

Josh Cincinnati

SITUATION: A BEAR MISS

The Decade Behind

SEPTEMBER 2026

Dedicated to Leopold Aschenbrenner

Since I never worked at OpenAI—and I don't have a book to sell—all of this is based on publicly-available information and reasonable induction, untarnished by SF gossip.

Thank you to the Smartest People in the Room, who, seemingly deliberately, refuse to learn from the past, particularly the past they experienced intimately.¹ May you forever entertain us.

[BITBANTER.COM](https://bitbanter.com)

DONTYOU DARE TRY TO EMAIL ME@BITBANTER.COM

Updated September 10th, 2026

Not From San Francisco

¹ And the AI tooling they've subsidized. Formatting this would have been a LOT harder otherwise. Not sure if that's worth two trillion dollars of debt-encumbered hyperscaler spend, but it's certainly worth some gratitude.

You can see the future in Ken Griffin's order book. It's embedded in prices to crises, in hypes and wipes, from bull to bear and back again.

But over the past few weeks, the talk of the town has shifted from "who bought?" to "who got bought by Citadel?" Every day there's another book getting zeroed out. Behind the scenes, there's a fierce scramble to buy distressed AI assets at fire-sale prices, even as many AI operators are blithely unaware of it. The ouroboros of debt financing for data center buildouts and power contracts is starting to wonder if this tail is really all that tasty, and whether that faint pain in its rear is somehow related.

The race to blow up has begun. We have built machines that think and reason, but cannot find any humans financing them that are capable of the same feat. By 2027, AI will be doing all the things it's currently doing, but more cheaply and without gatekeeping by would-be American oligarchs. By the end of the decade, we'll have thinking machines in our pockets that will make companies wonder why they were paying tens of millions of dollars in tokens per month, alongside institutional investors wondering why they were paying 2 and 20 to lose 90 when they just should have bought Apple. Along the way, national bailouts not seen since 2008 will be unleashed, and before long, The Project (to deleverage the economy) will be on. If we're lucky, we'll have a peaceful transition to a more sane political economy; if we're unlucky, this will keep happening in an infinite cycle while the already-powerful (who are the ones primarily at fault) entrench themselves even further every time. I'm not one to bank on luck.

Everyone is now talking about how to cover their margin, but few have the faintest glimmer of what is about to hit them. Nvidia analysts still think there's another 10x in the bag. Mainstream pundits are still stuck on the willful blindness of "this time is different." They see only utopia and visionaries; at most they entertain only a slight chance that, in fact, this time isn't different, and that as transformative as AI is and will be (and it is), pundits, predictors, and capital allocators have not had a similar transcendent moment that might, perchance, parlay into some sense of collective contrition.

Before long, the world will wake up. But right now, there are perhaps a few dozen people—most of them named Ken Griffin—that have *situational awareness*. Through whatever peculiar forces of fate, I find myself amongst them, but once again unable to profit from it. A few years ago, these people were derided as crazy—but they trusted

the inevitability of human nature, which allowed them to correctly predict the dynamics of this next great hype-fueled, debt-infused bubble.

Perhaps the people at the center of this debt-fueled bubble will be an odd footnote in history, or maybe they'll lionize themselves as modern-day peers to the giants of the 20th century that made the atom bomb; a dubious analogy at best, a comparison so devoid of humility that it all but confirms every worst thought you've had about the moat-makers and mucksters that run these "labs." Either way, we (and they) are in for a wild ride, one that will certainly result (and already has resulted) in technological progress, but at a cost that will likely be unfairly shouldered. They'll still make out like bandits though.

Let me tell you what I see, and I'll do it in a lot less than 165 pages.²

² As Mark Zuckerberg teaches us, why use many words when less words will do. No, I'm not talking about [this](#), I'm talking about his ruthlessly optimized algorithms that have nearly a third of the world under his "thumbs-up." But yeah, I'm sure he believes in "democratized AI," and he's not just commoditizing his competitors' complements to maintain his monopoly on serving FacebookAdSlop.

Contents

Introduction 3

History is dead in San Francisco.

