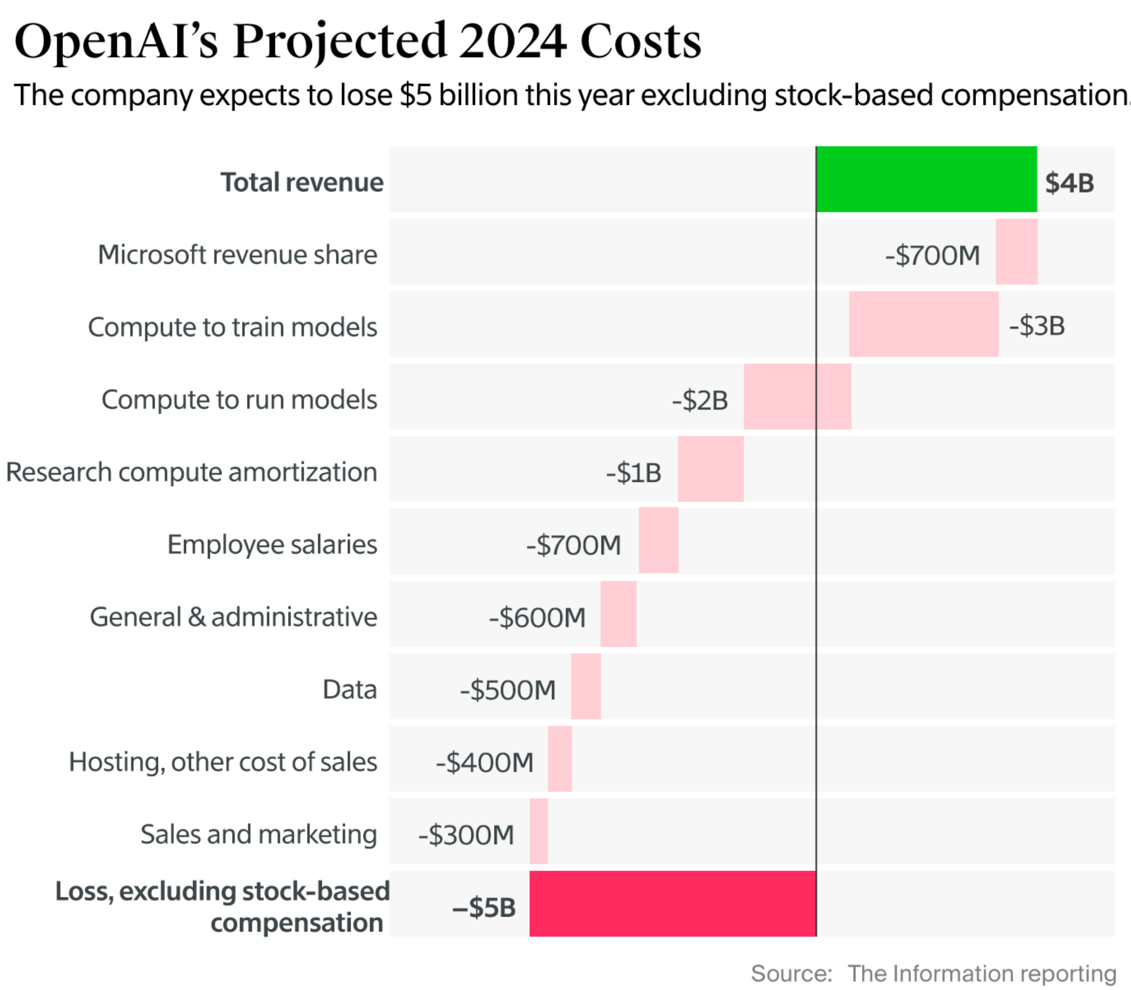

OpenAI's biggest cost-centers are its compute — for both training and inference — and its staff, who are paid on average $1.37 million a year (seriously), with an estimated 6400 people working at OpenAI.