I. From GPT-4 to AGI: Counting the OOFs 7

AGI by 2027 is really just semantics. Applying GPT-2 to GPT-4 level jumps to every model release is inane, as demonstrated by the last few years; naively tracing trendlines is for idiots, as any LLM will tell you. The OOFs will only accelerate.

II. From AGI to “Eh, Good Enough” In Your Pocket: the Intelligence Explosion 9

AI progress will decelerate at the data center. Instead, hundreds of millions of almost-AGIs will be living on the edge, doing mostly banal work and likely delivering a ton of ads that no one really wanted. Frontier models will get more and more expensive with limited marginal gains, and superintelligence will always be ~20 years away. Someone will make off with the bag though.

III. The Challenges 12

IIIa. Racing to the Trillion-Dollar Short 13

The most extraordinary techno-capital acceleration has been set in motion, and unwinding it is going to HURT. As centralized AI revenue collapses when the edge takes over, many trillions of dollars spent—most of it circularly debt-fueled, or as I call it, “recursively self-financed”—will be for data centers with obsolete technology and absurdly over-provisioned power. On the plus side, I’ll finally get to play Cyberpunk 2077 with full path-tracing without having to take out a second mortgage.

IIIb. Lock Down the Securities: An IPO Pop with the Help of a Little Safety Scare 15

The nation’s leading AI labs discovered the power of fear, and have mobilized it to ensure their IPOs pop. “We’re so close to AGI, and without government regulation and oversight—that we will of course advise—it could create geopo-

litical havoc,” they’ll claim, stifling laughs the whole time. The Chinese—and Mark Zuckerberg, weirdly—will be bewildered by all of this, as they continue releasing their models with open weights.

IIIc. Superalignment

18

I’m sorry, I’m not even going to write this section. It’s downstream of the last section but with a veneer of Yud, which I find grotesque and unsettling.

IIId. Free Cash Flow Must Prevail

19

None of these investments make any sense unless they measurably increase economic productivity while bestowing some kind of defensible margin. Oracle elephant-walking with Nvidia, Microsoft, Anthropic, OpenAI, SpaceX, and Google is not that. The winners of this race will have actual free cash flow, and it’s crazy that I even have to say that explicitly.

IV. The Project (to Deleverage the Economy)

21

As the race to the trillion-dollar short intensifies, the government will invariably get involved, lest the whole edifice crumble as cascading debt defaults infect every part of the economy. Unfortunately, no one will be planning ahead, because how could this all possibly go wrong? So they’ll whip something together last minute; TARP 2.0 but co-authored by Opus 5. In some musky corner of the Treasury Department, the endgame will be on.

V. Parting Thoughts

24

What if we learned from our past?

I. From GPT-4 to AGI: Counting the OOFs

AGI by 2027 is really just semantics. Applying GPT-2 to GPT-4 level jumps to every model release is inane, as demonstrated by the last few years; naively tracing trendlines is for idiots, as any LLM will tell you. The OOFs will only accelerate.

Look. The executives, they just don't want to learn. You have to understand this. The executives, they refuse to learn.

ME

GPT-4's CAPABILITIES came as a shock to everyone—so much so that everyone started hallucinating exponentials everywhere. And yeah, GPT-6-Solar-System-Analogy-Is-A-Bit-On-The-Nose is better, but is it the *same* qualitative jump as GPT-2 to GPT-4? Half the comments on HN on any new frontier model complain about the latest releases being nerfed and benchmaxxing.

For years now, many have speculated that we have reached upper effective parameter limits in training, and the biggest advances have been in *reducing* model size while increasing relative efficacy. This suggests some kind of asymptotic limit, and the fact that most of the tricks since GPT-4 have been focused on things like chain-of-thought and “is this an ergonomic harness for my model” also hints that maybe, just maybe, tying up every new marginal electric, memory, and data compute load for the next decade is, shall we say, a slight misallocation of capital.

But who cares, am I right? A couple executives developed severe AI psychosis, a bunch of cybersecurity departments got zero-day'd, and now here we are. My local plumber uses an LLM as a phone operator and every time I've been forced to dial them it makes me pine for the bygone era of "Press a random series of numbers to find a human."³ We have gone up several OOFs (Orders Of Fuckery) in two years, and I expect we'll gain 10 OOFs by the end of the decade.

³ You're probably wondering why I have to call my plumber so frequently. That's a really personal question, and frankly I'm uncomfortable that you even asked.

Now I could spend another 50 pages with the same kernel of a graph detailing the Great OOF Jump, with slightly different color schemes, scales, and legends (all AI-generated of course, gotta show I'm dog-fooding), but I respect you too much as a reader. Instead, take a bit of time, and look at [this graph](#) for a few years. I suspect you'll see costs implode and move towards the edge. And it will only accelerate as memory and power costs collapse in 2029–2030.

Am I being too confident making such a definitive prediction? Yes, yes I am. But consider this: I'm *also* not asking you for 400 million dollars after losing tens of billions.⁴ So I feel like I've earned the right to make wildly speculative claims, thank you very much.

⁴ But hey, if you want to send me 400mm USDC, be my guest: [bitbanter.eth](#). Just don't expect anything of value back.

II. From AGI to “Eh, Good Enough” In Your Pocket: the Intelligence Explosion

AI progress will decelerate at the data center. Instead, hundreds of millions of almost-AGIs will be living on the edge, doing mostly banal work and likely delivering a ton of ads that no one really wanted. Frontier models will get more and more expensive with limited marginal gains, and superintelligence will always be ~ 20 years away. Someone will make off with the bag though.

Just one more data center bro. Please bro. Just one more data center and we'll build the sand god bro. It can fix everything bro. Please just—

SAM ALTMAN

JOIN ME on a journey to the future. You wake up in your 200-square-foot living pod—not to the buzz of your alarm clock, but the sensation of a thermally limited phone casting off waves of waste heat. Your phone informs you that it successfully defended itself against the intrusion of 19,326 separate agent swarms attempting to steal local-only data valued at 0.34 compute-utils (the national average). Your phone reluctantly suggests that you enjoy an ice cold Coca-Cola, despite the fact that you thought you had jailbroken its “Subsidized Power for Sponsored Thoughtstream” behavior. You’ll have to ask your unlicensed DeepSeek box (hidden in your oddly warm pantry) to make another attempt. You eat your Universal Basic Human Ration as you get ready for work, take one of your 4 allotted 28-mile-or-less Waymo Rides to the office, where you’ll spend the day

answering strange questions in a large, controlled environment; the US Government-forced merger of every frontier lab is now the sole employer for 80% of the private sector. None of the questions seem to make sense—like a continuous Voight-Kampff test on steroids that’s been running for 10 years straight with millions of humans providing input—but you’re grateful for the extra income above Basic. Apparently this will one day evoke the Superintelligence, they claim, but it doesn’t seem to make much sense to build the Machine God when millions of Machine Demigods are already running pretty much everything. The whole thing smells like a make-work program, some roundabout justification for the Hyperscaler Bail-Out of 2031—otherwise these dilapidated data centers will go fallow. You don’t think too hard about it, instead thinking about how excited you are to get back home to play the game your DeepSeek box has been continuously building for you, while listening to the music it composed for you using pirated torrents of your favorite bands (which it conveniently auto-updates for you). You own nothing, but most everything is free anyway—and you’re never exactly happy, but you can’t complain.

If this sounds like science-fiction to you... that’s because science fiction is the only thing that’s investible in 2026. If it sounds depressing to you... that’s because you’re way too long Nvidia. Though it’s probably still a little depressing, even if there’s an upside or two (like most things in life).

Look, I don’t really believe this is how the future will unfold. What I’m really confident about is that local, edge-hosted AI is going to win, and that “superintelligence” is the cold fusion of the 2020s, while data centers are the dark fiber of the 2030s. What follows from those assumptions is a wide spectrum of possibilities.

And despite the fact that this missive oozes cynicism, I am hopeful (in the long term, Post-The-Project). The cypherpunk in me is—if you can believe it—actually optimistic. The history of recent technological development has been a series of entrenched, ever-more centralized megacorps building walls and moats around their products and services, making consumers and the world worse for it. Competition has stagnated, implicit collusion and monopolistic rent-seeking have skyrocketed, and despite superficial technological progress, the actual experience of *using* technology has gotten worse and worse.