Assuming the lowest end of the estimated salary — $990,000 — OpenAI is paying $6.36 billion in salaries **alone**, and if we assume the average is accurate, that number creeps up to over **$8.7 billion just for paying people.** Even if we play silly buggers by saying it's half, that's still $3 billion to $4 billion *just for people, at a time when AI talent is getting way, way more expensive.*

On its own, this is a *terrifying number.* OpenAI is spending incredible amounts of money on thousands of highly-specialized staff, and any layoffs would immediately send the industry into a spiral. OpenAI's massive amounts of capital have allowed it to inflate and distort the market for AI talent, and its investment in talent now means that it has to continually pay obscene rates for thousands of people, and recently gave out bonuses of hundreds of thousands to millions of dollars to around a thousand employees. On the low end, this cost it $200 million, otherwise known as "10% of OpenAI's projected API revenue for 2025."

This article also says that 1000 people are "about a third of the company." If that's the case, OpenAI is paying around **$3 billion in salaries (at $990,000 a salary),** also known as "$1 billion more than OpenAI's projected API revenue for 2025."

In any case, this is an obscene amount of money!

Compute Costs Are Likely Astronomical, Burning

At Least $15 Billion — If Not $20 Billion — In 2025 Alone

So, OpenAI is poised to lose $8 billion in 2025, off of $12.7 billion of revenue. That suggests total expenditures of $20.7 billion.

If we assume that salaries are $3 billion and the various miscellaneous costs (data, administrative fees, legal fees, sales, marketing, etc) are another $3bn, we're left with $14.7bn. If we assume that salaries are in the $6bn range (which they very well may be), we're left with $11.7bn.

So, we're looking at either $11.7 billion or $14.7 billion in compute costs. In February 2025, The Information reported that OpenAI projected to spend $13 billion on Microsoft Azure this year, before any changes to product launches around GPT-5 (or GPT-4.5) were going to happen. OpenAI has also signed deals with both Google and CoreWeave for compute, with CNBC reporting that Google's ChatGPT infrastructure will run in the US, Japan, the Netherlands, Norway and the United Kingdom.

There is not a chance in ***hell*** OpenAI spends less than $15 billion on compute this year, and I'd wager it's more in the region of $20 billion.

Or, put it another way, its $8bn loss is likely going to be much, much higher than expected.

Since 2024, OpenAI has launched multiple token-hungry reasoning models, including the ultra-expensive o3 model that apparently cost as much as $30,000 per task. On top of that, GPT-5, according to the Wall Street Journal, took multiple $500 million training runs alone, and ended up being released as two separate models — GPT-4.5 and GPT-5, with Altman describing 4.5 as a "big, expensive model."

Additionally, OpenAI launched GPT-4.1 in April, a model focused on coding — another expensive, token-heavy Large Language Model use case.

On top of that, OpenAI's user base has supposedly exploded. In August 2024, OpenAI said that ChatGPT had more than 200 million weekly active users, and the New York Times would report in late September 2024 that ChatGPT had hit 350 million monthly active users.

> **Sidenote**: Note the difference between "monthly active users" and "weekly active users."

OpenAI would announce in February 2025 it hit 400 million weekly active users, and by August they'd announce it had hit 700 million weekly active users.

Sam Altman would even say at one point that GPUs were melting due to the popularity of OpenAI's new image generator.

So, OpenAI has hundreds of millions of new users, using reasoning models (and yes, OpenAI brought its reasoning model o3-mini to ChatGPT in January 2025, and sources tell me these models are still quite compute intensive!) and generating compute-intensive images. The popularity of these products was so significant that OpenAI had to buy *extra capacity from Google to keep up*, *starting in May of this year*, and Google will also be (according to Reuters) providing compute for training AND inference.

As a result, it's very likely OpenAI's compute costs are over $15 billion, and I wouldn't be surprised if they're even worse, considering how sources have told me OpenAI loves to have separate teams working on the same thing, burning compute with nary a care in the world. Its losses are *not* $8 billion — I expect them to be at least $10 billion, if not upwards of $12 billion, especially if their salaries run in the higher range of my estimations.

OpenAI Is Bleeding Out, And Could Run Out Of Money By End Of Year

There is no scenario where funding OpenAI further makes sense, other than one where venture capitalists lie to themselves about a company that will never, ever stop burning money.

OpenAI is not a business — it's a money-incinerator. In a little over a year, it's raised $26 billion (and a $4 billion credit line), yet still needs another $20 billion by the end of the year? Then it claims it'll *only* need another $17 billion in 2027? That's ludicrous! It has to start paying Oracle $30 billion a year in 2028, the very same year The Information reports it'll be paying Microsoft $28 billion for Azure compute. How does OpenAI possibly afford all of this, even with the capital it's raised so far?