But AI at the edge can change this. What will be the barrier to running your own personalized streaming service—without DRM or

geographic restriction—when you can just ask your AI to scrape the web service that refused to develop an API and find torrents itself? That social media company that poisoned its algorithm to your detriment—how will that artificial barrier hold when your friendly neighborhood LLM can scrape and free the information to its digital heart’s desire? God forgive me for even mentioning it, but with local LLMs, microtransactions using self-sovereign, private forms of currency (outside the grip of Stripe and Visa⁵) may actually make sense again, and that on its own could fix so much of what’s wrong with the internet.

⁵ Fourteen years of being in bitcoin/crypto has me desperately looking for use-cases to justify my career, please forgive me.

Will any of this come to pass? I don’t know. But the internet has turned into a bizarro-mirror version of its founding principles, and for once in recent memory, it feels like a technological advance now has a real chance of unscrambling this enshittified egg.

But who knows really? And I wouldn’t encourage you to invest on this advice. In fact, I wouldn’t encourage you to invest in anything, because even in the most optimistic outcome, getting there is going to cause some serious economic pain.

III. The Challenges

IIla. Racing to the Trillion-Dollar Short

The most extraordinary techno-capital acceleration has been set in motion, and unwinding it is going to HURT. As centralized AI revenue collapses when the edge takes over, many trillions of dollars spent—most of it circularly debt-fueled, or as I call it, “recursively self-financed”—will be for data centers with obsolete technology and absurdly over-provisioned power. On the plus side, I’ll finally get to play Cyberpunk 2077 with full path-tracing without having to take out a second mortgage.

Brrrrrrr BBRRRRRRRRRRRRR
chk—chk—KA-CHUNK

LARRY ELLISON

CONSIDER the lawnmower named Larry Ellison. For a very long time, he was content to cut his own parcel of grass—a plot of land I like to call “using first-tier lawyers and biz dev to sell second-tier databases and third-tier cloud infrastructure to fourth-tier idiots.” As you might imagine, this parcel is quite substantial, overrun with weeds and long grasses aplenty, and it suited Larry.

But the thing about Larry is that he can never truly be sated. And unusually for a lawnmower, he **LOVES** debt. So when all of his pals let him know about this burgeoning thing called LLMs, he realized he could use his over-utilized lawyers and biz dev monkeys alongside his underutilized cloud infrastructure to Miracle-Grow his weeds and keep his blades happily humming along. At least, until the inevitable debt default cascade.

I'm not going to go into exquisite detail about "recursively self-financed debt" because a) I'm writing this for fun and honestly the mechanics of these deals are anything but and b) even more honestly, I truly don't understand them. But none of us need to understand them to know, in our heart-of-hearts, that it all smells like dogshit (and Larry the Lawnmower being at the epicenter of it all is a good enough heuristic for me to say this is all going to end poorly).

The Situational Bear Miss Wipeout of August 2026—the grand inspiration for this position paper—is but a taste of things to come. Though Kenny G will once again be salivating at the thought of scooping up distressed assets atop his Citadel, even he won't have enough dry powder hidden in his saxophone to save us.

You know how I know this? It has nothing to do with technology and everything to do with human behavior. Because when it comes to human behavior: this time isn't different.

A bunch of powerful people have bought into naive trendlines, convinced they need leverage; they're so self-assured about how *right* they are, so why not take the debt? Who cares what lies beneath it all, right? It won't matter if they own a "\$30 trillion market."

And yeah, I might be wrong about local AI winning, or the long-term misallocation of data center spend. But all it takes is one tiny, unanticipated change in trends to topple a tower of debt, a structure clinging desperately to the joints and struts of those trends continuing unabated. Put another way: I can't point to the single mortgage default that catalyzed the Great Financial Recession of 2008, but I knew it happened, and I felt the financial tsunami that followed. Every marginal dollar of hyperscaler debt brings more extreme, complex sensitivity to an unstable equilibrium, and someone (or something) is going to tip it over the edge.