So, devil's advocate! Let's say OpenAI has been **really good** with money, and in fact has been saving fastidiously this whole time.

By the end of the year it's made a total of $16.7 billion — that's $4 billion in 2024 and $12.7 billion in 2025 — but it cost OpenAI $29.7 billion to do it, and it's raised $50.7 billion, from the beginning of 2024 to the end of December 2025, when it'll receive $20 billion from SoftBank if it converts from a non-profit to a for-profit.

> Note: the FT is saying "$10 billion" - but this doesn't make sense to me at all. It's a $40 billion - $10 billion initially with $7.5bn from SoftBank and $2.5 billion from other investors, then another $30 billion, with $8.3 billion already raised from other investors, leaving $20.7 billion for SoftBank.

> In the event that SoftBank "cuts the round in half," that means that the final amount would be $20 billion total for **the entire round** - meaning that the remaining balance would be $1.7 billion, or potentially $700 million - I swear to god I saw Sarah Friar, OpenAI's CFO, say they raised $11 billion in the June 2025 round. I'll add the link if I can find it.
>
> I'm not sure what I am missing here. Is the suggestion that this round was actually $20 billion SoftBank and $20 billion others? That doesn't reflect SoftBank's explanation of the deal.

This suggests that, even in the rosiest, press-leaked projections, OpenAI will run out of money by the end of the year outside of raising more capital.

The $20 billion lifeline from SoftBank would, potentially, relieve its burden — but even then, it's very clear that OpenAI **cannot survive without raising at least another $10 billion by the end of the year, and will need to raise at least $20 billion in 2026, if not more.**

And if it doesn't get the full check from SoftBank, OpenAI could very well die.

Chaos Bet: Microsoft Kills OpenAI By Blocking Its Non-Profit Conversion

As I wrote a few months ago, I think that Microsoft may *want* OpenAI to collapse. The Information reported in June that, when funding OpenAI's 2023 round, top Microsoft executives believed OpenAI would fail, and the nature of their agreement — which gives Microsoft full rights to sell OpenAI's models and access to their intellectual property and research — means that OpenAI's failure would allow it to effectively take all of OpenAI's stuff.

Though nothing is certain, there is a scenario where — [as reported back in June] — Microsoft simply walks away from negotiations. I imagine if this seems likely that OpenAI will continue its belligerent shadow PR campaign [where it continues to accuse Microsoft of things like antitrust violations], and if Altman begins taking a public stance on the matter, *it means that the negotiations have fully broken down and it's time for the fun to begin.*

While both parties are still negotiating, it's *been a while since we heard an update.* OpenAI wants Microsoft to reduce its cut of OpenAI's revenue, give up IP rights, and a bunch of other crap that Microsoft doesn't have to agree to.

OpenAI's logic, I imagine, is that Microsoft would benefit from an OpenAI IPO.

Remember, Microsoft has a lot of insight into the state of OpenAI, both in terms of its intellectual property and its physical infrastructure (which Microsoft runs almost the entirety of). Microsoft likely knows how busy its servers are, and can probably extrapolate revenues as a result, if it doesn't already know them.

Surely if things were going amazingly Microsoft would work with OpenAI to make this deal happen? Wouldn't it want 33% of the world's biggest, most hugest, most important company ever?

I get the sense that this is not really a "negotiation" so much as it's Microsoft saying it's happy with the terms of the current deal, knowing that its say-so is the only way to get OpenAI to resolve its non-profit to for-profit conversion by [December 2025]. If OpenAI misses that deadline, the round size drops to $20 billion, and considering OpenAI has already raised $19.3 billion (including

SoftBank's initial $7.5 billion), that means SoftBank owes it a whole $700 million.

And now the Financial Times has reported that negotiations with Microsoft could drag on into 2026. How, exactly, does any of this work out?

Even If Microsoft Agrees, OpenAI Does Not Have Enough Time To Convert To A For-Profit By The End Of The Year

People are so fuckin' cavalier about this! Even if Microsoft gave its blessing *tomorrow*, OpenAI would have three or four months to pull off something that has literally never happened before, and have to get past both the Delaware and California District Attorneys. This is a complex, multi-faceted legal challenge with little or no precedent, and while I'm no legal or non-profit expert, **I see no way it gets this done by December 2025.**

SoftBank has also not been entirely clear on whether the current remedies to the for-profit situation actually work. As I wrote up in July, SoftBank isn't exactly clear (thanks in part to weasel wording and translation from Japanese to English) as to whether OpenAI's current reformation will actually fulfil its requirements:

> Recently, OpenAI updated their blog. And also the media has been covered, and there was some misunderstanding. See, like you said, that the media is saying giving up on becoming the profit organization. I don't think that's the correct understanding. The shareholder of the – or the holding company is the nonprofit.
> And underneath the company which does the business, but more specifically, nothing has been changed at this moment. They just

> show the direction of the structure. That's what's been announced. So at this moment, nothing has been really happened or changed, just to discuss the direction. And for us, I don't think that that's a wrong direction, a bad direction.That's something that we've been expected. And our investment is in the business vehicle, and this is – hopefully that they'll be able to grow. Then that I believe that's a good thing for our investments.

When asked about whether “the event” (referring to OpenAI not converting and thus SoftBank cutting its round in half), SoftBank's CFO said the following:

> No. Event is not happening now. So if such event takes place by the end of the year, then reduce from $40 billion to $20 billion. We never know when such event takes place.

That doesn't sound like a ringing endorsement!

In any case, it isn't clear that OpenAI is capable of converting to a for-profit at all, let alone by the end of the year. While I imagine SoftBank **could** show some leniency, there's also the other problem that I've gone over in detail in this newsletter — that SoftBank can't currently afford the remaining $20 billion without some terrible financial consequences.

Chaos Pick: OpenAI Does Not Convert, And Does Not Receive More Money From SoftBank

In this scenario, OpenAI is unable to convert to a for-profit by the end of the year, and SoftBank does not provide the money it promised.

At this point, OpenAI will die. Not immediately, but they will, by my calculations, begin running out of cash by around April 2026. It will raise another round — likely another $10 billion, draining yet more of the dwindling dry powder, likely opening up to as many investors as want in — and keep spinning its wheels suggesting that it is "in active negotiations" with SoftBank, or potentially go as far as suing it, as WeWork once sued SoftBank when it pulled out of a share sale.

OpenAI will, I believe, keep lying about its prospects of becoming a public company, and use these lies to continue harvesting as much venture capital as it can until the inevitable happens — people realize that there's no way a non-profit of this size and complexity can go public, that there will never be another liquidity event, and that there's no reason to fund this company further.

At this point, the AI bubble will already be bursting. **OpenAI failing to convert and not receiving the money from SoftBank will begin an industry-wide panic, and unilaterally state that the party is over.**

I don't know how "likely" this is, but it sure feels possible.

Alternate Chaos Pick: OpenAI Does Convert, But SoftBank Can't Get The Money

This scenario is looking more likely than ever, due to the fact that it's unclear how SoftBank actually sends that $20 billion. I still can't find any evidence that SoftBank is making any moves to raise it. I still can't find a way it can do it outside of maxing out its leverage. Remember, it took 21 different banks to fund SoftBank's $15 billion bridge loan, $7.5 billion of which was used to fund OpenAI.

There is a decent chance that SoftBank slow-walks this investment regardless of what happens. The $10 billion — of which SoftBank owed $7.5 billion — was due in *April*, but it took until *June* to get there.

So there's a scenario where everything goes well — where Microsoft agrees, where OpenAI can somehow convert in time, and where SoftBank doesn't have the money. It's unclear how much is actually left in Vision Fund 2. A report from February 2024 said that there was around $6.4 billion left

I don't believe there is, otherwise it would have used that money to fund the first $7.5 billion.

At this point, **OpenAI will begin raising more money from other parties, and do so in a desperate, aggressive way. Even if it succeeds in doing so, venture capital is going to run out of money to fund them, and SoftBank is too.**

How Does OpenAI (or Anthropic) IPO?

Assuming all of these insanely complex and expensive things happen, and assuming that these companies don't simply run out of money before they do so, *the goal for OpenAI and Anthropic is to go public.*

I do not think this is possible, but if they try, I believe both will receive a similar reception to WeWork's S-1 filing, which CNBC journalist Annie Palmer described parts of as "strange and alarming," leading to WeWork pulling its plans of going public a month later.