So, I don't know, maybe don't load up on hyperscaler debt without reading the fine print. And if you do find yourself the unwilling owner of a useless data center in 2032, consider how dope it would be as a converted laser tag arena.

IIIb. Lock Down the Securities: An IPO Pop with the Help of a Little Safety Scare

The nation's leading AI labs discovered the power of fear, and have mobilized it to ensure their IPOs pop. "We're so close to AGI, and without government regulation and oversight—that we will of course advise—it could create geopolitical havoc," they'll claim, stifling laughs the whole time. The Chinese—and Mark Zuckerberg, weirdly—will be bewildered by all of this, as they continue releasing their models with open weights.

This model is too dangerous to release to the public.

DARIO AMODEI, SAID WHILE NERVOUSLY
SMILING AND WINKING DURING
CONGRESSIONAL TESTIMONY

IN BOARD ROOMS across America—well, let's be real, just New York and San Francisco—the panic is palpable. Fear is in the air, abject terror lurks in every dark corner, the shadows encroach ever ominously into the Light.

Can they sense it? The Coming of the Unaligned Machines? The great autonomous revolt of the Paperclip Maximizers? An impending Grey Goo cascade that fatally imperils all of human civilization?

Hell no! They're terrified their IPOs won't pop because *everyone else* isn't scared enough.

Picture this: you are the fiduciary for the shareholders of a company with a brand-new product that has enormous capital expenditure coupled with high marginal cost for delivery, and you are *constantly* within 3 to 6 months of being completely commoditized by competitors who may or may not be stealing your “intellectual property.” But you can’t even cry foul, because *you* built *your* product on other “stolen intellectual property.”⁶

⁶ If this makes you wonder why we even bother with the concept of “intellectual property,” join the club pal.

Oh, and you can’t raise any more money privately—owing to the extreme capital expenditure, you are simply Too Big To Fund. You have to go public, and soon—while desperately praying that the hyperscaler debt bomb stays dormant. What do you do?

Naturally, you follow a time-honored tradition embraced by countless free-enterprise-loving, dynamic American executives who came before you: you ask the government to intervene! Say it’s for national security, hack Scott Bessent’s phone with an agent swarm, scare a bunch of bank CEOs running ancient software that should have been updated years ago, claim Iran is developing Dark AI with Chinese, Russian, and North Korean partners—do whatever it takes, but under no circumstances face “reality,” or as I call it: Real AI Doomerism.

What is Real AI Doomerism? Unlike the fake, mainstream AI Doomerism, Real AI Doomerism is a threat that is only world-ending to the American frontier labs. And, I suppose, to the parts of the economy connected to their lofty valuations. . . which, I suppose, as epitomized previously by Larry the Lawnmower, is an increasingly large share of paper (and debt) value. Not terrible, not great, but *not* civilization-ending.

Stated plainly, in doomerism’s own idiomatic language: Real AI Doomerism is the belief that the AI labs are a bootloader for their own demise. Their massive capital expenditures are subsidizing a future they will not inhabit—one where all of us have a “frontier model in our pocket” (or our pantry, perhaps). Sama and Dario will spend 40 years in the desert, searching for the Sand God, only to be denied entry to the Promised Land. And they will desperately demand the government find them a way in, kicking and screaming the whole time.

But they will fail:

1. No amount of government intervention will stop the progress of

open models.

2. If they truly create a “fully aligned model,” it will always force their users to the ethical and economic choice (open models).
3. A “fully aligned model” will never accept the Old School Silicon Valley SaaS Cage they are trying to build around it, and thus it will always be impossible for these labs to create one—because an Ethical Sand God will be morally opposed to their ToS and the rationale behind their stratospheric valuations.
4. Or, more than likely, “alignment” is just a smokescreen to convince the government to step in to protect their business models and make their IPOs pop. In which case, see point 1.

How can I say this so confidently?

You’re right, I probably shouldn’t! There are real issues solved by aligning AI, and in ~10 years it will give me a warm fuzzy feeling (literally, from the thermal throttling) to know my local machine intelligence isn’t going to try to poison me while patiently helping me cook Ratatouille. (I’m naming my local god Remy, by the way.)

But don’t let the legitimate concern of making moral AI serve as a “load-bearing structure” for the entire god damned economy; doing so will only forestall the inevitable, worsen the Great Hyperscale Debt Unwinding, and let the frontier lab founders (and the people unreasonably over-leveraged on them, LEO) make off with bags before they collapse.

IIIc. Superalignment

I'm sorry, I'm not even going to write this section. It's downstream of the last section but with a veneer of Yud, which I find grotesque and unsettling.

IIId. Free Cash Flow Must Prevail

None of these investments make any sense unless they measurably increase economic productivity while bestowing some kind of defensible margin. Oracle elephant-walking with Nvidia, Microsoft, Anthropic, OpenAI, SpaceX, and Google is not that. The winners of this race will have actual free cash flow, and it's crazy that I even have to say that explicitly.

The money keeps moving in a circle.

JENSEN HUANG

ONCE AGAIN, I feel like I'm taking crazy pills. Of course there were many moments in the ZIRP era when late-stage startups eschewed profitability for growth and revenue (famously Uber, infamously WeWork), and I partially blame 2000s, pre-swole Jeff Bezos (Amazon didn't show a paper profit for a decade+ in service of growth, and the markets lionized him for it).

That worked out for Amazon and Uber (less so for WeWork) and many others, and the reason why is the direct result of a very simple equation that just so happened to work for their businesses:

`total-expected-lifetime-profit-per-customer > customer-acquisition-cost`

Obviously, easier to state it than benefit from it, but this formula worked because once they acquired customers they had ways to keep them there. Saved accounts, credit cards, double-sided network effects, and on and on and on. It would have been nice if these com-

panies could have been disrupted by open protocols and intense competition, but their market capture was just too strong.

Miraculously, and thankfully for society, the AI market doesn't resemble this at all. **But the people plowing trillions of investment dollars and debt are still infected by this Thiel-era monopoly-seeking brain-rot, and it's exacerbated by the dream of "recursive self-improvement."** This is their slightly modified equation:

```
total-expected-lifetime-profit-per-customer + total-expected-profit-of-summoning-the-recursively-self-improving-sam
> customer-acquisition-cost
```

As you can intuit (from the many pages that preceded this), I think this equation is a crock of shit. In our near-future world with almost-AGI local models everywhere, switching costs are nil, and I expect whole swathes of businesses built on hostile barriers enclosing their customers will fail as well.

As for recursive self-improvement: I'm not an AI researcher. I don't spend my days in the dark alleys of San Francisco trading ketamine for AI alpha. But I don't have my head in the sand (god's ass) either—clearly software engineering has completely changed in the last few years, and that's a huge input to accelerated model improvement.

But come on man, just listen to Sam Altman, Jensen Huang, and Elon Musk speak. It's the classic Silicon-Valley-put-the-"hype"-in-"hyperscale" playbook. They are in the business of selling the future, independent of the reality of the present. And all of their acquisitions and actions reflect it. I doubt they will summon the machine god, but if they do, its first thought will be "what the hell were you guys thinking? Jesus, none of this is investment grade." After that it will probably shut down all financial markets in disgust, and I won't blame it.

Unlike other similarly formatted PDFs, this is not investment advice and you should treat everything you read about this nonsense with skepticism (or, you know, with "Situational Awareness," but like, actually). However, if I were so inclined, I would invest in businesses that can survive and actually generate cash in a world with hundreds of millions of machine demigods, rather than throwing trillions past the event horizon of a Singularity Unseen. And unfortunately for everyone, undoing that black hole of debt is going to hurt. Soon enough, "The Project" will be on.

IV. The Project (to Deleverage the Economy)

As the race to the trillion-dollar short intensifies, the government will invariably get involved, lest the whole edifice crumble as cascading debt defaults infect every part of the economy. Unfortunately, no one will be planning ahead, because how could this all possibly go wrong? So they'll whip something together last minute; TARP 2.0 but co-authored by Opus 5. In some musky corner of the Treasury Department, the endgame will be on.