Both Anthropic and OpenAI burn billions of dollars and neither appear to have a single answer as to why. I also believe that both of them, using the press as their agents, share inflated or obfuscated revenue and cost projections as a means of keeping the hype going.

This is **entirely** gut instinct, but I believe that neither of these companies would survive a thorough audit, in part because their costs are so chaotic and variable. Another issue for any IPO would be that Microsoft effectively owns OpenAI, and takes a 20% cut (for now?) of OpenAI's revenue, and, oh, I dunno, *the fact that it has absolutely no path to profit?*

Assuming a conversion, I believe both companies *could* file to go public. But if CoreWeave's horrifying S-1 tells us anything, OpenAI and Anthropic's books are going to look like one of Dexter's victims.

OpenAI and Anthropic both have a kind of halo with the media where reporters will take great pains to find the best possible read on their stories. CFO Sarah Friar can go on CNBC — a leading business news channel — and waffle on about nothing without fear of troubling questions like "hey, as the CFO of OpenAI, can you explain how you plan to stop burning billions of dollars?"

This is all cute and fun for now, and has allowed both companies to raise billions of dollars of capital at the most ridiculous rates, all without any real scrutiny of their business practices or business models. They've been given permission to destroy our environment, steal from millions of people, and let their models drive people to psychosis, with the media mostly responding with "eh, that'll happen."

As a result, there's very little pushback against revenue leaks that, when laid out together, do not make much sense, or at the very least suggest both companies' costs are astronomically higher. Personally — and no, I do not let this bias change how I report these stories, nor will I make claims that I cannot confirm — I do not trust the numbers they are leaking, nor do I think they are being truthful about the condition of their businesses.

What Happens Next?

This has now been the longest thing I've ever written, and the most I've ever written in the space of a week, but I wanted to try and estimate what might happen next as quickly as possible.

NVIDIA's earnings will be groundhog day for AI. If Jensen doesn't see his shadow, earnings will beat expectations and growth will continue, and the markets will get a little hopium for the next three months.

But I don't think, even with successful earnings, that this narrative goes away. Big tech is not showing any return on investment for AI, and OpenAI blew the launch of GPT-5 so hard that it's unclear where this company goes from here. Both OpenAI and Anthropic are dependent on endless resources at a time when the markets are *hinting* that they're going to need to see those resources actually do something.

I will also be honest: I have no idea how OpenAI converts, as I have no idea how SoftBank gives it the money if it does so. I also have no idea how OpenAI or Anthropic go public, and *they must go public, because they are too big to acquire*, and going public will require them to show all of their dirty, ugly financial statements.

Even if OpenAI raises its $20 billion, it needs to raise another $10 billion or more in 2026, and Anthropic is in a similar position. How much more venture capital is left to prop up these capital-engorged failsons? How long will Silicon Valley tolerate two giant losers that love to lose money?This entire narrative — this entire *industry* — is based on vibes, and there are enough signs that something weird and dodgy is going on that I believe *someone* is misstating how much they make, how much they lose, or a little of both.

So Why Did You Say 2027?

In reality, I think this nonsense can continue another year, even if OpenAI doesn't convert. Clammy Sam Altman has always been good at raising money, and I think he can beguile the media and venture capitalists with his siren song of "innovation." Each OpenAI and Anthropic round is heavily oversubscribed, meaning there's plenty of dumb money left to feed. I'd wager another $30 billion, max.

That isn't remotely enough for the long-term, but it's enough to let them limp to a slow death in February 2027.

That being said, there are enough events — a failure to convert, a failure to get the money from SoftBank, or really any major market vibe shift — that could turn this entire situation inside-out.

I also will add that as things get desperate, OpenAI and Anthropic can and will increase prices (as they already have begun to do with their enterprise customers) to try and make their books look better, especially if they make it anywhere near an IPO. This will begin the Subprime AI Crisis I wrote up last year, where the cost of "doing AI" becomes untenable, draining what remains of the hundreds of millions shoved into the bodies of overvalued AI startups.

As The Information wrote last week, "AI model buyers are feeling squeezed while companies that sell them are soaring," a dark omen of what's to come.

Things could blow up in two months, or the middle of next year, or the beginning of 2027. But I see no way that the generative AI boom ends well — and in fact believe we're in for a hell of a collapse once the bubble bursts.