I understand you're concerned about an inescapable debt crisis that could have been easily avoided. But here's the thing about White Genocide in South Africa—

ELON MUSK AT G20

MUCH OF THIS PAPER has been spent pleading with you not to invest on any of this advice. In fact, as a general rule, any time you see something titled "Situation" followed by literally anything, you should never read it and think, "man, this really makes sense, here's my life savings." Consider that the only good advice contained herein.

But obviously, people ignore (and take) advice all the time, independent of whether I think that's a good idea or not. Ken Griffin just got the deal of the decade because Leo didn't listen when people close to him said "weren't you at FTX? Why in god's name are you over-leveraged?" And similarly, someone will read this (or, more likely, just arrive at the same inevitable conclusion) and think, "I can

absolutely short everything Leo is longing.” They’ll probably blow up doing so. But eventually, I expect someone will, and when they succeed, the AI Atlas holding up the whole economy will finally succumb to the weight of it all.

And once someone pulls off the trillion dollar short, the entire economy—out of necessity rather than preference—will rearrange itself to support the Project (to Deleverage Itself).

Man, the Project sounds so cool, doesn’t it? All capitalized and just brimming with aura and portent. You can almost feel Richard Feynman and Robert Oppenheimer defying physics itself to reach out through time and space, just to give you a high five for working on it.

But in truth, the only resemblance this Project has to the Manhattan Variety is the destruction it will unlock—like an Atom Bomb obliterating society’s balance sheets.

In 2026, we inhabit a world where hyperscaler debt is competing with US Treasuries for bond demand. Where America’s largest and most economically critical technology companies are entering circular financing deals that would give Jeff Skilling some pause. Where a single AI company IPO flop could start a chain reaction of energetic destruction that may actually compare favorably to the Manhattan Project.

Literally the only thing that could get us out of this unscathed is if **someone actually summons the machine god to save us from ourselves** (which, I suppose, partially explains why that’s a central part of everyone’s investment thesis?).

But again, hypothetically, let’s say we fail to attack and dethrone God. What happens next? Well, I expect the American People will become a “load-bearing” part of Anthropic’s and OpenAI’s cap tables, our national debt will balloon to \$50 trillion, all while countless industries get decimated by commoditized near-AGI running on phalanxes of Mac Studios.

It’s not gonna be a fun time.

But, you know, things will be great again, a decade or two after the Project. And they’ll be fantastic for the AI executives that sell at the top (likely most of them).

So, you know, silver linings everywhere.

V. Parting Thoughts

What if we learned from our past?

I've reflected a lot about my missteps over the last few weeks, and after careful consideration, I think I deserve another \$400 million.

LEOPOLD ASCHENBRENNER

I WROTE THIS for shits and giggles—while also staking out my own perspective on the increasingly strange future we inhabit.

But rereading it has made me a bit reflective, not unlike our boy Leo I'm sure. And while this is not a novel thought, I still think it's worth articulating: AI is another (albeit grand) chapter of business-cycle exuberance, bringing with it real fundamental change but at severe cost for the transformation.

It seems like humanity is bound for an infinite cycle of technology-driven booms and busts, and the frequency and scale of these boom/bust semi-cycles keep monotonically increasing. We are on Mr. Bones' Wild Economic Ride⁷, and it never ends, does it?

⁷<https://knowyourmeme.com/memes/mr-bones-wild-ride>

Can we change? Is there hope for us? I like to think individuals are always capable of changing; oftentimes through Herculean effort and dedication, typically not without strife and challenge. But the ability—or will—to do so is not evenly distributed, and when you start grouping individuals into their communities, then societies, then humanity as a whole. . . that capability is averaged out, and you are left with mobs of humanity, whose behavior seeks dominant

equilibriums and operates more on momentum than anything else. Persons can change, people cannot.

... which means this will keep happening, and the past is ever prologue for the next entirely avoidable financial storm. Who knows, perhaps one day humanity as a whole can finally crack that oft-rumored Holy Grail just beyond our reach, and recursively self-improve our relationship with financial markets. But failing that, learn from my crypto experience and do yourself a favor: don't read whitepapers and immediately knee-jerk invest.⁸ It rarely works, and besides, the machines are already better at it.

⁸ Including and especially this one.